

Release Notes

Published
2023-07-13

Contrail Networking Release 21.4.L2

Table of Contents

[Introduction | 1](#)

[New and Changed Features | 1](#)

[Supported Platforms in Contrail Networking Release 21.4 | 4](#)

[Container Tags | 7](#)

[Juju Charms Versions | 8](#)

[Known Behavior | 12](#)

[Resolved Issues | 50](#)

[Deprecated Features | 50](#)

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 21.4.L2 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/contrail21/information-products/topic-collections/release-notes/index.html.

New and Changed Features

IN THIS SECTION

- [New and Changed Features in Contrail Networking Release 21.4.L2 | 2](#)
- [New and Changed Features in Contrail Networking Release 21.4.L1 | 2](#)
- [New and Changed Features in Contrail Networking Release 21.4 | 3](#)

The features listed in this section are new or changed as of Contrail Networking Release 21.4.L2. A brief description of each new feature is included.

New and Changed Features in Contrail Networking Release 21.4.L2

Fast Forward Upgrade Support for Environments Using Red Hat Openstack Release 13 and Contrail Networking Release 1912.L4

Starting in Contrail Networking Release 21.4.L2, you can perform a Fast Forward Upgrade (FFU) from an environment using Contrail Networking Release 1912.L4 in Red Hat Openstack Release 13 to an environment using Contrail Networking Release 21.4.L2 in Red Hat Openstack Release 16.2.4. The FFU provides users with an opportunity to upgrade between long-life versions of Red Hat OpenStack that are multiple Openstack versions apart from one another while also upgrading Contrail Networking using the same process.

For information on performing this FFU, see [Fast Forward Upgrade: Updating Contrail Networking 1912.L2 and Red Hat OpenStack 13 to Contrail Networking 21.4.L2 and Red Hat Openstack 16.2](#).

BGP as a Service (BGPaaS) Support for Remote Compute External Controllers

Starting in Contrail Networking Release 21.4.L2, you can configure BGP as a Service (BGPaaS) in remote compute clusters. You configure BGPaaS in a remote compute cluster in the same manner that you would configure outside of a remote compute cluster.

For information on configuring BGPaaS in Contrail Networking, see [BGP as a Service](#).

For additional information on remote compute in Contrail Networking, see [Remote Compute](#).

New and Changed Features in Contrail Networking Release 21.4.L1

Support to Apply Import Routing Policies to MP-BGP Routes

Starting in Contrail Networking Release 21.4.L1, you can use Contrail Networking to apply import routing policies to MP-BGP routes only. This enhancement allows routing policies to match routes from MP-BGP traffic from the MPLS backbone network without also matching routes from BGPaaS. This enhancement works by changing how a routing policy identifies routes when the protocol in the routing policy is set as BGP.

For additional information, see [Routing Policy](#).

Remote Compute Enhancements in Contrail Command

Starting in Contrail Networking Release 21.4.L1, you can better monitor and configure BGP Routers and vRouters associated with remote compute routing clusters in the Contrail Command graphical user interface.

For more information on these enhancements, see the *Monitoring and Configuring BGP Routers for Remote Compute in Contrail Command* and *Viewing the Virtual Router Connected to Remote Compute Clusters in Contrail Command* sections in [Remote Compute](#).

New and Changed Features in Contrail Networking Release 21.4

Support for Automatically Deploying Remote Compute Using RHOSP/TripleO

Starting in Contrail Networking 21.4, you can deploy remote compute automatically using RHOSP/TripleO for edge use cases.

For more information, see [Remote Compute](#).

Graceful Restart or Long-Lived Graceful Restart Support for a EVPN Type 2 Route

Starting in Contrail Networking Release 21.4, the graceful restart or long-lived graceful restart features support the EVPN Type 2 routes to resume the sessions and relearn the routes, reducing overall network impact.

For more information, see [Configuring Graceful Restart and Long-lived Graceful Restart](#).

UBI 8 Support for Contrail Container Images

Starting in Contrail Release 21.4, the Red Hat Universal Base Image 8 (ubi8) is used as the base for Contrail container images.

For more information, see [Understanding Contrail Containers](#).

Upgrade Contrail Software

Starting in Contrail Networking Release 21.4, you can upgrade Red Hat OpenStack Platform (RHOSP) from RHOSP 13 or RHOSP 16.1 to RHOSP 16.2 by leveraging Red Hat Fast Forward Upgrade (FFU) procedure while simultaneously upgrading Contrail Networking from Release 1912.L4 or 2011.L3 to Release 21.4.

For more information, see [Upgrading Contrail Networking Release 1912.L4 or 2011.L3 with RHOSP 13 or RHOSP 16.1 to Contrail Networking Release 21.4 with RHOSP 16.2](#).

Support for Hosts using Red Hat Virtualization (RHV) in Red Hat OpenStack 16 Environments

Starting in Contrail Networking Release 21.4, Contrail Networking has been enhanced to operate in Red Hat OpenStack 16 environments with hosts using Red Hat Virtualization (RHV).

For more information, see the [Setting Up the Infrastructure \(Contrail Networking Release 21.4 or Later\)](#) in the *Using Contrail Networking with OpenStack* section of the [Contrail Networking Installation and Upgrade Guide](#).

Supported Platforms in Contrail Networking Release 21.4

[Table 1 on page 4](#) lists the orchestrator releases and the corresponding operating systems and kernel versions supported for all Contrail Networking Release 21.4 releases through Release 21.4.L2.

Table 1: Supported Platforms

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
Contrail Networking Release 21.4.L2	Kubernetes 1.24.X	Juju Charms	Ubuntu 18.04.5—Linux Kernel Version 4.15.0-206-generic MaaS Version: 2.6.2
	Kubernetes 1.21.4/1.23.15 KubeVirt 0.48.0	Ansible	CentOS 7.9—Linux Kernel Version 3.10.0-1160.66.1.el7.x86_64 Ansible version : 2.7.11 Docker version: 18.03.1-ce/20.10.21

Table 1: Supported Platforms (Continued)

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	Openstack Queens	Ansible	CentOS 7.9—Linux Kernel Version 3.10.0-1160.66.1.el7.x86_64 Ansible version: 2.7.11 Docker version: 18.03.1-ce
	Canonical OpenStack Queens	Juju Charms	Ubuntu 18.04.5—Linux Kernel Version 4.15.0-206-generic MaaS Version: 2.6.2
	Canonical OpenStack Ussuri	Juju Charms	Ubuntu 20.04.5—Linux Kernel Version 5.4.0-139-generic MaaS Version: 3.2.6
	Red Hat OpenStack 16.2.4	RHOSP 16 Director	RHEL 8.4—Linux Kernel Version 4.18.0-305.65.1.el8_4.x86_64 Director Image: rhosp-director-images-16.2-20221124.2.el8ost.noarch Red Hat Content Sync Date: Dec 15, 2022 at 08:30 AM
	Red Hat OpenStack 16.2.3	RHOSP 16 Director	RHEL 8.4—Linux Kernel Version 4.18.0-305.49.1.el8_4.x86_64 Director Image: rhosp-director-images-x86_64-16.2-20220609.1.el8ost.noarch Red Hat Content Sync Date: Nov 03, 07:19 PM
Contrail Networking Release 21.4.L1	Kubernetes 1.23	Juju Charms	Ubuntu 18.04.5—Linux Kernel Version 4.15.0-200-generic MaaS Version: 2.6.2
	Kubernetes 1.21.4	Ansible	CentOS 7.9—Linux Kernel Version 3.10.0-1160.66.1.el7.x86_64 Ansible version : 2.7.11 Docker version: 18.03.1-ce
	Openstack Rocky	Ansible	CentOS 7.9—Linux Kernel Version 3.10.0-1160.66.1.el7.x86_64 Ansible version: 2.7.11 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	CentOS 7.9—Linux Kernel Version 3.10.0-1160.66.1.el7.x86_64 Ansible version: 2.7.11 Docker version: 18.03.1-ce
	Canonical OpenStack Train	Juju Charms	Ubuntu 18.04.5—Linux Kernel Version 4.15.0-200-generic MaaS Version: 2.6.2
	Canonical OpenStack Ussuri	Juju Charms	Ubuntu 20.04.5—Linux Kernel Version 5.4.0-135-generic MaaS Version: 2.4.2
	Red Hat OpenStack Platform 16.2.3	RHOSP 16 Director	RHEL 8.4—Linux Kernel Version 4.18.0-305.25.1.el8_4.x86_64 Director Image: rhosp-director-images-16.2-20211112.1.el8ost.noarch.rpm Red Hat Content Sync Date : Jan 11th, 2022, 05:21 PM PST RHEL 8.4—Linux Kernel Version 4.18.0-305.49.1.el8_4.x86_64 Director Image: rhosp-director-images-x86_64-16.2-20220609.1.el8ost.noarch Red Hat Content Sync Date : Nov 03 2022, 07:19 PM
Contrail Networking Release 21.4	Kubernetes 1.22	Juju Charms	Ubuntu 18.04.5—Linux Kernel Version 4.15.0-187-generic MaaS Version: 2.6.2
	Kubernetes 1.21.4	Ansible	CentOS 7.9—Linux Kernel Version 3.10.0-1160.66.1.el7.x86_64 Ansible version : 2.7.11 Docker version: 18.03.1-ce"
	Openstack Rocky	Ansible	CentOS 7.9—Linux Kernel Version 3.10.0-1160.66.1.el7.x86_64 Ansible version: 2.7.11 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	CentOS 7.9—Linux Kernel Version 3.10.0-1160.66.1.el7.x86_64 Ansible version: 2.7.11 Docker version: 18.03.1-ce"
	Canonical OpenStack Train	Juju Charms	Ubuntu 18.04.5—Linux Kernel Version 4.15.0-187-generic MaaS Version: 2.6.2
	Canonical OpenStack Ussuri	Juju Charms	Ubuntu 20.04.3—Linux Kernel Version 5.4.0-120-generic MaaS Version: 2.4.2
	Red Hat OpenStack Platform 16.2	RHOSP 16 Director	RHEL 8.4—Linux Kernel Version 4.18.0-305.12.1.el8_4.x86_64 Director Image: rhosp-director-images-16.2-20210902.2.el8ost.noarch Red Hat Content Sync Date : October 06, 2021, 05:21 PM PST RHEL 8.4—Linux Kernel Version 4.18.0-305.25.1.el8_4.x86_64 Director Image: rhosp-director-images-16.2-20211112.1.el8ost.noarch.rpm Red Hat Content Sync Date : Jan 11th, 2022, 05:21 PM PST

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 21-based software, see [README Access to Contrail Registry 21.X](#).

Juju Charms Versions

[Table 2 on page 9](#) lists the charms versions supported by Contrail Networking Release 21.4 through Release 21.4.L2.

The published versions of charms are available in the Charmhub on the Juju website. A link to the Charmhub page for each charm is provided in the table.

Charms can also be downloaded from the [tf-charms page](#) on the Tungsten Fabric Github page.

Table 2: Supported Charms Versions

Contrail Networking Release	Charms Version
Contrail Networking Release 21.4.L2	<p><i>Contrail Openstack</i> Supported Version: 37 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-openstack</p> <p><i>Contrail Agent</i> Supported Version: 39 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-agent</p> <p><i>Contrail Controller</i> Supported Version: 36 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-controller</p> <p><i>Contrail Analytics</i> Supported Version: 35 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-analytics</p> <p><i>Contrail AnalyticsDB</i> Supported Version: 34 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-analyticsdb</p> <p><i>Contrail Keystone Auth</i> Supported Version: 31 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-keystone-auth</p> <p><i>Contrail Command</i> Supported Version: 22 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-command</p> <p><i>Contrail Kubernetes Master</i> Supported Version: 30 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-kubernetes-master</p> <p>Contrail Kubernetes Node Supported Version: 30 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-kubernetes-node</p>

Table 2: Supported Charms Versions (*Continued*)

Contrail Networking Release	Charms Version
Contrail Networking Release 21.4.L1	<p><i>Contrail Openstack</i> Supported Versions: 35 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-openstack</p> <p><i>Contrail Agent</i> Supported Versions: 37 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-agent</p> <p><i>Contrail Controller</i> Supported Version: 34 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-controller</p> <p><i>Contrail Analytics</i> Supported Version: 33 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-analytics</p> <p><i>Contrail AnalyticsDB</i> Supported Version: 32 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-analyticsdb</p> <p><i>Contrail Keystone Auth</i> Supported Version: 30 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-keystone-auth</p> <p><i>Contrail Command</i> Supported Version: 30 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-command</p> <p><i>Contrail Kubernetes Master</i> Supported Version: 28 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-kubernetes-master</p> <p>Contrail Kubernetes Node Supported Version: 26 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-kubernetes-node</p>

Table 2: Supported Charms Versions (*Continued*)

Contrail Networking Release	Charms Version
Contrail Networking Release 21.4	<p><i>Contrail Openstack</i> Supported Versions: 34 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-openstack</p> <p><i>Contrail Agent</i> Supported Versions: 36 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-agent</p> <p><i>Contrail Controller</i> Supported Version: 33 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-controller</p> <p><i>Contrail Analytics</i> Supported Version: 32 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-analytics</p> <p><i>Contrail AnalyticsDB</i> Supported Version: 31 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-analyticsdb</p> <p><i>Contrail Keystone Auth</i> Supported Version: 29 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-keystone-auth</p> <p><i>Contrail Command</i> Supported Version: 19 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-command</p> <p><i>Contrail Kubernetes Master</i> Supported Version: 27 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-kubernetes-master</p> <p>Contrail Kubernetes Node Supported Version: 24 Charmhub Page: https://charmhub.io/juniper-os-software-contrail-kubernetes-node</p>

Known Behavior

IN THIS SECTION

- [Known Behavior in Contrail Networking Release 21.4.L2 | 12](#)
- [Known Behavior in Contrail Networking Release 21.4.L1 | 21](#)
- [Known Behavior in Contrail Networking Release 21.4 | 30](#)
- [Known Behavior in Contrail Networking Release 21.3.1 | 37](#)
- [Known Behavior in Contrail Networking Release 21.3 | 43](#)

This section lists known limitations with this release.

Known Behavior in Contrail Networking Release 21.4.L2

OpenStack

- CEM-30003 When a project having resources is renamed in Openstack, the new name of the project is not updated in Contrail Networking Release 21.4.L2. As a result, the project name in Openstack is different from Contrail Networking Release 21.4.L2. However, if the project does not have any resources, when any resource will be created in the respective project, the project name and the resource name are the same in Contrail Networking Release 21.4.L2 as in Openstack.
- CEM-29229 An IPU upgrade from Contrail Networking Release 21.4 to Contrail Networking Release R21.4.L1 may fail with the rabbitmq container removing state.

Workaround: Complete the following procedure:

1. Enter the **pcs resource cleanup** command on all OpenStack nodes and ensure all resources are up.
2. Increase the timeout to 600 seconds by running the **pcs resource op defaults update timeout=600s** command.
3. Enter the **openstack overcloud update converge** command with all parameters.

4. After the update convergence is complete, you may have to clean up some of the converge failures.

Re-enter the **openstack overcloud update converge** command with all parameters to clean up these converge failures.

- CEM-29059 A fast forward upgrade (FFU) from Contrail Networking Release 1912 to Contrail Networking Release 21.4.L1 is not supported. The upgrade attempt fails while upgrading the controller.
- CEM-26599 In deployments using remote compute, the time to create a virtual machine (VM) on a remote site can take an extended amount of time. The VM creation time can take over 6 minutes.
- CEM-27083 During deployment, a core error sometimes occurs with the Contrail Collector service. The Contrail Collector service automatically restarts and no functional impact is observed when this issue is experienced.
- 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.

DPDK and SR-IOV

- CEM-29841 In environments using DPDK, BGP sessions are not coming up due to issues with wrong next hop programming in packet flows.
- 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-29739 In deployments deployed using Ansible, XMPP might go down due to issues with vDNS server scaling.

Workaround: Configure less than 200 vDNS servers.

- CEM-28861 On a compute node using an Intel 3000 smartNIC, the deployment sometimes fails due to issues related to initiating the NIC. The log shows a series of mbuf errors.

Workaround: Restart the contrail-vrouter-dpdk container on the failing node and rerun the deployment. You can also delete the problematic compute node and redeploy the cluster.

- CEM-11163 In Fortville X710 NIC: With TX and RX buffers performance degrade is observed as mbufs gets exhausted.

General Routing

- CEM-30058 In environments using managed physical network function (PNF) devices with an SRX device, the PNF service chain is not pushing the configuration to the SRX device. The PNF service chain, therefore, does not work.

Workaround: No known workaround in environments using managed PNF devices. The issue is not experienced in environments using unmanaged PNF devices.

- CEM-30023 In remote compute environments using BGP, routes are not getting populated for the control node and for the BGP primary and secondary peers.
- CEM-30002 The Contrail Collector is unexpectedly failing.

Workaround: No known workaround. The Contrail Collector services are restarted and return online automatically after the failure. This issue has no impact on traffic forwarding performance.

- CEM-29723 In environments deployed using Ansible, the Contrail Collector is unexpectedly failing.

Workaround: No known workaround. The Contrail Collector services are restarted and return online automatically after the failure. This issue has no impact on traffic forwarding performance.

- CEM-28889 BGP as a Service (BGPaaS) does not work across multiple pods in the same VM due to BGPaaS sessions flapping.

Workaround: Always configure BGPaaS to be deployed with a single pod in a single VM.

- CEM-29152 A configuration push for a logical router interconnect is failing when a routing policy is configured.
- CEM-28914 The Contrail vRouter is dropping packets from a virtual machine (VM) to a bare metal server (BMS) on an external network that is extended to a Juniper Networks MX series router.

Workaround: Disable Reverse Path Forwarding (RPF).

- CEM-27502 To enable QOS on Intel XXV710 NICs, you must first perform the following procedure on the compute node:

1. Reboot the compute node.

2. Enter the **systemctl start lldpad** command.

3. Enter the **systemctl status lldpad dcbtool sc ens1f0 dcb on** command.

You should know be able to enable QoS on the NIC.

If you are unable to enable QoS on the NIC after performing step 3, proceed to step 4.

4. Enter the **dcbtool sc dcbx v:force-cee** command.

- CEM-25388 Updating the DNS server attribute associated with VN does not work. To update, add, or remove the DNS server details in a VN, the user can delete and re-create the VN with desired DNS server details.
- CEM-20477 Incase of EVPN service chain, the service chain with transparent firewall does not establish connection between left and right VNs.
- 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-29418 Hitless software upgrade attempts are failing with a "MODULE FAILURE\nSee stdout/stderr for the exact error" error.

Workaround : QFX series switches can be upgraded manually by logging into the console of the device. You can achieve a hitless upgrade by putting the QFX series switch that is the target of the upgrade into maintenance mode.

- CEM-29154 The Zero Touch Provisioning (ZTP) process fails to onboard an MX Series router due to an IP address not getting assigned to an FXP interface.
- CEM-28971 An attempt to onboard a fabric using Zero Touch Provisioning (ZTP) is not completing. The control and collector processes are in the initializing state and the RabbitMQ procedure is continuously restarting on one control node.
- CEM-28941 Configurations are not getting pushed to the fabric device.

Workaround: Restart RabbitMQ and the DM container.

- 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

- CEM-29109 In Kubernetes environments using the Tungsten Fabric operator (tf-operator), the Contrail Query Engine goes down unexpectedly while deployed.

Workaround: No known workarounds. The Query Engine restarts automatically.

- CEM-29067 In a Kubernetes environment deployed using tf-operator, the Contrail collector process fails during upgrades. The Contrail collector services restart and come back online automatically after the failure.
- JCB-187287 High Availability provisioning of Kubernetes master is not supported.

General

- CEM-30015 During a Fast Forward upgrade from Contrail Networking Release 1912.L4 to 21.4.L2, a RHOSP-upgrade-Kernel crash might be observed on a kernel compute node.
- CEM-29197 A virtual port group (VPG) cannot be deleted after removing a virtual machine interface (VMI). Workaround: Delete the VMI instead of removing the VMI reference from the VPG.
- CEM-29163 The Contrail Agent sometimes fails in remote compute environments the during overcloud load creation process. The agent automatically recovers and restarts.
- CEM-28044 The Contrail Networking upgrade process might backup Config DB nodes in the wrong order.

Workaround: Use updated utility to backup the data instead. Alternatively, manually inspect the backup file, create an ordered list of `contrail_config_database_node` UUIDs (in the correct order). Pass those to the `commandutil` during restore.

The steps to create an ordered list of Config DB nodes are as follows:

1. Download `commandutil` binary from <https://webdownload.juniper.net/swdl/dl/secure/site/1/record/165079.html>.
 2. Copy `commandutil` to `/tmp` folder.
 - `chmod u+x /tmp/commandutil`
 - `docker cp /tmp/commandutil contrail_command:/usr/bin/commandutil_21.4.L2`
 - `docker exec contrail_command commandutil_21.4.L2 convert --intype rdbms --outtype yaml --out /etc/contrail/db.yml -c /etc/contrail/command-app-server.yml; mkdir -p ~/backups; mv /etc/contrail/db.yml ~/backups/`
 - CEM-27916 In environments that were deployed using Juju Charms and are using Contrail Networking with LXD, nodes that are using LXD are not restarting after a Contrail Controller reboot. The reboot triggers an IP tables issue that is not resolvable by Contrail Networking. This IP tables issue is described in the [Unable to reach host gateway from lxd container](#) article on the [Charmhub Discourse](#) page.
- Workaround: Enable forward accept on the nodes using LXD by entering the **`sudo iptables -P FORWARD ACCEPT`** command.
- CEM-25524 In lab scenario with K8s deployment with one controller and one compute, while inter namespace traffic was enabled, `contrail-controller` core was observed very rarely and the system recovered itself after the core. This scenario did not happen in HA environments, which is the recommended deployment for production.

- CEM-25331 On a scaled environment, restarting a controller node may keep the schema transformer service restarting for few times.
- CEM-25109 In the VPG page in UI, VLAN information does not get reflected for internal SRIOV VPG. This is only an intermittent behaviour.
- 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-30038 The Contrail Web application does not display BFD Neighbors information and Appformix for some of the QFX devices.

Workaround: Remove the show system core-dumps command from the Appformix Monitoring command list.

- 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

Known Behavior in Contrail Networking Release 21.4.L1

OpenStack

- CEM-30003 When a project having resources is renamed in Openstack, the new name of the project is not updated in Contrail Networking Release 21.4.L1. As a result, the project name in Openstack is different from Contrail Networking Release 21.4.L1. However, if the project does not have any

resources, when any resource will be created in the respective project, the project name and the resource name are the same in Contrail Networking Release 21.4.L1 as in Openstack.

- CEM-29229 An IPU upgrade from Contrail Networking Release 21.4 to Contrail Networking Release R21.4.L1 may fail with the rabbitmq container removing state.

Workaround: Complete the following procedure:

1. Enter the **pcs resource cleanup** command on all OpenStack nodes and ensure all resources are up.
2. Increase the timeout to 600 seconds by running the **pcs resource op defaults update timeout=600s** command.
3. Enter the **openstack overcloud update converge** command with all parameters.
4. After the update convergence is complete, you may have to clean up some of the converge failures.

Re-enter the **openstack overcloud update converge** command with all parameters to clean up these converge failures.

- CEM-29059 A fast forward upgrade (FFU) from Contrail Networking Release 1912 to Contrail Networking Release 21.4.L1 is not supported. The upgrade attempt fails while upgrading the controller.
- CEM-26599 In deployments using remote compute, the time to create a virtual machine (VM) on a remote site can take an extended amount of time. The VM creation time can take over 6 minutes.
- CEM-27083 During deployment, a core error sometimes occurs with the Contrail Collector service. The Contrail Collector service automatically restarts and no functional impact is observed when this issue is experienced.
- 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.

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-28861 On a compute node using an Intel 3000 smartNIC, the deployment sometimes fails due to issues related to initiating the NIC. The log shows a series mbuf errors.

Workaround: Restart the contrail-vrouter-dpdk container on the failing node and rerun the deployment. You can also delete the problematic compute node and redeploy the cluster.

- CEM-11163 In Fortville X710 NIC: With TX and RX buffers performance degrade is observed as mbufs gets exhausted.

General Routing

- CEM-28889 BGP as a Service (BGPaaS) does not work across multiple pods in the same VM due to BGPaaS sessions flapping.

Workaround: Always configure BGPaaS to be deployed with a single pod in a single VM.

- CEM-29152 A configuration push for a logical router interconnect is failing when a routing policy is configured.
- CEM-28914 When Reverse Path Forwarding (RPF) is enabled in a topology where a virtual machine (VM) on a compute node is sending packets to a bare metal server (BMS) over an external network extended to an MX Series router, reply packets are dropped in the Contrail vRouter. The error messages indicate that the RPF lookup is failing.

Workaround: Disable RPF.

- CEM-27502 To enable QoS on Intel XXV710 NICs, you must first perform the following procedure on the compute node:

1. Reboot the compute node.
2. Enter the **systemctl start lldpad** command.
3. Enter the **systemctl status lldpad** command.

You should now be able to enable QoS on the NIC.

If you are unable to enable QoS on the NIC after performing step 3, proceed to step 4.

4. Enter the **dcbttool sc dcbx v:force-cee** command.

- CEM-25388 Updating the DNS server attribute associated with VN does not work. To update, add, or remove the DNS server details in a VN, the user can delete and re-create the VN with desired DNS server details.
- CEM-20477 In case of EVPN service chain, the service chain with transparent firewall does not establish connection between left and right VNs.
- 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 to a BGPaaS object. You can attach the BFD 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-29418 Hitless software upgrade attempts are failing with a "MODULE FAILURE\nSee stdout/stderr for the exact error" error.

Workaround : QFX series switches can be upgraded manually by logging into the console of the device. You can achieve a hitless upgrade by putting the QFX series switch that is the target of the upgrade into maintenance mode.

- CEM-29154 The Zero Touch Provisioning (ZTP) process fails to onboard an MX Series router due to an IP address not getting assigned to an FXP interface.

- CEM-28971 An attempt to onboard a fabric using Zero Touch Provisioning (ZTP) is not completing. The control and collector processes are in the initializing state and the RabbitMQ procedure is continuously restarting on one control node.
- CEM-28941 Configurations are not getting pushed to the fabric device.

Workaround: Restart RabbitMQ and the DM container.

- 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

- CEM-29067 In a Kubernetes environment deployed using tf-operator, the Contrail collector process fails during upgrades. The Contrail collector services restart and come back online automatically after the failure.
- JCB-187287 High Availability provisioning of Kubernetes master is not supported.

General

- CEM-29197 A virtual port group (VPG) cannot be deleted after removing a virtual machine interface (VMI). Workaround: Delete the VMI instead of removing the VMI reference from the VPG.
- CEM-29163 The Contrail Agent sometimes fails in remote compute environments the during overcloud load creation process. The agent automatically recovers and restarts.

- CEM-28044 The Contrail Networking upgrade process might backup Config DB nodes in the wrong order.

Workaround: Use updated utility to backup the data instead. Alternatively, manually inspect the backup file, create an ordered list of contrail_config_database_node UUIDs (in the correct order). Pass those to the commandutil during restore.

The steps to create an ordered list of Config DB nodes are as follows:

1. Download commandutil binary from <https://webdownload.juniper.net/swdl/dl/secure/site/1/record/165079.html>.
 2. Copy commandutil to /tmp folder.
 - `chmod u+x /tmp/commandutil`
 - `docker cp /tmp/commandutil contrail_command:/usr/bin/commandutil_21.4.L2`
 - `docker exec contrail_command commandutil_21.4.L2 convert --intype rdbms --outtype yaml --out /etc/contrail/db.yml -c /etc/contrail/command-app-server.yml; mkdir -p ~/backups; mv /etc/contrail/db.yml ~/backups/`
- CEM-27916 In environments that were deployed using Juju Charms and are using Contrail Networking with LXD, nodes that are using LXD are not restarting after a Contrail Controller reboot. The reboot triggers an IP tables issue that is not resolvable by Contrail Networking. This IP tables issue is described in the [Unable to reach host gateway from lxd container](#) article on the [Charmhub Discourse](#) page.
- Workaround: Enable forward accept on the nodes using LXD by entering the **sudo iptables -P FORWARD ACCEPT** command.
- CEM-25524 In lab scenario with K8s deployment with one controller and one compute, while inter namespace traffic was enabled, contrail-controller core was observed very rarely and the system recovered itself after the core. This scenario did not happen in HA environments, which is the recommended deployment for production.
 - CEM-25331 On a scaled environment, restarting a controller node may keep the schema transformer service restarting for few times.
 - CEM-25109 In the VPG page in UI, VLAN information does not get reflected for internal SRIOV VPG. This is only an intermittent behaviour.
 - 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

Known Behavior in Contrail Networking Release 21.4

OpenStack

- CEM-30003 When a project having resources is renamed in Openstack, the new name of the project is not updated in Contrail Networking Release 21.4. As a result, the project name in Openstack is different from Contrail Networking Release 21.4. However, if the project does not have any resources, when any resource will be created in the respective project, the project name and the resource name are the same in Contrail Networking Release 21.4 as in Openstack.
- CEM-26599 In deployments using remote compute, the time to create a virtual machine (VM) on a remote site can take an extended amount of time. The VM creation time can take over 6 minutes.
- CEM-27083 During deployment, a core error sometimes occurs with the Contrail Collector service. The Contrail Collector service automatically restarts and no functional impact is observed when this issue is experienced.
- 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.

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-11163 In Fortville X710 NIC: With TX and RX buffers performance degrade is observed as mbufs gets exhausted.

General Routing

- CEM-27502 To enable QOS on Intel XXV710 NICs, you must first perform the following procedure on the compute node:
 1. Reboot the compute node.
 2. Enter the **systemctl start lldpad** command.
 3. Enter the **systemctl status lldpad** command.
 You should know be able to enable QoS on the NIC.
 If you are unable to enable QoS on the NIC after performing step 3, proceed to step 4.
 4. Enter the **dcdbtool sc dcbx v:force-cee** command.
- CEM-25388 Updating the DNS server attribute associated with VN does not work. To update, add, or remove the DNS server details in a VN, the user can delete and re-create the VN with desired DNS server details.
- CEM-20477 Incase of EVPN service chain, the service chain with transparent firewall does not establish connection between left and right VNs.
- 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

- CEM-27916 In environments that were deployed using Juju Charms and are using Contrail Networking with LXD, nodes that are using LXD are not restarting after a Contrail Controller reboot. The reboot triggers an IP tables issue that is not resolvable by Contrail Networking. This IP tables issue is described in the [Unable to reach host gateway from lxd container](#) article on the [Charmhub Discourse](#) page.

Workaround: Enable forward accept on the nodes using LXD by entering the **sudo iptables -P FORWARD ACCEPT** command.

- CEM-25524 In lab scenario with K8s deployment with one controller and one compute, while inter namespace traffic was enabled, contrail-controller core was observed very rarely and the system recovered itself after the core. This scenario did not happen in HA environments, which is the recommended deployment for production.
- CEM-25331 On a scaled environment, restarting a controller node may keep the schema transformer service restarting for few times.
- CEM-25109 In the VPG page in UI, VLAN information does not get reflected for internal SRIOV VPG. This is only an intermittent behaviour.
- 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

Known Behavior in Contrail Networking Release

21.3.1

OpenStack

- 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.

OpenShift

- CEM-25482 While upgrading Contrail in Openshift based environments, contrail-collector core is observed during the upgrade process. As this happens once while the system is in the upgrade window and the service auto recovers, no functionality impact is expected.
- 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-23810 When Generic Segmentation Offload (GSO) is enabled, IPv6 packets with header lengths of 128 bytes or more are dropped.

Workaround: Disable GSO.
- 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-11163 In Fortville X710 NIC: With TX and RX buffers performance degrade is observed as mbufs gets exhausted.

General Routing

- CEM-25388 Updating the DNS server attribute associated with VN does not work. To update, add, or remove the DNS server details in a VN, the user can delete and re-create the VN with desired DNS server details.
- CEM-20477 In case of EVPN service chain, the service chain with transparent firewall does not establish connection between left and right VNs.
- 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

- CEM-27916 In environments that were deployed using Juju Charms and are using Contrail Networking with LXD, nodes that are using LXD are not restarting after a Contrail Controller reboot. The reboot triggers an IP tables issue that is not resolvable by Contrail Networking. This IP tables issue is described in the [Unable to reach host gateway from lxd container](#) article on the [Charmhub Discourse](#) page.

Workaround: Enable forward accept on the nodes using LXD by entering the **sudo iptables -P FORWARD ACCEPT** command.

- CEM-25538 Configuring BGPaaS can bring kernel agent state to initializing.

As a workaround, configure the iptables on compute again (as given in CEM-23911) and restart the agent.
- CEM-25524 In lab scenario with K8s deployment with one controller and one compute, while inter namespace traffic was enabled, contrail-controller core was observed very rarely and the system

recovered itself after the core. This scenario did not happen in HA environments, which is the recommended deployment for production.

- CEM-25331 On a scaled environment, restarting a controller node may keep the schema transformer service restarting for few times.
- CEM-25109 In the VPG page in UI, VLAN information does not get reflected for internal SRIOV VPG. This is only an intermittent behaviour.
- 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

Known Behavior in Contrail Networking Release 21.3

OpenStack

- 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.

OpenShift

- CEM-25482 While upgrading Contrail in Openshift based environments, contrail-collector core is observed during the upgrade process. As this happens once while the system is in the upgrade window and the service auto recovers, no functionality impact is expected.
- 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 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-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-11163 In Fortville X710 NIC: With TX and RX buffers performance degrade is observed as mbufs gets exhausted.

General Routing

- CEM-25388 Updating the DNS server attribute associated with VN does not work. To update, add, or remove the DNS server details in a VN, the user can delete and re-create the VN with desired DNS server details.
- CEM-20477 Incase of EVPN service chain, the service chain with transparent firewall does not establish connection between left and right VNs.

- 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

- CEM-27916 In environments that were deployed using Juju Charms and are using Contrail Networking with LXD, nodes that are using LXD are not restarting after a Contrail Controller reboot. The reboot triggers an IP tables issue that is not resolvable by Contrail Networking. This IP tables issue is described in the [Unable to reach host gateway from lxd container](#) article on the [Charmhub Discourse](#) page.

Workaround: Enable forward accept on the nodes using LXD by entering the **sudo iptables -P FORWARD ACCEPT** command.

- CEM-25538 Configuring BGPaaS can bring kernel agent state to initializing.
As a workaround, configure the iptables on compute again (as given in CEM-23911) and restart the agent.
- CEM-25524 In lab scenario with K8s deployment with one controller and one compute, while inter namespace traffic was enabled, contrail-controller core was observed very rarely and the system recovered itself after the core. This scenario did not happen in HA environments, which is the recommended deployment for production.
- CEM-25331 On a scaled environment, restarting a controller node may keep the schema transformer service restarting for few times.
- CEM-25109 In the VPG page in UI, VLAN information does not get reflected for internal SRIOV VPG. This is only an intermittent behaviour.
- 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

Resolved Issues

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

[Resolved Issues in Contrail Networking Release 21.4.L2](#)

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 Features

The following features are deprecated for the Contrail fabric managements use case in Contrail Networking Release 21.3.

- VNF for BMS workloads. See [Creating VNF Service Chains for Inter-LR Traffic](#).
- Collapsed Spine topology
- MX Series Routers acting as server leaf
- AR-Replicator
- PIM support on Unmanaged PNF use case

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