

Contrail Networking Release 1910

Published
2025-08-18

RELEASE

Table of Contents

Introduction

New and Changed Features

Supported Platforms Contrail Networking Release 1910

Known Behavior

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 1910 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/contrail19/information-products/topic-collections/release-notes/index.html.

New and Changed Features

IN THIS SECTION

- [Encryption Support Between Analytics API Servers and Client Servers | 2](#)
- [Enhanced Routing Policies to Support Modification of Secondary Routes in Virtual Networks | 2](#)
- [Support for Aborting Ongoing Jobs in Contrail Networking Fabric | 2](#)
- [Support for Integrating VMware with Contrail Networking Fabric | 2](#)
- [Support for Monitoring DPDK Enabled Bond Interfaces | 3](#)
- [Support for Trunk Networking Between Contrail Networking and Neutron | 3](#)
- [Support for QFX5120-32C Devices in Lean-Spine Deployments | 3](#)
- [Viewing Overlay Route Information | 4](#)

The features listed in this section are new or changed as of Contrail Networking Release 1910. A brief description of each new feature is included.

Encryption Support Between Analytics API Servers and Client Servers

Starting with Contrail Networking Release 1910, the connection between Analytics API servers and Client servers is encrypted with SSL. The Clients servers connect to the Analytics API server through the REST API Port. In earlier releases, the connection between Analytics API server and the Clients servers was not encrypted, which could pose a security threat.

For more information, see [Encryption Between Analytics API Servers and Client Servers](#).

Enhanced Routing Policies to Support Modification of Secondary Routes in Virtual Networks

Starting with Contrail Networking Release 1910, virtual network routing policies are automatically applied to secondary routes. This feature is especially useful as a mechanism to modify routes imported from MP-BGP, including routes that are imported from the MPLS network, using routing policies.

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

Support for Aborting Ongoing Jobs in Contrail Networking Fabric

Starting with Contrail Networking Release 1910, you can use the Contrail Command user interface (UI) to abort an ongoing fabric job at any stage. In earlier releases, an ongoing fabric job could not be aborted from the Contrail Command UI.

For more information, see [Aborting Ongoing Fabric Jobs](#).

Support for Integrating VMware with Contrail Networking Fabric

Contrail Networking Release 1910 supports integrating VMware with Contrail Networking fabric. A dedicated Contrail vCenter Fabric Manager (CVFM) plugin is deployed for this integration. This plugin is deployed when you install Contrail Command. You can then enable this plugin when you provision Contrail Command. However, if you do not enable this plugin during provisioning, you can enable the plugin from the **Infrastructure>Cluster** page of the Contrail Command UI.



NOTE: VMware-Contrail Networking Fabric integration supports only brownfield device discovery. Greenfield device discovery is not supported.

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

Support for Monitoring DPDK Enabled Bond Interfaces

Starting with Contrail Networking Release 1910, you can use the Contrail Command user interface (UI) to monitor and view the details of master and slave devices that are members of a bond interface. This feature is available for device systems configured with Data Plane Development Kit (DPDK). In earlier releases, you could only see a bond interface in the Contrail Command UI.

For more information, see [Monitoring Bond Interfaces in DPDK Enabled Devices](#).

Support for Trunk Networking Between Contrail Networking and Neutron

Starting with Contrail Networking Release 1910, Contrail integrates with Neutron trunk port APIs, which enables trunk networking between Contrail and Neutron instances. Trunk networking uses trunk extension that is used to multiplex incoming and outgoing packets from multiple Neutron logical networks using a single Neutron logical port. A trunk extension is integrated in Neutron as a collection of Neutron logical ports. In the trunk extension that is implemented, Contrail introduces logical entities defined by OpenStack Trunk API to provide backend support for Neutron Trunk Port API. The Neutron Trunk Port object maps to Contrail Virtual Port Group (VPG) object, which was designed for handling non-LCM BMS workflow and multi-VLAN support.

Support for QFX5120-32C Devices in Lean-Spine Deployments

Starting in Contrail Networking Release 1910, a QFX5120-32C device can be used in lean-spine deployments.

For more information on supported hardware platforms and associated roles, see [Supported Hardware Platforms and Associated Node Profiles and Roles](#).

Viewing Overlay Route Information

Starting with Contrail Networking Release 1910, you can view control node and compute node information using Contrail Command. To view control node information and to see a graphical representation of the overlay routes navigate to the **Infrastructure > Cluster > Cluster Nodes** page. The cluster infrastructure components, including the numbers of control nodes, compute nodes, analytics nodes, config nodes, and database nodes currently operational and also virtual networks are displayed. Click the control and compute node tabs for more information on the nodes. Alternatively, navigate to the **Monitoring > Dashboards** page to view more information as well as graphical representations of control and compute nodes.

For more information, see [Viewing Overlay Routes](#).

Supported Platforms Contrail Networking Release 1910

[Table 1 on page 4](#) lists the orchestrator releases and the corresponding operating systems and kernel versions supported by Contrail Networking Release 1910.

Table 1: Supported Platforms

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
Contrail Networking Release 1910	Kubernetes 1.12.9	Ansible	<ul style="list-style-type: none"> CentOS 7.7—Linux Kernel Version 3.10.0-1062.1.1 Docker version: 18.03.1-ce
	OpenShift 3.11	Ansible	<ul style="list-style-type: none"> RHEL 7.7—Linux Kernel Version 3.10.0-1062.121

Table 1: Supported Platforms (Continued)

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	OpenStack Rocky	Ansible	<ul style="list-style-type: none"> CentOS 7.7—Linux Kernel Version 3.10.0-1062.1.1 <p>Ansible version: 2.5.2</p> <p>Docker version: 18.03.1-ce</p>
	OpenStack Queens	Ansible	<ul style="list-style-type: none"> CentOS 7.7—Linux Kernel Version 3.10.0-1062.1.1 <p>Ansible version: 2.5.2</p> <p>Docker version: 18.03.1-ce</p>
		Juju Charms	<ul style="list-style-type: none"> Ubuntu 18.04.2—Linux Kernel Version 4.15.0-48-generic <p>MaaS Version: 2.4.2</p>
		Helm	<ul style="list-style-type: none"> Ubuntu 16.04.3—Linux Kernel Version 4.4.0-165-generic <p>Docker version: 17.03.2-ce</p> <p>Helm version: 2.7.2</p> <p>Kubernetes version: 1.9.3</p>
	Red Hat OpenStack Platform 13.0.8	RHOSP 13 director	<ul style="list-style-type: none"> RHEL7.7—Linux Kernel Version 3.10.0-1062.1.2
	VMware vCenter 6.7	Ansible	<ul style="list-style-type: none"> ESX version 6.5 <p>CentOS VM version running vRouter: CentOS 7.7—Linux Kernel Version 3.10.0-1062.1.1</p>

Table 2: Supported AppFormix Release

Contrail Networking Release	AppFormix Release	Operating System
Contrail Networking Release 1910	AppFormix 3.1.6	CentOS 7.7

Known Behavior

This section lists known limitations with this release.

- CEM-9278 The sFlow stats for the BMS added after initial provisioning of a cluster is not displayed. As a workaround, to enable sFlow stats for the BMS added post initial provisioning, execute the following:

1. Add the host as Remote Host in AppFormix UI.

Go to AppFormix Swagger API (**Settings > API Documentation > Link to AppFormix Documentation**).

Click **Show/Hide** to get the API Details.

Go to /Hosts POST API.

Set X-Auth-Type as OpenStack and fill the X-Auth-Token with Keystone token. Specify the following in the body:

```
{
  "HostName": "10.84.23.38", <<< Fill your IP
  "AgentBaseUrl": "",
  "Name": "b3s38", <<< Fill the hostname
  "HostType": "kvm",
  "LinkCapacity": "10G",
  "Source": "remote.host",
  "AutomaticInstanceDiscovery": false,
  "ServerId": "b3s38", << fill the hostname
  "Metadata": {}
}
```

Send POST request.

2. Once a device is added in the UI, go to **Settings > Network Devices**. Select the Network Device which you want to add to BMS.

Go to **Edit** section, set **LLD** to **Disabled**, select **SNMP**, click **Next** and set **snmp community string** and click **Save**.

Go to **Edit Connection Info > Continue**, select the **Network Device** and then **Add the Target Device** as BMS and set the interface on Network Device which is connected to this BMS and click **Save**.

Go to Contrail Command UI, the BMS stats can be seen.

- CEM-9201 AS_PATH has duplicate ASN. For example consider a scenario where you create a routing-policy and update as-path with value of 788882. Attach the routing policy to the VN and create a VM in the VN and check for routes in VRF table in the controller. In the VRF table, for a VM-IP, for the route from compute (XMPP Interface), as_path is set as 78888. For the same route, there is an update from another controller, for which as_path is seen as "78888 78888". Here AS_PATH has duplicate ASN and the routing-policy is applied for update from the same sub-cluster.

As a workaround, specify protocol as bgp/xmpp/etc while configuring routing policy so that it is not applied again.

- CEM-8834 RHOSP13 upgrade from release 1909 to release 1910 fails due to https://bugzilla.redhat.com/show_bug.cgi?id=1758703. As a workaround, contact Redhat Support referencing the Bugzilla ID and get the patch. Apply the patch on the overcloud and execute the following commands:

```
openstack overcloud upgrade run --nodes $nodes --playbook upgrade_steps_playbook.yaml
openstack overcloud upgrade run --nodes $nodes --playbook deploy_steps_playbook.yaml
```

- CEM-8717 RHOSP13 Install and upgrade fails due to https://bugzilla.redhat.com/show_bug.cgi?id=1751338. Perform one of the following workarounds to continue with upgrade or install.

- For satellite-based installation

Contact Redhat Support referencing the Bugzilla ID and get the python-paunch package and add it to the Satellite and proceed with install as usual.

- For non-satellite based installation from Redhat CDN

Contact Redhat Support referencing the Bugzilla ID and get the python-paunch package and apply it to the overcloud as follows:

1. Subscribe the overcloud node to RHEL.

2. `enable repos "sudo subscription-manager repos --enable=rhel-7-server-rpms --enable=rhel-7-server-rh-common-rpms --enable=rhel-ha-for-rhel-7-server-rpms --enable=rhel-7-server-openstack-13-rpms --enable=rhel-7-server-extras-rpms"`

3. Extract the **python-paunch*tar** file and remove it.
4. Check if the following three files are present. Note the version is for representative purposes only.

```
[heat-admin@overcloud-novacompute-1 ~]$ ls -lrt
total 228
-rw-rw-r-. 1 heat-admin heat-admin 33584 Sep 13 19:51 python-paunch-2.5.0-6.el7ost.noarch.rpm
-rw-rw-r-. 1 heat-admin heat-admin 168960 Sep 13 19:52 python-paunch-doc-2.5.0-6.el7ost.noarch.rpm
-rw-rw-r-. 1 heat-admin heat-admin 20828 Sep 13 19:52 python-paunch-tests-2.5.0-6.el7ost.noarch.rpm
[heat-admin@overcloud-novacompute-1 ~]$
```

5. Install hot fix.

```
sudo yum localinstall python-paunch-*
```

- CEM-8701 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`. User should avoid using `enterprise_style=True` if the fabric object will onboard BMS LCM instance.
- CEM-7874 User defined alarms may not be generated, when third stunnel/Redis service instance is down after the first two instances were restarted.
- CEM-5334 The multi cloud gateway on the cloud will allow traffic from only a vRouter or Controller nodes to reach to the On-Prem cluster. So in case of deployment where the On-Prem open stack cluster need to be extended to the K8s cluster on the cloud, the k8s master must be defined in one of the vRouters on the cloud.

- CEM-5284 Cloud Compute/vrouter nodes will not be listed in the cluster-nodes/compute node page, all nodes/computes will be listed in the servers page
- CEM-5282 When Azure cloud is extended to On-Prem cluster running on RHEL hosts, contrail-status shows vRouters running on Azure as initializing, though the services are up. This is due to the Red Hat issue <https://access.redhat.com/solutions/2766251>.
- CEM-5141 For deleting compute nodes, the UI workflow will not work. Instead, update the instances.yaml with "ENABLE_DESTROY: True" and "roles:" (leave it empty) and run the following playbooks.

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

For example:

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

- CEM-5043 VNI update on a LR doesnt update the RouteTable. As a workaround, delete the LogicalRouter and create a new LogicalRouter with the new VNI.
- CEM-5042 Adding new subnet on an already provisioned VPC is not supported. If all the subnets are added during initial bringup of VPC, nodes can be added incrementally to the subnets anytime.
- CEM-5041 Provisioning of Region or VPC objects only on the cloud without any nodes is not supported. Add atleast one node while provisioning Region/VPC.
- CEM-5024 Current multi cloud provisioning does not enable the On-prem TOR to exchange public cloud subnets with the On-Prem controllers. The user need to add static routes on the controllers to all the public cloud subnets.

- CEM-4943 After deleting and reprovisioning public cloud infra, though the nodes get deleted from the cloud, the API server and Kubernetes will have stale entries for the deleted objects. To clean up the stale entries, run the following housekeeping scripts:

1. Log in to the command container.
2. Navigate to the **contrail-multi-cloud** folder.

```
cd /usr/share/contrail/contrail-multi-cloud/
```

3. Run the following script.

```
TF_STATE=/root/contrail-multi-cloud/terraform.tfstate INVENTORY=inventories/inventory.yml  
TOPOLOGY=/root/contrail-multi-cloud/topology.yml ./housekeeper.sh
```



NOTE: If you run the script after provisioning, ensure that TF_STATE is the backup file. For example:

```
TF_STATE=/root/contrail-multi-cloud/terraform.tfstate.backup  
INVENTORY=inventories/inventory.yml TOPOLOGY=/root/contrail-multi-cloud/  
topology.yml ./housekeeper.sh
```

- CEM-4941 The multicloud gateway on the public cloud cannot be shared across different subnets. Each subnet must have its own gateway.
- CEM-4865 Provisioning of Contrail Controllers on public cloud is not supported. Controllers need to be provisioned On-prem.
- CEM-4467 On DPDK computes, sometimes VM creation fails with "Connection is closed" error. The issue is not related to any of the contrail components. It is related to systemd-machined service in registering VMs. As a workaround, restart the systemd-machined service to fix the issue.
- CEM-4381 Contrail Fabric device manager tasks can fail if one or more Contrail API servers is down. Contrail-status on the Contrail config nodes can be used to determine if this situation occur.
- CEM-4370 After creating a PNF Service Instance, the fields like PNF eBGP ASN*, RP IP Address, PNF Left BGP Peer ASN*, Left Service VLAN*, PNF Right BGP Peer ASN*, Right Service VLAN* cannot be modified. If there is a need to modify these values, delete and re-create the Service Instance with intended values.

- CEM-4190 IPtables rules are not updated on MC-GW nodes. As a workaround, you must configure IPtables on the on-premise MC-GW nodes with INPUT and FORWARD and default ACCEPT policy.
- CEM-3959 BMS movement across TORs is not supported. To move BMS across TORs the whole VPG need 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-3324 Users cannot provision Contrail Cluster entirely in Public cloud. Contrail Cluster need to be On-Prem and vRouters can be extended to public cloud.
- JCB-204796 In a Helm-based provisioned cluster, VM launch fails if MariaDB replication is set to >1.
- JCB-202874 After deleting a vRouter chart with DPDK, the NICS do not rebind to the host in Helm.
- JCB-190956 While creating ironic-provision, service address in the subnet must be pointing to openstack ironic node ip/kolla internal vip.
- JCB-187320 On a DPDK compute vif `list -rate core-dumps` with traffic.
- JCB-187287 High Availability provisioning of Kubernetes master is not supported.
- JCB-186493 When a snapshot of an active VM fails, shutdown the VM before generating the snapshot.
- JCB-184837 After provisioning Contrail by using a Helm-based provisioned cluster, restart nova-compute container.
- JCB-184776 When the vRouter receives the head fragment of an ICMPv6 packet, the head fragment is immediately enqueued to the assembler. The flow is created as hold flow and then trapped to the agent. If fragments corresponding to this head fragment are already in the assembler or if new fragments arrive immediately after the head fragment, the assembler releases them to flow module. Fragments get enqueued in the hold queue if agent does not write flow action by the time the assembler releases fragments to the flow module. A maximum of three fragments are enqueued in the hold queue at a time. The remaining fragments are dropped from the assembler to the flow module.

As a workaround, the head fragment is enqueued to assembler only after flow action is written by agent. If the flow is already present in non-hold state, it is immediately enqueued to assembler.
- JCB-177787 In DPDK vRouter use cases such as SNAT and LBaaS that require netns, jumbo MTU cannot be set. Maximum MTU allowed: <=1500.
- JCB-177541 When you receive an error message during Kolla provisioning, rerunning the code will not work. In order for the provisioning to work, restart provisioning from scratch.
- JCB-171466 Metadata SSL works only in HA deployment mode.

- JCB-163773 A false alarm for config service is generated when config and configdb services are installed on different nodes. Ignore the false alarm.
- JCB-162927 SR-IOV with DPDK co-existence deployment is not supported using contrail-helm-deployer.