

Contrail Networking Release 1911

Published
2025-08-18

RELEASE

Table of Contents

Introduction

New and Changed Features

Supported Platforms Contrail Networking Release 1911

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 1911 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

- [Adding Leaf or Spine Devices to a Greenfield Fabric | 2](#)
- [Configuring AppFormix Flows from Contrail Command UI | 2](#)
- [Support for Google Cloud Platform | 2](#)
- [Support for Packer | 2](#)
- [Device Functional Groups and Grouping Fabric Devices | 3](#)
- [Support for Increased vRouter Next Hop Limit and Monitoring Next Hop and MPLS Labels Usage | 3](#)
- [Support for ZTP in HA Cluster | 3](#)

The features listed in this section are new or changed as of Contrail Networking Release 1911. A brief description of each new feature is included.

Adding Leaf or Spine Devices to a Greenfield Fabric

Starting with Contrail Networking Release 1911, you can add new leaf or spine devices to expand an existing greenfield fabric. The feature is especially useful when you do not add all the required devices to the fabric on day one and want to add devices to the fabric at a later point. You can add new devices to a fabric by uploading a YAML file that contains the device information.

For more information, see [Adding a Leaf or Spine Device to an Existing Fabric Using ZTP](#).

Configuring AppFormix Flows from Contrail Command UI

Starting with Contrail Networking Release 1910, AppFormix flow is integrated in the Contrail Command UI. You can create and attach telemetry profiles to devices in a Contrail-managed datacenter fabric and view telemetry information for those devices. Contrail Command now acts as a single pane of glass where you can access the features of both Contrail Networking and AppFormix, providing you a unified telemetry experience.

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

Support for Google Cloud Platform

Starting with Contrail Networking Release 1911, you can configure and monitor cloud environments created in Google Cloud Platform (GCP) using Contrail Command. To provision a GCP environment in Contrail Command, navigate to the **Multi Cloud > Create** page and select **GCP** from the **Type of Cloud** list.

For more information, see [Deploying Contrail Multicloud with Contrail Command](#).

Support for Packer

Starting with Contrail Networking Release 1911, a Contrail Command plugin allows Packer to provision images from cloud service providers for use with Contrail Networking. The Packer tool automates the image creation process to create reusable, provisioned images that install more quickly and more reliably than images that are created in non-automated environments.

The Packer tool automatically operates in Contrail Command; no user configuration is required or possible. For more information on Packer, see [Why Use Packer?](#).

Device Functional Groups and Grouping Fabric Devices

Starting with Contrail Networking Release 1911, for certain operations, like OS version upgrade and role assignment, you can assign properties to a user-defined group of devices called device functional groups (DFGs). These properties are applied while provisioning fabric devices using Zero Touch Provisioning (ZTP) or during device Return Material Authorization (RMA). Device functional groups are assigned to devices in the `fabric_ztp.yml` file under **Device Info** used during fabric creation in the UI.

You can view existing groups on the **Device Function Groups** tab of the **Infrastructure > Fabrics** page. To create a new group, click **Create** on the top right corner of the page. For more information, see [Grouping Fabric Devices and Roles Using Device Functional Groups](#).

Support for Increased vRouter Next Hop Limit and Monitoring Next Hop and MPLS Labels Usage

Starting with Contrail Networking Release 1911, the next hop value in the vRouter is increased to 32 bits. By default, the vRouter creates 512K next hops and it supports up to 1 million next hops. In release 1911, you can also configure a watermark limit in vRouter agent configuration file, which enables you to monitor the usage and availability of next hops and Multiprotocol Label Switching (MPLS) labels. In earlier releases, Contrail vRouter supported 16 bits next hop value, which enabled it to create a maximum of only 65,536 next hops.

For more information, see [Contrail vRouter Next Hop Configuration](#).

Support for ZTP in HA Cluster

Contrail Networking Release 1911 supports High Availability (HA) cluster to manage fabrics. With the introduction of this high availability scenario, the DHCP server (dnsmasq) runs only during the zero-touch-provisioning (greenfield onboarding) process. After the fabric onboarding process is complete, the config files that are generated by the device manager and that are applied to dnsmasq, are deleted. After the files are deleted, the dnsmasq will not serve any more clients on the ZTP network. Starting in release 1911, lease file records are maintained in an external storage called Cassandra database.

For more information, see [Using HA Cluster to Manage Fabric](#).

Supported Platforms Contrail Networking Release 1911

Table 1 on page 4 lists the orchestrator releases and the corresponding operating systems and kernel versions supported by Contrail Networking Release 1911.

Table 1: Supported Platforms

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
Contrail Networking Release 1911	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 3.10.0-1062.4.1
	OpenStack Rocky	Ansible	<ul style="list-style-type: none"> CentOS 7.7—Linux Kernel Version 3.10.0-1062.1.1 Ansible version: 2.5.2 Docker version: 18.03.1-ce
	OpenStack Queens	Ansible	<ul style="list-style-type: none"> CentOS 7.7—Linux Kernel Version 3.10.0-1062.1.1 Ansible version: 2.5.2 Docker version: 18.03.1-ce
		Juju Charms	<ul style="list-style-type: none"> Ubuntu 18.04.2—Linux Kernel Version 4.15.0-48-generic MaaS Version: 2.4.2

Table 1: Supported Platforms *(Continued)*

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
		Helm	<ul style="list-style-type: none"> Ubuntu 16.04.3—Linux Kernel Version 4.4.0-165-generic Docker version: 17.03.2-ce Helm version: 2.7.2 Kubernetes version: 1.9.3
	Red Hat OpenStack Platform 13.0.9	RHOSP 13 director	<ul style="list-style-type: none"> RHEL7.7—Linux Kernel Version 3.10.0-1062.4.1
	VMware vCenter 6.7	Ansible	<ul style="list-style-type: none"> ESX version 6.5 CentOS VM version running vRouter: CentOS 7.7—Linux Kernel Version 3.10.0-1062.1.1

[Table 2 on page 5](#) lists the AppFormix release to use with Contrail Networking Release 1911.

Table 2: AppFormix Release

Contrail Networking Release	AppFormix Release	Operating System
Contrail Networking Release 1911	3.1.9 - AppFormix	CentOS 7.7
	1.0.5 - AppFormix Flows	

Known Behavior

This section lists known limitations with this release.

- CEM-21728 Updating a BGPaaS resource (from the heat template) using Heat fails with *"HTTP Status: 500 Content: KeyError: 'uuid'"* error message.
- CEM-10199 In public cloud deployments, after deleting the public cloud, the snapshots are left in the cloud. To clear them, the user has to log in to the respective cloud console (AWS/Azure/GCP) and deregister the AMLs and delete the snapshots from there.
- CEM-9979 During upgrade of DPDK computes deployed with OOO Heat Templates in RHOSP environment, vRouter coredumps are observed. This is due to the sequence in which the services are started during upgrade and does not have impact on cluster operation.
- CEM-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-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-openshift-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-5141 For deleting compute nodes, the UI workflow will not work. Instead, update the `instances.yaml` with `"ENABLE_DESTROY: True"` and `"roles:"` (leave it empty) and run the following playbooks.

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

For example:

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

...

...

- CEM-5043 VNI update on a LR doesn't update the RouteTable. As a workaround, delete the LogicalRouter and create a new LogicalRouter with the new VNI.
- CEM-5041 Provisioning of Region or VPC objects only on the cloud without any nodes is not supported. Add at least 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 needs 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.