Contrail Release 4.1.3 Release Notes

Release 4.1.3
March 2019

Contents

Introduction

Juniper Networks Contrail 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 4.1.3 of Juniper Networks Contrail. They describe new features, limitations, and known problems.

These release notes are displayed on the Juniper Networks Contrail Documentation Web page at https://www.juniper.net/documentation/en_US/contrail4.1/information-products/topic-collections/release-notes/index.html.

New and Changed Features

The features and enhancements listed in this section are new or changed as of Contrail Release 4.1. A brief description of each new feature is included.

New and Changed Features in Contrail Release 4.1.3

There are no new features in Contrail Release 4.1.3.

New and Changed Features in Contrail Release 4.1.2

The feature listed in this section is new as of Contrail Release 4.1.2.

Support for SmartNIC from Netronome

Contrail Release 4.1.2 supports Netronome SmartNIC. You can use Juju to deploy Contrail Release 4.1 with Netronome SmartNICs. The Netronome SmartNIC improves Contrail SDN performance, by saving host resources and providing a stable, high-performance infrastructure.

The Netronome SmartNIC has all server-side networking features, such as overlay networking based on MPLS over UDP/GRE and VXLAN. It supports DPDK, SR-IOV, and Express Virtio (XVIO) for data plane acceleration.

New and Changed Features in Contrail Release 4.1.1

The feature listed in this section is new as of Contrail Release 4.1.1.

Support for Flat Provider Network on SR-IOV Virtual Functions

Contrail Release 4.1.1 supports configuration of VLAN ID 0 on single-root I/O virtualization (SR-IOV) virtual functions to allow multiple VLAN traffic to a virtual machine (VM) running over a single SR-IOV interface.

Support for SR-IOV, DPDK and vRouter on RHEL

Contrail Release 4.1.1 supports SR-IOV, Data Plane Developer Kit (DPDK) and the Contrail vRouter kernel module on Red Hat Enterprise Linux (RHEL) operating systems.

New and Changed Features in Contrail Release 4.1

The features listed in this section are new as of Contrail Release 4.1.

Using Huge Pages to Facilitate vRouter Hash Table Handling

To facilitate vRouter handling of flow and bridge tables at bootup, Contrail Release 4.1 requires the user to enable huge pages (1G in Linux), so that sufficient contiguous memory is available to the vrouter module. Huge page allocation and usage for the vrouter is in the kernel space. Enable huge pages at installation to use this feature.

Simple Underlay Connectivity without Gateway

For simple enterprise use cases and public cloud environments, it is possible to directly route packets using the IP fabric network without using an SDN gateway.

The following features can be enabled when using this method:

- Network policy support for IP fabric
- Security groups for VMs and containers on IP fabric
- Security groups for vhost0 interface, to protect compute node or bare metal server applications
- Support for service chaining, if policy dictates that traffic goes through a service chain.

See *Simple Underlay Connectivity without Gateway*.

Contrail Support for SR-IOV on RHEL

Starting in Release 4.1, Contrail supports single root I/O virtualization (SR-IOV) on Red Hat Enterprise Linux (RHEL) operating systems. Contrail Release 3.0 through Release 4.0 supported SR-IOV on Ubuntu systems only.

For more information, see *Configuring Single Root I/O Virtualization (SR-IOV)*.

Bidirectional Forwarding and Detection Health Check over Virtual Machine Interfaces

Contrail Release 4.1 supports BFD-based health check for VMIs.

Health check for VMIs is already supported in earlier releases as poll-based checks with ping and curl commands. When enabled, these health checks run periodically, once every

few seconds. Consequently, failure detection times can be quite large and are always in seconds.

Health checks based on the BFD protocol can provide failure detection and recovery in sub-second intervals, because applications are notified immediately upon BFD session state changes.

See *Service Instance Health Checks*.

Bidirectional Forwarding and Detection Health Check for BGPaaS

Contrail Release 4.1 adds support for BFD-based health check for BGP as a Service (BGPaaS) sessions.

The BFD-based health check over VMIs, also introduced in Contrail Release 4.1, cannot be directly used for a BGPaaS session, because the session shares a tenant destination address over a set of VMIs, with only one VMI active at any given time.

When configured, any time a BFD-for-BGP session is detected as down by the health checker, corresponding logs and alarms are generated.

To enable this health check, configure the **ServiceHealthCheckType** property and associate it with a bgp-as-a-service configuration object. This can also be accomplished in the Contrail WebUI.

See *Service Instance Health Checks*.

Health Check of Transparent Service Chain

Contrail Release 4.1 enhances service chain redundancy by implementing an end-to-end health check for the transparent service chain. The service health check monitors the status of the service chain and if there is a failure, the control node no longer considers the service chain as a valid next hop, triggering traffic failover.

A segment-based health check is used to verify the health of a single instance in a transparent service chain. The user creates a service-health-check object, with type **segment-based**, and attaches it to either the left or right interface of the service instance. The service health-check packet is injected to the interface to which it is attached. When the packet comes out of the other interface, a reply packet is injected on that interface. If health check requests fail after 30-second retries, the service instance is considered unhealthy and the service VLAN routes of the left and right interfaces are removed. When the agent receives health-check replies successfully, it adds the retracted routes back on both interfaces, which triggers the control node to start reoriginating routes to other service instances on that service chain.

See *Service Instance Health Checks*.

More Efficient Flow Queries

Flow queries are now analyzed on a 7-tuple basis, enabling more efficient flow queries by focusing on elements more important for analysis, and de-emphasizing lesser elements. More efficient queries enable load reduction and allow application of security policy.

An enhanced security framework is implemented to manage connectivity between workloads, or VMIs. Each VMI is tagged with the attributes of Deployment, App, Tier, and Site, and the user specifies security policies for VMIs using the values of these tags.

The existing **FlowLogData** is replaced by **SessionEndpointData**, and a **SessionAggregate** map provides statistics about the flow sessions and the security tags. Session data can belong to either Sampled or Logged Flows. **SessionAggregates** are sent to configurable destinations, including collector, local log, and syslog.

RBAC for Analytics API and WebUI—Beta

Role-based access control (RBAC) for analytics API provides the ability to access UVE and query information based on the permissions of the user for the UVE or queried object. Previously, the analytics API supported authenticated access only for the cloud-admin role. However, to display network monitoring for tenant pages in the UI, the analytics API now supports RBAC (similar to that of the config API) so that tenants can view information about the networks for which they have the read permissions. Tenants will not be able to view system logs and and flow logs, which are only viewable by the cloud-admin role. A non-admin user will be able to see only non-global UVEs.

In the **/etc/contrail/contrail-analytics-api.conf**, the section **DEFAULTS**, the parameter **aaa_mode** now supports **rbac** as one of the values.

See *Role-Based Access Control for Analytics*.

Security Policy Enhancements

As the Contrail environment has grown and become more complex, it has become harder to achieve desired security results with the existing network policy and security group constructs. The Contrail network policies have been tied to routing, making it difficult to express security policies for environments such as cross sectioning between categories, or having a multi-tier application supporting development and production environment workloads with no cross environment traffic.

Contrail 4.1 introduces new firewall security policy objects, including the following enhancements:

- Routing and policy decoupling—introducing new firewall policy objects, which decouples policy from routing.
- Multi dimension segmentation—segment traffic and add security features, based on multiple dimensions of entities, such as Application, Tier, Deployment, Site, UserGroup.
- Policy portability—security policies can be ported to different environments, such as ‘from development to production’, ‘from pci-complaint to production’, ‘to bare metal environment’ and ‘to container environment’.

See *Security Policy Enhancements*.

Allocation of Service Instance IP

In service chaining version 2, for scaling up, the **contrail-svc-monitor** allocates a service instance IP address from the same subnet currently in use. If the scaling is not required, the IP is wasted, from a limited pool of IPs.

Starting with Contrail 4.1, any new service instance allocates IPs from a different subnet, by using a fixed value for ther IP, allocated from 0.0.0.0/8 and ::ffff/104 for IPv4 and IPv6.

Existing service instances retain use of the previous method of allocating IPs; new instances make use of the new allocation method.

Long-Lived Graceful Restart for XMPP

Contrail Release 4.1 introduces support for long-lived graceful restart (LLGR) with XMPP helper mode. Previous versions of Contrail provided only the BGP helper mode. Graceful restart and long-lived graceful restart can be enabled using the Contrail web UI or by using the provision_control script.

In the web UI, you can control the helper modes at **Configure > Infrastructure Global Config > Edit BGP Options**, see Figure 1 on page 7.

Figure 1: Edit BGP Options Page

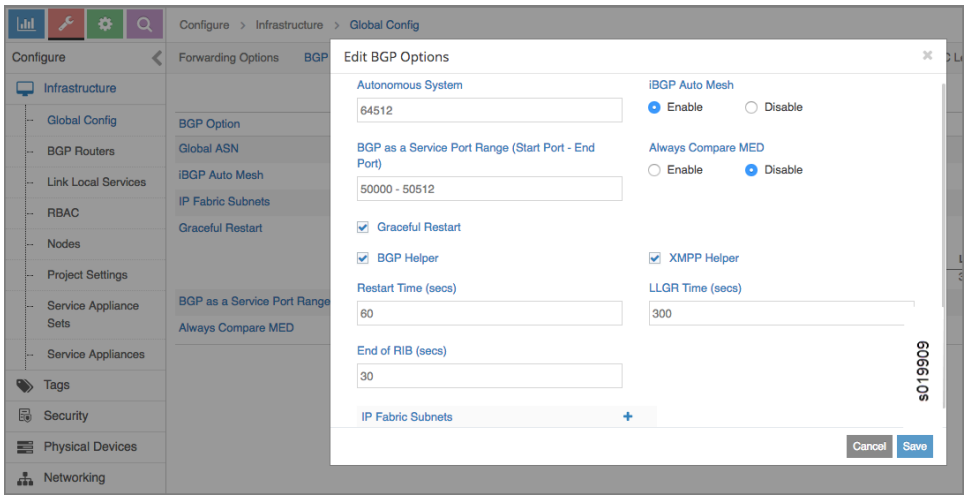

The helper modes can also be enabled via schema, and can be disabled selectively in a Contrail control node for BGP or XMPP sessions by configuring **gr_helper_disable** in the **/etc/contrail/contrail-control.conf** configuration file.

For more information, see *Configuring Graceful Restart and Long-lived Graceful Restart*.

Proxy Encryption of Interactions of vRouter and Nova API

OpenStack allows VMs to access metadata by sending an HTTP request to the link local address 169.254.169.254. The request is proxied to Nova API and HTTP header fields are added, which Nova uses to identify the source instance and respond with appropriate metadata. In Contrail, the vRouter is the proxy, trapping the metadata requests, adding the header fields, and sending the requests to the Nova API server. Previously, these requests were not encrypted, posing a security risk.

In Contrail 4.1, SSL is used to encrypt the HTTP interactions between the Contrail vRouter and Nova API.

To enable this encryption on the Nova side, add the following configuration in the default section of the **nova.conf** file.

```
enabled_ssl_apis = metadata
nova_metadata_protocol = https
nova_metadata_insecure = False
ssl_cert_file = cert.pem
ssl_key_file = privkey.pem
ssl_ca_file = cacert.pem
```

To enable this encryption on the Contrail vrouter agent, add the following configuration in the **METADATA** section of **contrail-vrouter-agent.conf**.

```
metadata_use_ssl = True
metadata_client_cert = client_cert.pem
metadata_client_key = client_key.pem
metadata_ca_cert = cacert.pem
```

Contrail provisioning is updated to populate the configuration files and to copy the certificate files to the appropriate paths.

Contrail EVPN-VXLAN Support Using QFX Series Switches

Contrail Release 4.1 enables you to use Ethernet VPN (EVPN) with Virtual Extensible LAN protocol (VXLAN) encapsulation when you have an environment that includes both virtual and bare metal devices. MX Series routers use EVPN-VXLAN encapsulation to provide both Layer 2 and Layer 3 connectivity for end stations within a Contrail virtual network (VN).

Two types of encapsulation methods are used in virtual networks:

- MPLS-over-GRE (generic routing encapsulation) is used for Layer 3 overlay virtual network routing between Contrail and MX Series routers.
- EVPN-VXLAN is used for Layer 2 overlay virtual network connectivity between virtual machines on Contrail, bare-metal servers attached to QFX Series switches, and their respective Layer 3 gateway configured on the QFX Series switch. Subsequently, inter-VXLAN routing between virtual machines and bare-metal servers, and between bare-metal servers on different VXLAN network identifiers (VNIs), is performed on the QFX Series switch.

For more information, see *EVPN-VXLAN Support for Bare Metal Devices and QFX Device Configuration*.

Supported Platforms Contrail 4.1

Table 1 on page 9 lists the operating system versions and the corresponding Linux or Ubuntu kernel versions supported by Contrail Release 4.1.

Table 1: Supported Platforms

Contrail Release	Orchestrator Release	Operating System and Kernel Versions
Contrail Release 4.1.3	OpenStack Ocata	• RHEL 7.5—Linux kernel version 3.10.0-862.11.6 and Linux kernel version 3.10.0-957 (RHOSP11) • Ubuntu 16.04.2—Linux kernel version 4.4.0-116-generic • VMware vCenter 6.0, 6.5—Ubuntu 16.04.2 kernel version 4.4.0-62-generic
	OpenStack Newton	• RHEL 7.5—Linux kernel version 3.10.0-862.11.6 • RHEL 7.6—Linux kernel version 3.10.0-957 (RHOSP10) • Ubuntu 16.04.2—Linux kernel version 4.4.0-116-generic
	OpenStack Mitaka	• Ubuntu 14.04.5—Linux kernel versions 3.13.0-142-generic and 4.4.0-116-generic
Contrail Release 4.1.2	Kubernetes 1.7.5	• Ubuntu 16.04.2—Linux kernel version 4.4.0-116-generic
	OpenStack Ocata	• RHEL 7.5—Linux kernel version 3.10.0-862.11.6 and Linux kernel version 3.10.0-957 (RHOSP11) • Ubuntu 16.04.2—Linux kernel version 4.4.0-116-generic • VMware vCenter 6.0, 6.5—Ubuntu 16.04.2 kernel version 4.4.0-62-generic
	OpenStack Newton	• RHEL 7.5—Linux kernel version 3.10.0-862.11.6 • Ubuntu 16.04.2—Linux kernel version 4.4.0-116-generic
	OpenStack Mitaka	• Ubuntu 14.04.5—Linux kernel versions 3.13.0-142-generic and 4.4.0-116-generic
Contrail Release 4.1.1	Kubernetes 1.7.5	• Ubuntu 16.04.2—Linux kernel version 4.4.0-116-generic
	Openshift 3.6	• RHEL 7.5—Linux kernel version 3.10.0-862.3.2
	OpenStack Ocata	• RHEL 7.5—Linux kernel version 3.10.0-862.3.2 (RHOSP11) • Ubuntu 16.04.2—Linux kernel version 4.4.0-116-generic • VMware vCenter 6.0, 6.5—Ubuntu 16.04.2 kernel version 4.4.0-62-generic
	OpenStack Newton	• RHEL 7.5—Linux kernel version 3.10.0-862.3.2 (RHOSP10) • Ubuntu 16.04.2—Linux kernel version 4.4.0-116-generic
	OpenStack Mitaka	• Ubuntu 14.04.5—Linux kernel versions 3.13.0-142-generic and 4.4.0-116-generic

Table 1: Supported Platforms (continued)

Contrail Release	Orchestrator Release	Operating System and Kernel Versions
Contrail Release 4.1	Kubernetes 1.7.5	• Ubuntu 16.04.2—Linux kernel version 4.4.0-62-generic
	Openshift 3.6	• RHEL 7.4—Linux kernel version 3.10.0-693
	OpenStack Ocata	• RHEL 7.4—Linux kernel version 3.10.0-693 (RHOSP11) • Ubuntu 16.04.2—Linux kernel version 4.4.0-62-generic • VMware vCenter 6.0, 6.5—Ubuntu 16.04.2 kernel version 4.4.0-62-generic
	OpenStack Newton	• RHEL 7.4—Linux kernel version 3.10.0-693 (RHOSP10) • Ubuntu 16.04.2—Linux kernel version 4.4.0-62-generic
	OpenStack Mitaka	• Ubuntu 14.04.5—Linux kernel versions 3.13.0-110-generic and 4.4.0-34-generic

NOTE: In Contrail Release 4.0 and later, if the stock kernel version of your Ubuntu system is other than the required version, you can upgrade the kernel for all nodes in the cluster by using the following parameter in cluster.json for Server Manager or SM-Lite provisioning or testbed.py.

```
{
  "cluster" : [{
    "parameters" : {
      "provisioning" :  {
        "contrail" : {
          "kernel_upgrade" : true
        }
      }
    }
  }]
}
```

Known Behavior

This section lists known limitations with this release.

Known Behavior in Contrail Release 4.1.3

- Redhat BZ 1684271: Installation of Contrail Release 4.1.3 release on Red Hat 7.6 does not work. Contrail Release 4.1.3 is tested on Red Hat 7.5.
- 1797358 - In a Netronome based Smart NIC deployment, gateway less forwarding feature does not work with nfp_p1 interface. This works with nfp_p0 interface.
- 1796812 - In Netronome based Smart NIC deployments, Launching a VM directly using net-ids is not supported. Port-ids need to be used instead.
- 1681680 When the DPDK vRouter fragments packets before sending them on the wire, the reassembly of the fragments on the receiver might time out in some cases.
- 1694343 In DPDK vRouter use-cases (SNAT, LBaaS) that require **netns** to be launched, do not set Jumbo frames. Use MTU <= 1500 bytes.
- 1705795 On an RHOSP10 provisioned cluster, if the vrouter-agent gets restarted, vhost0 interface does not come up. The **service supervisor-vrouter restart** command brings the service back online.
- 1709974 TSN support in RHOSP-based clusters are supported upto RHOSP10. As a workaround, deploy the stack with computes and DPDK first. Then change the **VrouterPhysicalInterface**, add the TSN nodes, and update the stack.
- 1711256 Project isolation is not supported in nested mode. In nested mode, Namespaces-isolation results in a virtual-network creation and doesn't create a new project.
- 1716308 When the head fragment is received in the vRouter, the head fragment is enqueued to the assembler immediately upon arrival. 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. If agent does not write flow action by the time the assembler releases fragments to the flow module, fragments get enqueued in the hold queue. As a maximum of only three fragments are enqueued in the hold queue, rest of the fragments from the assembler get dropped in the flow module. This leads to the whole packet being dropped on the receive side leading to the first packet loss.
- 1718807 In OVSDB case, the routes are exported from the ToR Agent where the SG is appropriately updated so that inter-virtual network traffic doesn't require any explicit SG to be configured for it to pass. When TOR is peering with control node, this SG has to be explicitly configured and this behavior is expected.
- 1720990 With policy-based mirroring with ECMP destinations, one of the destination vRouter drops packet with invalid NH.
- 1721620 VNC API sends an update of all subfields in a field, like virtual_network_properties and not the updated subfields, like allow_transit of vn_properties alone, hence matching against all the subfields due to which you might not be able to update a subfield. As a workaround, perform the following steps:

  1. Delete the attributes which are not updated from the field class. For example, all the attributes of virtual_network_properties except allow_transit.

2. Call vn_obj.set_virtual_network_properties().
3. Use RestApi or Contrail-UI instead of vnc_api.

- 1724357 While provisioning a RHOSP10 cluster with DPDK nodes, the DPDK node power state goes offline during introspection stage. As a workaround, delete the DPDK nodes from Ironic configuration and add them with the right configuration.

  1. **ironic node-delete** ***node-name***
  2. Create a JSON file configuration as shown in the following example for all DPDK nodes:

```
{
  "nodes": [
    {
      "mac": [
        "90:e2:ba:4c:67:3d"
      ],
      "name": "compute3-dpdk",
      "capabilities" : "profile:compute-dpdk",
      "pm_user": "admin",
      "pm_addr": "10.87.122.164",
      "pm_password": "admin",
      "pm_type": "pxe_ipmitool"
    }
  ]
}
```

  3. **openstack baremetal import --json path to .json**
  4. **openstack baremetal introspection bulk start**

- 1728802 Session logging: Incorrect VN information seen for sessions on transparent SI VMIs. As a workaround, ensure that VLAN NH inherits the policy status from its associated interface.
- 1729059 You must use Ansible version 2.3 to install Contrail using **contrail-ansible** for Kubernetes and OpenShift deployments.
- 1735057 When bringing up Contrail cluster on Red Hat container, manually install docker-py on all the target nodes.

  - To install Pip, use the following command:

```
wget https://bootstrap.pypa.io/get-pip.py
python get-pip.py
```

  - To install docker-py, use the following command:

```
pip install docker-py
```

Known Behavior in Contrail Release 4.1.2

- 1794702 - vRouter crashes if network policy contains large number of Community Tags. Recommended number for 4.1.2 code is less than 50.
- 1797358 - In a Netronome based Smart NIC deployment, gateway less forwarding feature does not work with nfp_p1 interface. This works with nfp_p0 interface.
- 1796812 - In Netronome based Smart NIC deployments, Launching a VM directly using net-ids is not supported. Port-ids need to be used instead.
- 1681680 When the DPDK vRouter fragments packets before sending them on the wire, the reassembly of the fragments on the receiver might time out in some cases.
- 1694343 In DPDK vRouter use-cases (SNAT, LBaaS) that require **netns** to be launched, do not set Jumbo frames. Use MTU <= 1500 bytes.
- 1705795 On an RHOSP10 provisioned cluster, if the vrouter-agent gets restarted, vhost0 interface does not come up. The **service supervisor-vrouter restart** command brings the service back online.
- 1709974 TSN support in RHOSP-based clusters are supported upto RHOSP10. As a workaround, deploy the stack with computes and DPDK first. Then change the **VrouterPhysicalInterface**, add the TSN nodes, and update the stack.
- 1711256 Project isolation is not supported in nested mode. In nested mode, Namespaces-isolation results in a virtual-network creation and doesn't create a new project.
- 1716308 When the head fragment is received in the vRouter, the head fragment is enqueued to the assembler immediately upon arrival. 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. If agent does not write flow action by the time the assembler releases fragments to the flow module, fragments get enqueued in the hold queue. As a maximum of only three fragments are enqueued in the hold queue, rest of the fragments from the assembler get dropped in the flow module. This leads to the whole packet being dropped on the receive side leading to the first packet loss.
- 1718807 In OVSDB case, the routes are exported from the ToR Agent where the SG is appropriately updated so that inter-virtual network traffic doesn't require any explicit SG to be configured for it to pass. When TOR is peering with control node, this SG has to be explicitly configured and this behavior is expected.
- 1720990 With policy-based mirroring with ECMP destinations, one of the destination vRouter drops packet with invalid NH.
- 1721620 VNC API sends an update of all subfields in a field, like virtual_network_properties and not the updated subfields, like allow_transit of vn_properties alone, hence matching against all the subfields due to which you might not be able to update a subfield. As a workaround, perform the following steps:

  1. Delete the attributes which are not updated from the field class. For example, all the attributes of virtual_network_properties except allow_transit.

2. Call vn_obj.set_virtual_network_properties().
3. Use RestApi or Contrail-UI instead of vnc_api.

- 1724357 While provisioning a RHOSP10 cluster with DPDK nodes, the DPDK node power state goes offline during introspection stage. As a workaround, delete the DPDK nodes from Ironic configuration and add them with the right configuration.

  1. **ironic node-delete** ***node-name***
  2. Create a JSON file configuration as shown in the following example for all DPDK nodes:

```
{
  "nodes": [
    {
      "mac": [
        "90:e2:ba:4c:67:3d"
      ],
      "name": "compute3-dpdk",
      "capabilities" : "profile:compute-dpdk",
      "pm_user": "admin",
      "pm_addr": "10.87.122.164",
      "pm_password": "admin",
      "pm_type": "pxe_ipmitool"
    }
  ]
}
```

  3. **openstack baremetal import --json path to .json**
  4. **openstack baremetal introspection bulk start**

- 1728802 Session logging: Incorrect VN information seen for sessions on transparent SI VMIs. As a workaround, ensure that VLAN NH inherits the policy status from its associated interface.
- 1729059 You must use Ansible version 2.3 to install Contrail using **contrail-ansible** for Kubernetes and OpenShift deployments.
- 1735057 When bringing up Contrail cluster on Red Hat container, manually install docker-py on all the target nodes.

  - To install Pip, use the following command:

```
wget https://bootstrap.pypa.io/get-pip.py
python get-pip.py
```

  - To install docker-py, use the following command:

```
pip install docker-py
```

Known Behavior in Contrail Release 4.1.1

- 1681680 When the DPDK vRouter fragments packets before sending them on the wire, the reassembly of the fragments on the receiver might time out in some cases.
- 1694343 In DPDK vRouter use-cases (SNAT, LBaaS) that require **netns** to be launched, do not set Jumbo frames. Use MTU <= 1500 bytes.
- 1705795 On an RHOSP10 provisioned cluster, if the vrouter-agent gets restarted, vhost0 interface does not come up. The **service supervisor-vrouter restart** command brings the service back online.
- 1709974 TSN support in RHOSP-based clusters are supported upto RHOSP10. As a workaround, deploy the stack with computes and DPDK first. Then change the **VrouterPhysicalInterface**, add the TSN nodes, and update the stack.
- 1711256 Project isolation is not supported in nested mode. In nested mode, Namespaces-isolation results in a virtual-network creation and doesn't create a new project.
- 1716297 Provisioning fails for OpenStack HA with SMLite on Ocata.
- 1716308 When the head fragment is received in the vRouter, the head fragment is enqueued to the assembler immediately upon arrival. 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. If agent does not write flow action by the time the assembler releases fragments to the flow module, fragments get enqueued in the hold queue. As a maximum of only three fragments are enqueued in the hold queue, rest of the fragments from the assembler get dropped in the flow module. This leads to the whole packet being dropped on the receive side leading to the first packet loss.
- 1718807 In OVSDB case, the routes are exported from the ToR Agent where the SG is appropriately updated so that inter-virtual network traffic doesn't require any explicit SG to be configured for it to pass. When TOR is peering with control node, this SG has to be explicitly configured and this behavior is expected.
- 1720990 With policy-based mirroring with ECMP destinations, one of the destination vRouter drops packet with invalid NH.
- 1721564 Contrail Ocata: ansible_hostname does not handle "-" in hostname. As a workaround, copy the cert and key files with the expected name without the "-" in the same path as **/etc/contrail_smgr/puppet/ssl/**.
- 1721620 VNC API sends an update of all subfields in a field, like virtual_network_properties and not the updated subfields, like allow_transit of vn_properties alone, hence matching against all the subfields due to which you might not be able to update a subfield. As a workaround, perform the following steps:

  1. Delete the attributes which are not updated from the field class. For example, all the attributes of virtual_network_properties except allow_transit.
  2. Call vn_obj.set_virtual_network_properties().

3. Use RestApi or Contrail-UI instead of vnc_api.

- 1724357 While provisioning a RHOSP10 cluster with DPDK nodes, the DPDK node power state goes offline during introspection stage. As a workaround, delete the DPDK nodes from Ironic configuration and add them with the right configuration.

  1. **ironic node-delete** ***node-name***
  2. Create a JSON file configuration as shown in the following example for all DPDK nodes:

```
{
  "nodes": [
    {
      "mac": [
        "90:e2:ba:4c:67:3d"
      ],
      "name": "compute3-dpdk",
      "capabilities" : "profile:compute-dpdk",
      "pm_user": "admin",
      "pm_addr": "10.87.122.164",
      "pm_password": "admin",
      "pm_type": "pxe_ipmitool"
    }
  ]
}
```

  3. **openstack baremetal import --json path to .json**
  4. **openstack baremetal introspection bulk start**

- 1728802 Session logging: Incorrect VN information seen for sessions on transparent SI VMIs. As a workaround, ensure that VLAN NH inherits the policy status from its associated interface.
- 1729059 You must use Ansible version 2.3 to install Contrail using **contrail-ansible** for Kubernetes and OpenShift deployments.
- 1735057 When bringing up Contrail cluster on Red Hat container, manually install docker-py on all the target nodes.
  - To install Pip, use the following command:

```
wget https://bootstrap.pypa.io/get-pip.py
python get-pip.py
```

  - To install docker-py, use the following command:

```
pip install docker-py
```

Known Behavior in Contrail Release 4.1

- 1735874 Kubernetes: Analytics services fail when 3 node HA setup is brought up using single yaml.
- 1681680 When the DPDK vRouter fragments packets before sending them on the wire, the reassembly of the fragments on the receiver might time out in some cases.
- 1694343 In DPDK vRouter use-cases (SNAT, LBaaS) that require **netns** to be launched, do not set Jumbo frames. Use MTU <= 1500 bytes.
- 1705795 On an RHOSP10 provisioned cluster, if the vrouter-agent gets restarted, vhost0 interface does not come up. The **service supervisor-vrouter restart** command brings the service back online.
- 1709974 TSN support in RHOSP-based clusters are supported upto RHOSP10. As a workaround, deploy the stack with computes and DPDK first. Then change the **VrouterPhysicalInterface**, add the TSN nodes, and update the stack.
- 1711256 Project isolation is not supported in nested mode. In nested mode, Namespaces-isolation results in a virtual-network creation and doesn't create a new project.
- 1716297 Provisioning fails for OpenStack HA with SMLite on Ocata.
- 1716308 When the head fragment is received in the vRouter, the head fragment is enqueued to the assembler immediately upon arrival. 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. If agent does not write flow action by the time the assembler releases fragments to the flow module, fragments get enqueued in the hold queue. As a maximum of only three fragments are enqueued in the hold queue, rest of the fragments from the assembler get dropped in the flow module. This leads to the whole packet being dropped on the receive side leading to the first packet loss.
- 1718807 In OVSDB case, the routes are exported from the ToR Agent where the SG is appropriately updated so that inter-virtual network traffic doesn't require any explicit SG to be configured for it to pass. When TOR is peering with control node, this SG has to be explicitly configured and this behavior is expected.
- 1719430 While upgrading OpenShift cluster from build 31 to 32 in Redhat base OS, contrail kube-manager fails to come up for permission due to an issue with **contrail-kube-manager.log**. As a workaround, change the permission of the log file inside **contrail-kube-manager** docker and restart the service.
- 1720118 Configuration of Allowed Address Pair (AAP) with prefix length less than 24 is not allowed.
- 1720990 With policy-based mirroring with ECMP destinations, one of the destination vRouter drops packet with invalid NH.
- 1721564 Contrail Ocata: ansible_hostname does not handle "-" in hostname. As a workaround, copy the cert and key files with the expected name without the "-" in the same path as **/etc/contrail_smgr/puppet/ssl/**.

- 1721620 VNC API sends an update of all subfields in a field, like virtual_network_properties and not the updated subfields, like allow_transit of vn_properties alone, hence matching against all the subfields due to which you might not be able to update a subfield. As a workaround, perform the following steps:

  1. Delete the attributes which are not updated from the field class. For example, all the attributes of virtual_network_properties except allow_transit.
  2. Call vn_obj.set_virtual_network_properties().
  3. Use RestApi or Contrail-UI instead of vnc_api.

- 1722877 There is no automated provisioning method through director for SRIOV.
- 1724357 While provisioning a RHOSP10 cluster with DPDK nodes, the DPDK node power state goes offline during introspection stage. As a workaround, delete the DPDK nodes from Ironic configuration and add them with the right configuration.

  1. **ironic node-delete** ***node-name***
  2. Create a JSON file configuration as shown in the following example for all DPDK nodes:

```
{
  "nodes": [
    {
      "mac": [
        "90:e2:ba:4c:67:3d"
      ],
      "name": "compute3-dpdk",
      "capabilities" : "profile:compute-dpdk",
      "pm_user": "admin",
      "pm_addr": "10.87.122.164",
      "pm_password": "admin",
      "pm_type": "pxe_ipmitool"
    }
  ]
}
```

  3. **openstack baremetal import --json path to .json**
  4. **openstack baremetal introspection bulk start**

- 1728802 Session logging: Incorrect VN information seen for sessions on transparent SI VMIs. As a workaround, ensure that VLAN NH inherits the policy status from its associated interface.
- 1729059 You must use Ansible version 2.3 to install Contrail using **contrail-ansible** for Kubernetes and OpenShift deployments.
- 1733684 ContrailSecurity: Addressgroup match performs the OR function of both subnet and label rather than the AND function.
- 1734110 Provisioning SSL for metadata fetch does not work if the OpenStack SKU is Ocata.

- 1734790 In RBAC-enabled deployments, SM-Lite doesn't enable RBAC in api-paste.ini for Neutron. As a workaround, after provisioning the cluster, manually update **/etc/neutron/api-paste.ini** as described in RBAC.
- 1735054 RBAC for analytics alarms doesn't work.
- 1735057 When bringing up Contrail cluster on Red Hat container, manually install docker-py on all the target nodes.
  - To install Pip, use the following command:

    ```
    wget https://bootstrap.pypa.io/get-pip.py
    python get-pip.py
    ```

  - To install docker-py, use the following command:

    ```
    pip install docker-py
    ```

- 1735081 While deploying using Contrail Networking package, the dependency for liblua5.3-0 must be resolved explicitly by the user.
- 1735590 In Kubernetes and OpenShift-based deployments when we crate SNAT router and extend cluster-network to that SNAT router host is losing all connectivity.

  As a workaround, if you want to use the SNAT feature in Contrail, disassociate the ip-fabric-cluster-network-default policy and delete it.
- 1773294 In a RHOSP deployment incase of DPDK vRouter crash the vRouter does not restart itself. As a workaround, use the **service supervisor-vrouter restart** command.

Resolved Issues

This section lists limitations that are resolved with this release.

Resolved Issues in Contrail Release 4.1.3

Contrail Release 4.1.3 addresses issues in BFD health check and stability issues in the vRouter.

Resolved Issues in Contrail Release 4.1.2

You can research limitations that are resolved with Contrail Release 4.1.2 in Launchpad at:

https://launchpad.net/juniperopenstack/+milestone/r4.1.2.0 .

Resolved Issues in Contrail Release 4.1.1

You can research limitations that are resolved with Contrail Release 4.1.1 in Launchpad at:

https://launchpad.net/juniperopenstack/+milestone/r4.1.1.0 .

Resolved Issues in Contrail Release 4.1

You can research limitations that are resolved with Contrail Release 4.1 in Launchpad at:

https://launchpad.net/juniperopenstack/+milestone/r4.1.0.0-fcs .

Deploying Contrail Release 4.1 with Netronome SmartNICs by Using Juju

Prerequisites

Make sure that the following requirements are met:

- MaaS Server (if the cloud type is maas)—Install MaaS and Juju on this server
- Juju Controller—Bootstrap the Juju controller. This can be on a VM or a bare metal server.
- A repository—To get Netronome and patched OpenStack packages. The repository can be a virtual machine.
- Contrail controller—Use Ubuntu 16.04 xenial
- OpenStack—OpenStack can be on the same node as contrail-controller. Use the patched packages to support virtio-forwarder
- Compute node with Agilio SmartNIC—Use Ubuntu16.04 xenial

Deploying Contrail Release 4.1 with Netronome SmartNIC

Follow these steps to deploy Contrail Release 4.1 with Netronome SmartNIC:

1. Download the Netronome build package (Agilio vRouter r4.1.2—release notes build 82—(2018/10/25)) from https://support.netronome.com .

2. Download the Contrail package (contrail-networking-docker_4.1.2.0-<build>_xenial.tgz) at https://support.juniper.net/support/downloads/?p=contrail#sw .

   

   NOTE: **You must use Ubuntu 16.04 Xenial for the deployment.**

3. Run the following script to rename the NFP interfaces.

```
Extract Netronome build tar file - >
Netronome_R4.1.2_build_<build_nr>_Juju/Maas-Commissioningscript/
00-maas-01-dagilio
```

Ubuntu-16.04-ga (kernel-4.4) package does not contain NFP driver. So the NFP driver needs to be installed separately. If deployed using MaaS, run the following command in the **/etc/maas/preseeds/curtin_userdata** directory to install the NFP driver.

```
netronome_01_driver_add_target: curtin in-target -- bash -c '[[ "$(lspci -Dnnd
 19ee:4000)" ]] && aptget install --quiet --assume-yes agilio-nfp-driver-dkms
 agilio-nic-firmware || true'
```

The sample **bundle.yaml** file for the deployment is available in the Netronome .tar file in the following path:

Netronome_R4.1.2_build_26_Juju/agilio-vrouter/
example-docker-ocata-agilio-bundle.yaml

4. Change the Kernel options setting for enabling Huge Pages as shown below:

```
kernel_opts='intel_iommu=on iommu=pt default_hugepagesz=2M hugepagesz=2M
hugepages=<nr_hugepages>'
```

5. Enable hardware acceleration by installing the following Agilio packages.

```
agilio-nfp-driver-dkms, agilio-nic-firmware
```

6. Bootstrap Juju controller.

7. Run the following command to clone Contrail Charm from Github.

```
git clone https://github.com/Juniper/contrail-charms
```

8. Agilio-vrouter uses **virtio-forwarder** as **vnic-type** to enable hardware acceleration. You need to patch OpenStack packages to use **virtio-forwarder** as **vnic-type**. For this, you need to specify the location of the patched OpenStack packages in the **bundle.yaml** file.

   The following are the patched OpenStack Charm packages to be picked from the local directory.

```
charm: ./charm-nova-cloud-controller, charm: ./charm-neutron-api
```

```
charm: ./charm-nova-compute, charm: ./agilio-vrouter
```

You can find the patches in https://drive.google.com/drive/folders/1XtFA3rEkYeXl13xAqTrwHD6jO8JcIBWI .

Use the following script provided in the Netronome build tar file:

```
Extract Netronome_R4.1_build_<build_nr>.tar-> debs/package_builder ->
./create-openstackpackages.sh
```

9. Deploy the cluster using the following command.

```
juju deploy ./bundle.yaml
```

10. Run the following commands to attach the Contrail packages to Juju.

```
juju attach contrail-controller contrail-controller=contrail-controller.tar.gz
juju attach contrail-analytics contrail-analytics=contrail-analytics.tar.gz
juju attach contrail-analyticsdb contrail-analyticsdb.tar.gz
```

11. Use the following commands to view the status or to log into any of the contrail/openstack units after deployment:

```
openstack units after the deployment:
juju status
```

12. Identify the name of the unit or IP address from the **juju status** output:

```
ssh into contrail-controller:
Juju ssh accel-nova-compute/0 ß0 is unit number
```

You can find the logs in the **/var/log/juju** directory of the unit.

If logs do not indicate failure of any of the above steps, and all the units are deployed and are ready for use, make sure that **contrail-api** is reachable from **neutron-server**. If control data subnet is not same as the subnet used for the provisioning, you may need to configure a static route.

Launching a VM after Deploying Contrail Release 4.1 with Netronome SmartNIC

Follow these steps to launch a VM after deployment.

1. Create a flavour with the following metadata for spawning hardware accelerated VMs.

```
--property hw:mem_page_size = 2048 or --property hw:mem_page_size = 1048576
```

In this release, spawning a VM without creating a port is not supported. You need to create a port with **virtio-forwarder** as **vnic-type**. Then you need to use this port to launch a hardware accelerated VM.

```
vnic_type virtio-forwarder
```

2. Use OpenStack dashboard or nova boot to spawn the VM.

Caveats

This section lists the known issues with Netronome SmartNIC deployment.

- **nfp_p1** interface is not recommended to be used because of the following issue:

  https://bugs.launchpad.net/juniperopenstack/trunk/+bug/1797358

- Launching a VM directly using net-ids is not supported. You need to use port-ids. See https://bugs.launchpad.net/juniperopenstack/r4.1/+bug/1796812 for more information.

Upgrading Contrail 4.0 to 4.1

This section provides the process for upgrading an existing Contrail Release 4.0 system to Contrail Release 4.1.

- Upgrade Assumptions on page 23
- Upgrade Procedure on page 23

Upgrade Assumptions

This upgrade procedure assumes the following.

- The initial cluster (4.0.x) was provisioned using Server Manager.
- The OpenStack SKU is the same in the “from” and “to” versions.
- A backup has been made of the analytics database, see *Backing Up Contrail Databases Using JSON Format*.

Upgrade Procedure

1. Make a backup of the analytics database, because the upgrade procedure removes the analytics database information, see *Backing Up Contrail Databases Using JSON Format*.

2. Add the new Contrail 4.1 Debian image to the Server Manager JSON used for provisioning.

   server-manager add image –f contrail_image.json

3. Upgrade the cluster by reprovisioning the cluster with the new image.

   - For an all-in-one, single-node demo system:

     server-manager provision—-cluster_id <all_in_one_cluster> combined_image_mainline

   - For a multinode system:

server-manager provision —-cluster_id <multi_node> combined_image_mainline

4. Monitor progress of the provisioning by observing cluster status or log entries.
   - Cluster status: **server-manager display server --cluster_id <cluster_id> --select "id,ip_address,roles,status"**
   - Log entries: /var/log/contrail-server-manager/debug.log

NOTE: **Log entries from the previous version are lost in the upgrade process.**

For more upgrade instructions, see:

- Upgrade Procedure for RHOSP-based Contrail 4.1.2 to Contrail 4.1.3
- Upgrade Procedure for RHOSP-based Contrail 4.1.1 to Contrail 4.1.2
- Upgrade Procedure for RHOSP-based Contrail 3.2.x to Contrail 4.1

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the Juniper Networks TechLibrary site, and do one of the following:

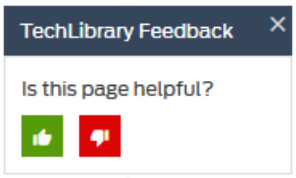

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- E-mail—Send your comments to **techpubs-comments@juniper.net**. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: https://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: https://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at https://www.juniper.net/cm/.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see https://www.juniper.net/support/requesting-support.html.

Revision History

March 2019—Revision 5, Contrail 4.1.3

January 2019—Revision 4, Contrail 4.1.2

November 2018—Revision 3, Contrail 4.1.2

June 2018—Revision 2, Contrail 4.1.1

November 2017—Revision 1, Contrail 4.1