

Contrail Release 4.0 Release Notes

Release 4.0
June 2017

Contents

Introduction	2
New and Changed Features	2
Alarms History	2
Allowed Address Pair	2
Containerized Contrail	3
Contrail Integration with Kubernetes	3
Mapping VLAN Tags from Physical NIC to a VMI (NIC-Assisted mirroring)	4
TCP Segment Offload (TSO) for DPDK vRouter	4
Platform Support	4
OpenStack Application Support	4
Supported Platforms Contrail 4.0	4
Known Behavior	5
Resolved Issues	8
Deprecated Items	8
Contrail Discovery Service and IF-MAP	8
Configuration with Testbed.py	9
Fabric Commands	9
Upgrading Contrail 3.2 to 4.0	9
Overview of Changes in Contrail Release 4.0	9
Assumptions	10
Using a Server Manager Script to Upgrade from Contrail 3.2 to 4.0	10
Running the Upgrade Script	10
Steps for Upgrading Contrail 3.2 to 4.0 (Without using Server-Manager for upgrade)	10
Documentation Feedback	12
Requesting Technical Support	13
Self-Help Online Tools and Resources	13
Opening a Case with JTAC	13
Revision History	14

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 4.0 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 http://www.juniper.net/techpubs/en_US/contrail4.0/information-products/topic-collections/release-notes/index.html.

New and Changed Features

The features listed in this section are new or changed as of Contrail Release 4.0. A brief description of each new feature is included.

- [Alarms History on page 2](#)
- [Allowed Address Pair on page 2](#)
- [Containerized Contrail on page 3](#)
- [Contrail Integration with Kubernetes on page 3](#)
- [Mapping VLAN Tags from Physical NIC to a VMI \(NIC-Assisted mirroring\) on page 4](#)
- [TCP Segment Offload \(TSO\) for DPDK vRouter on page 4](#)
- [Platform Support on page 4](#)
- [OpenStack Application Support on page 4](#)

Alarms History

New fields in the Contrail Web user interface at **Monitor > Alarms > Dashboard** display alarms history, including alarms that were set or reset. You can also use a **contrail-status** query to view the alarms history. The **contrail-status** also displays a history of added, updated, and removed information for UVEs in Contrail.

See [Alarms History](#).

Allowed Address Pair

An allowed address pair extension is an OpenStack feature supported by Contrail. By default, there is no way to specify additional MAC and IP address pairs that are allowed to pass through a port in Neutron, because ports are locked down to their MAC address and the fixed IPs associated with their port for anti-spoofing reasons. This locking can sometimes prevent protocols such as VRRP from providing a high availability failover strategy. Using the allowed address pair extension enables additional MAC and IP address pairs to be allowed through ports in Neutron.

For more information, see [Service Chain Version 2 with Port Tuple](#).

Containerized Contrail

Starting with Contrail Release 4.0, some of the Contrail subsystems are delivered as Docker containers that group together related functional components.

The default containerized components include:

- contrail-controller
- contrail-analytics
- contrail-analyticsdb
- contrail-lb (loadbalancer)

Key features of the new architecture of Contrail containers:

- All of the containers are multiprocess Docker containers.
- Each container has an INI-based file that holds the configurations for all of the applications running in that container.
- Each container is self-contained, with minimal external orchestration needs.
- The user toolset `contrailctl` is used to manage the container configuration files.

To install containerized Contrail, all users need to use Server Manager or Server Manager Lite, even if they will not be using Server Manager for other provisioning functions.

For more information, see:

- [Introduction to Containerized Contrail Modules](#)
- [Installing Containerized Contrail Clusters Using Server Manager](#)
- [Installing Containerized Contrail Using Server Manager Lite \(SM-Lite\)](#)
- [Configuring Services Within Contrail Containers](#)
- [High Availability for Containerized Contrail](#)

Contrail Integration with Kubernetes

Contrail Release 4.0 supports a Container Network Interface (CNI) for integrating Contrail with the Kubernetes automation platform. Kubernetes, also called K8s, is an open-source platform for automating deployment, scaling, and operations of application containers across clusters of hosts, providing container-centric infrastructure. It provides a portable platform across public and private clouds. Kubernetes supports a pluggable framework called Container Network Interface (CNI) for most of the basic network connectivity. Contrail Release 4.0 provides support for CNI for Kubernetes.

For more information, see [Contrail Integration with Kubernetes](#).

Mapping VLAN Tags from Physical NIC to a VMI (NIC-Assisted mirroring)

Contrail Release 4.0 has the ability to mirror specific traffic to a traffic analyzer or to a physical probe using the Network interface card (NIC) instead of the vRouter to mirror packets. When NIC-assisted mirroring is enabled, ingress packets to be mirrored sent from a VM are routed to the NIC with a configured VLAN tag. The NIC is configured for VLAN port-mirroring and mirrors any packet with the VLAN tag.

See [Mapping VLAN Tags from a Physical NIC to a VMI \(NIC-Assisted Mirroring\)](#).

TCP Segment Offload (TSO) for DPDK vRouter

The ability to provide TCP segment offload (TSO) has been added to DPDK vRouters. The TSO feature is already available on non-DPDK routers. The TSO feature helps provide increased performance support for TCP-based large data applications, such as Web servers. The TSO feature enables segmentation and reassembly of large TCP segments to occur in lower layers of the stack, thereby boosting performance.

Platform Support

Contrail Release 4.0 includes the following platform support:

- Kubernetes 1.6 on Ubuntu 16.04.2
- Kubernetes 1.6 on CentOS 7.2
- Openstack Mitaka on Ubuntu 14.04.5
- Openstack Newton on Ubuntu 16.04.2

OpenStack Application Support

Contrail Release 4.0 supports integration with Fuel. Fuel is an open source deployment and management tool for OpenStack to streamline and accelerate the process of deploying, testing, and maintaining various configurations of OpenStack at scale.

Supported Platforms Contrail 4.0

[Table 1 on page 5](#) lists the operating system versions and the corresponding Linux or Ubuntu kernel versions supported by Contrail Release 4.0.

Table 1: Supported Platforms

Contrail Release	OpenStack Release	Operating System and Kernel Versions
Contrail Release 4.0	OpenStack Newton	<ul style="list-style-type: none"> • Ubuntu 16.04.2—Linux kernel version 4.4-62 • Red Hat 7.3—Linux kernel version 3.10.0-514.6.2 • VMware vCenter 5.5, 6.0—Ubuntu 14.04.4 kernel version 3.13.0-106-generic
	OpenStack Mitaka	<ul style="list-style-type: none"> • Ubuntu 14.04.5—Linux kernel versions 3.13.0-106-generic and 4.4.0-34-generic • VMware vCenter 5.5, 6.0—Ubuntu 16.04.2 kernel version 4.4.0-62-generic



NOTE: In Contrail Release 4.0 and later, if the stock kernel version of your Ubuntu system is older 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. Bug numbers are listed and can be researched in [Launchpad.net](https://bugs.launchpad.net/juniperopenstack) at <https://bugs.launchpad.net/juniperopenstack>.

- 1623695 In case of RBAC enabled clusters, user should create network-ipam in their own tenant configuration instead of using the default network-ipam for which the user doesn't have permissions.
- 1624148 In case of RBAC enabled clusters, service instance automatically created by the system on behalf of a user will not be visible in the UI.
- 1650420 In case of RBAC enabled clusters, objects created through LBaaS plugin are created with Neutron ownership.
- 1675224 VMs in non-HA cluster in SHUTOFF are in shutdown state after upgrade. As a workaround, set `resume_guests_state_on_host_boot = True` in the `nova.conf` file of the compute node, for the guest VMs to be resumed.
- 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.

- 1689761 SSL communication to Keystone service cannot be configured using Server Manager
- 1690108 In case of Contrail HA deployment, the lb container must not be installed on same host with other contrail containers.
- 1690904 TLS configured LBaaS needs Barbican service on OpenStack node and Barbican client on compute nodes. As a workaround , perform the following steps:
 1. Install the Barbican service on OpenStack node.
 2. Install the Barbican client on all compute nodes.
 3. Configure the **/etc/contrail/contrail-lbaas-auth.conf** file with the required authentication information on all compute nodes. For example, following is a sample content of the **/etc/contrail/contrail-lbaas-auth.conf** file:

```
[BARBICAN]
admin_tenant_name=service
admin_user=neutron
admin_password=<neutron-service-passwd>
auth_url=http://<keystone-service-ip>:<keystone-service-port>/v2.0
region=RegionOne
```

- 1691245 Kernel crashes when enabling LBaaS on Linux kernel version 4.4 and DPDK Hypervisor.
- 1691862 Server Manager requires internet access for installation of certain packages such as Cobbler, Puppet, DHCP, bind, and tftp.
- 1692067 Requests to Contrail Web-UI are not load balanced through contrail-lb. To access Web-UI, you must connect using the system IP on which the controller Docker runs.
- 1693590 Cinder volume creation fails after provisioning. As a workaround, restart the cinder-volume service manually on all OpenStack nodes.
- 1694343 In DPDK vRouter use-cases (SNAT, LBaaS) that require netns to be launched, do not set Jumbo frames. Use MTU <= 1500 bytes.
- 1694519 On a 16.04.2 Newton cluster, even in case of non-DPDK vRouter, the "contrail-status" will show status for "contrail-vrouter-dpdk" and during a restart of vRouter service, a core might appear. This does not have any functionality impact. As a workaround, to avoid seeing the core, remove the **/lib/systemd/system/contrail-vrouter-dpdk.service** file from the non-DPDK compute nodes.
- 1694849 On Ubuntu 16.04-based Docker containers sometimes the ProcessStatus alarm is not activated even if a process is not running. As a workaround, restart the node manager process based on the process that is not running on the respective container. For example, for the controller Docker, either the contrail-control-nodemgr or contrail-config-nodemgr process must be restarted. On the analytics Docker, contrail-analytics-nodemgr process needs to be restarted.
- 1694851 Provisioning of non-HA cluster using Server-Manager UI does not work in Contrail Release 4.0. However, CLI and JSON-based provisioning work for all use-cases.

- 1694857 After upgrade of cluster running Contrail Release 3.2 to 4.0, sometimes the Neutron server process does not start automatically. As a workaround, start the Neutron server explicitly after upgrade.
- 1695544 LBaaS cannot be configured from Horizon UI. Use CLIs instead of UI.
- 1695559, 1695503 Usage of uppercase characters in variable values in JSONs might stall installation. Some of the known variables where upper case characters cannot be used are server host names and image id.
- 1695566 CEPH storage integration with Ubuntu 16.04.2 Newton is not supported. Integration with Ubuntu 14.04.5 Mitaka is supported.
- 1695584 Server Manager provisioning-support for SSL communication for Keystone service is not available.
- 1695662 In Ubuntu 16.04.2 Newton cluster, the rabbitmq issue described in [Rabbitmq boot failure with "tables_not_present"](#) might be seen. As a workaround, perform the following steps:
 1. Stop the rabbitmq service on the node which is not clustered .


```
service stop and epmd -kill
```

In a three node setup, mostly two of nodes are clustered together and the third one is single. It's also possible that all the rabbitmq nodes are single too.
 2. Remove the `/var/lib/rabbitmq/mnesia` directory.
 3. Re-start rabbitmq on the node.
 4. Check Nova service-list output and restart all the services which are marked down.
- 1695741, 1694368 Ceilometer provisioning using Server-Manager will not work. If the user enables provisioning knob to enable ceilometer, sometimes provisioning will not get completed.
- 1695770 SSL communication to API-server provisioned using Server Manager will not work.
- 1695792 For use-cases where VMs on a compute node bypass vRouter through SRIOV, the nova.conf should have options to exclude interface used by vVouter. The Contrail Release 4.0 Ansible-based provisioning code does not add this. You need to explicitly set it as follows:
 1. Edit nova.conf in all OpenStack nodes and add the following lines.


```
scheduler_available_filters = nova.scheduler.filters.all_filters
scheduler_default_filters = RetryFilter, AvailabilityZoneFilter, RamFilter,
ComputeFilter, ComputeCapabilitiesFilter, ImagePropertiesFilter,
PciPassthroughFilter
```
 2. Restart Nova scheduler service.
- 1695842 Server Manager provisioning support for vcenter-only and vcenter-as-compute use-cases is not available. Use fab commands to install vcenter-only and vcenter-as-compute use-cases.

- 1695845 In Contrail Release 4.0, on a DPDK compute, the VMs cannot bypass vrouter through SRIOV. This use-case works in non-DPDK computes.
- 1695856 Server Manager Ansible-based provisioning support for Keystone V3 for containerized deployment is not available.
- 1695859 Server Manager does not have Ansible-based provisioning support for secure communication for rabbitmq service.
- 1695938 For use-cases that requires VMs in computes bypass vRouter using SRIOV, nova.conf should be configured with a white list of VFs that Nova can use for attaching VMs. The Contrail Release 4.0 Ansible-based provisioning code does not add this. You need to explicitly update the nova.conf similar to the following example and restart the nova-compute service on the compute nodes:

```
pci_passthrough_whitelist = { "address": "0000:81:10.0", "physical_network":  
    "physnet1" }  
pci_passthrough_whitelist = { "address": "0000:81:10.2", "physical_network":  
    "physnet1" }  
pci_passthrough_whitelist = { "address": "0000:81:10.4", "physical_network":  
    "physnet1" }  
pci_passthrough_whitelist = { "address": "0000:81:10.6", "physical_network":  
    "physnet1" }  
pci_passthrough_whitelist = { "address": "0000:81:11.0", "physical_network":  
    "physnet1" }  
pci_passthrough_whitelist = { "address": "0000:81:11.2", "physical_network":  
    "physnet1" }  
pci_passthrough_whitelist={"devname":"p1p2", "physical_network":"physnet2"}
```

Resolved Issues

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

<https://launchpad.net/juniperopenstack/+milestone/r4.0.0.0-fcs>.

Deprecated Items

The following features have been deprecated in Contrail Release 4.0.

Contrail Discovery Service and IF-MAP

Starting with Contrail Release 4.0, the existing centralized Contrail discovery service is replaced with a distributed method of allocating service resources. Any discovery APIs will cease to work. The replacement is a distributed resource allocation list of service nodes, maintained in each module of the system.

Additionally, information previously managed by IF-MAP has been changed to items managed in CONFIGDB.

For more information, see [Distributed Service Resource Allocation with Containerized Contrail](#).

Configuration with Testbed.py

In previous versions of Contrail, a **testbed.py** was used to initiate configurations at installation. With Contrail Release 4.0, the use of **testbed.py** is limited to only those installations that are managed by SM-Lite. Most users will edit JSON files for the initial configuration at installation. Installation using JSON files can be performed with Server Manager or SM-Lite.

Fabric Commands

Contrail Release 4.0 has limited support for Fabric commands. Installations that were previously accomplished with Fabric commands are now accomplished by means of Server Manager or SM-Lite deployments.

Upgrading Contrail 3.2 to 4.0

Contrail Release 4.0 presents a number of differences from earlier versions of Contrail, the most notable is that many Contrail services are now run within containers.

This section provides the process for upgrading an existing Contrail Release 3.2 system to Contrail Release 4.0. You can perform an upgrade by using a convenient script, or by performing the upgrade manually. Both procedures are included in this topic.

- [Overview of Changes in Contrail Release 4.0 on page 9](#)
- [Using a Server Manager Script to Upgrade from Contrail 3.2 to 4.0 on page 10](#)
- [Steps for Upgrading Contrail 3.2 to 4.0 \(Without using Server-Manager for upgrade\) on page 10](#)

Overview of Changes in Contrail Release 4.0

For previous releases of Contrail, through Contrail Release 3.2.x, Contrail could be provisioned by using fab commands or by using the Contrail Server Manager. Starting with Contrail Release 4.0, many Contrail services have been containerized, and provisioning can only be accomplished by using the Contrail Server Manager.

Additionally, the use of fab commands is no longer supported in Contrail Release 4.0.

You can perform the upgrade by using the steps presented here, or by using a script available in Server Manager that automates the steps.

Significant differences between Contrail 3.2 and 4.0 include:

- A number of Contrail services are run in containers. Container default names in Contrail Release 4.0 include:
 - contrail-controller
 - contrail-analytics
 - contrail-analyticsdb
 - contrail-lb (load-balancer)

For more information about Contrail containers, see *Introduction to Containerized Contrail Modules*.

Assumptions

The upgrade procedure assumes that RabbitMQ server and Neutron server are running on the OpenStack node by default.

Using a Server Manager Script to Upgrade from Contrail 3.2 to 4.0

A script is available in Server Manager that runs the steps provided in *Steps for Upgrading Contrail 3.2 to 4.0*.

The script location is: `/opt/contrail/server_manager/inplace_upgrade.py`.

Information regarding using the script:

- The script uses the **openstack-config** utility, you can install the `contrail-setup` package to get the utility.
- The timeout for cluster provisioning completion is 3600 seconds, you can increase the timeout for a cluster with a large number of nodes.

The default timeout parameter in the script:

PROV_TIMEOUT = 3600

- The script execution log is saved at `/var/log/contrail/inplace_upgrade.log`.

Running the Upgrade Script

1. Add the image, cluster, and servers to Server Manager running Contrail Release 4.0 code.
2. Run the script.

```
user@node:/opt/contrail/server_manager# python inplace_upgrade.py -h
usage: inplace_upgrade.py [-h] cluster_name image_name

positional arguments:
  cluster_name  cluster to be upgraded.The cluster and servers should be
                existing in SM
  image_name    version to upgrade to

optional arguments:
  -h, --help    show this help message and exit
```

Steps for Upgrading Contrail 3.2 to 4.0 (Without using Server-Manager for upgrade)

This procedure lists the steps needed for upgrading Contrail 3.2 to 4.0.

This procedure is an in-service upgrade from a running 3.2 system that will temporarily run in parallel with the new 4.0 system.

The services and roles referred to in this procedure could be running in separate nodes. Run the procedure commands in the nodes that are running the relevant roles or services.

1. Stop the listed services.

- supervisor-analytics
- supervisor-support-service
- supervisor-database
- contrail-database
- supervisor-webui
- supervisor-config
- supervisor-control
- haproxy (all nodes)
- redis-server (analytics/webui node)
- memcached (config nodes)
- neutron-server (config nodes)
- zookeeper (config nodes)

Example stop commands include:

```
service supervisor-analytics stop
```

```
service supervisor-support-service stop
```

2. Stop the epmd process. The epmd process is a RabbitMQ service that is usually running on a config node.

3. Use a **status** command to verify that none of the services that were stopped in Step 1 are still running.

Example status commands include:

```
service supervisor-analytics status
```

```
service supervisor-support-service stop status
```

4. Prepare the cluster.json and server.json, as required for Server Manager provisioning for Contrail Release 4.0, see *Installing Containerized Contrail Clusters Using Server Manager*.

5. Provision the new 4.0 cluster and wait for provisioning to be completed.

6. Check the status of provisioning.

```
server-manager status server --cluster_id <cluster-id>
```

7. To prevent port conflict with Cassandra between the two versions of Contrail, on the 3.2 config node(s), edit the Cassandra configuration file at **/etc/cassandra/cassandra.yaml** to change the listening port from 9160 to 29160.

When finished, start the Cassandra service.

```
rpc_port: 29160
```

```
service cassandra start
```

8. Identify the new (4.0) and old (3.2) IP address lists for Cassandra.
 - a. Connect to the controller container.
 - b. Edit **/usr/lib/python2.7/dist-packages/contrail_issu/issu_contrail_config.py** and add the old Cassandra list and the new Cassandra list:

```
'old_cassandra_address_list': '192.xxx.xxx.102:29160',
```

```
'new_cassandra_address_list': '10.xx.5.xxx:9161',
```

It is sufficient to provide a single ip:port of the Contrail 3.2 Cassandra node.

9. Run a Cassandra sync between the old and new config nodes, to copy the cassandra config keyspaces from the release 3.2 nodes to the Contrail 4.0 Cassandra database.

```
contrail-issu-pre-sync
```

Wait for the command to complete.

10. Shut down the Cassandra service in the 3.2 controller.
11. The upgrade of the compute nodes includes a reboot of the compute nodes. The reboot turns off the guest virtual machine instances. To restart the virtual machine instances, log in to the Openstack node and restart the VMs:

```
nova start <instance-uuid>
```

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 rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/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 <http://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: <http://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <http://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://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 <http://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 <http://www.juniper.net/support/requesting-support.html>.

Revision History

June 2017—Revision 2, Contrail 4.0

June 2017—Revision 1, Contrail 4.0

Copyright © 2017 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.