

# Release Notes for Contrail Release 2.21

**Release 2.21**  
**October 2015**

## **Contents**

Introduction . . . . .	2
New and Changed Features . . . . .	2
Discovery Clients Honor Discovery Server Response . . . . .	2
Discovery Client Library Provides Publish Reevaluation . . . . .	3
Layer 3 Only Forwarding Mode . . . . .	3
Improved ISSU Support in QFX Series TOR Switches . . . . .	4
Supported Platforms . . . . .	4
Known Behavior . . . . .	4
Known Issues . . . . .	5
Upgrading Contrail Software from Release 2.00 or Greater to Release 2.20 . . . . .	11
Documentation Updates . . . . .	14
Documentation Feedback . . . . .	14
Requesting Technical Support . . . . .	14
Self-Help Online Tools and Resources . . . . .	14
Opening a Case with JTAC . . . . .	15
Revision History . . . . .	15

## 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 2.21 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/contrail2.20/information-products/topic-collections/release-notes/index.html](http://www.juniper.net/techpubs/en_US/contrail2.20/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 2.20. A brief description of each new feature is included.

- [Discovery Clients Honor Discovery Server Response on page 2](#)
- [Discovery Client Library Provides Publish Reevaluation on page 3](#)
- [Layer 3 Only Forwarding Mode on page 3](#)
- [Improved ISSU Support in QFX Series TOR Switches on page 4](#)
- [Supported Platforms on page 4](#)

### Discovery Clients Honor Discovery Server Response

Once the load-balancer is triggered, the discovery server responds with a new set of servers in response to the subscriber. Clients must honor the new set of services, irrespective of the current connection being up or down.

The following are rules to honoring the new set of published services:

- For active-active connections, if both services are disrupted, the client applies only one service at a time.
- If the discovery server responds with a smaller subset of services, the stale services are not cleaned up.
- If the discovery server responds with a NULL list, the clients continue to function with the older list.
- Clients now support resubscribing if the connections are detected as being DOWN. For faster convergence, the client waits at least 3 intervals (heart-beat time (=15secs)) to trigger a resubscribe request .

The introspect command is provided to view the discovery list and the current active connections.



**NOTE:** Discovery servers ensure only one set of services are changed at any given time, but the client also needs to ensure it handles the new set gracefully.



**NOTE:** The above feature applies to subscribers that maintain active-active connections.

## Discovery Client Library Provides Publish Reevaluation

The discovery publish reevaluation enhancement provides the following:

- The client library provides the ability to publish a service with the oper-state up or down.
- The application must register a callback so the service can be reevaluated and the oper-state can be updated accordingly.
- Reevaluation of a published service is triggered every heartbeat interval.
- The discovery server must honor the operational state of published services and provide a new set of services in response to the next subscribe request.

Two new states are introduced, the admin-state and the oper-state which are displayed in the [discovery-server:5998](#) URL.

Clients can change the oper-state based on the internal state of services which enables the discovery server to deallocate or reallocate resources.

The admin-state can only be changes using the python CLI to pull services permanently DOWN.

## Layer 3 Only Forwarding Mode

Contrail Release 2.21 and later supports the integrated routing and bridging (IRB) model in which the traffic flowing within the same subnet is bridged while the traffic across subnets and virtual networks is routed. This mode is referred to as Layer 2 plus Layer 3 (**Layer2 + Layer3**) and it is the default mode in which a cluster operates and a virtual network is created. In Layer 2 plus Layer 3 mode, ARP requests are flooded. The Contrail solution minimizes ARP flooding by having the vRouter proxy for ARP where possible.

In some cases it might not be desirable to have this type of flooding. It might also be desirable to run the cluster or virtual network in the same mode used in previous releases. To support this, Contrail Release 2.21 also supports a **Layer 3 only** mode and a **Layer 2 only** mode.

When Layer 3-only mode is used:

- The IP routes learned from an IP VPN are used to route the traffic irrespective of whether the hosts are in the same subnet.

- Traffic coming from the fabric that uses MPLS over GRE (MPLSoGRE) or MPLS over UDP (MPLSoUDP) encapsulation, is routed at Layer 3.
- An EVPN IP address to MAC address binding is not used. Instead, any packet with a destination MAC address that is not the VRRP MAC address and is not an IP routable packet, is bridged.

The forwarding mode can be configured at the global level to set the default or it can be set at the virtual network level to override the default setting and set the mode for a specified virtual network. To configure the global forwarding mode in the Contrail Controller, select **Configure -> Global Config -> Forwarding Mode**. To configure the forwarding mode of a virtual network in the Contrail Controller, select **Configure -> Networking -> Create >Network -> Advanced Options -> Forwarding Mode**.

## Improved ISSU Support in QFX Series TOR Switches

A Contrail TOR Agent retains the MAC addresses learned from a TOR switch for five minutes after the OVSDB protocol connection between the TOR switch and TOR Agent goes down. This avoids churn in route distribution for short connection drops and helps in supporting ISSU on a QFX Series switch. For more information about the improved ISSU support, see the Release Notes: Junos<sup>®</sup> OS Release 14.1X53-D30 for QFX Switches.

## Supported Platforms

Contrail Release 2.21, is supported on the OpenStack Juno and Icehouse releases. Juno is supported on Ubuntu 14.04.2 and Centos 7.1.

Contrail networking is supported on Red Hat RHOSP 5.0, which is supported only on OpenStack Icehouse

In Contrail Release 2.21, support for VMware vCenter 5.5. vCenter is limited to Ubuntu 14.04.2 (Linux kernel version: 3.13.0-40-generic).

Other supported platforms include:

- CentOS 6.5 (Linux kernel version: 2.6.32-358.el6.x86\_64)
- CentOS 7.1 (Linux kernel version: 3.10.0-229.el7)
- Redhat 7/RHOSP 5.0 (Linux kernel version: 3.10.0-299.el7.x86\_64)
- Ubuntu 12.04.04 (Linux kernel version: 3.13.0-34-generic)
- Ubuntu 14.04. (Linux kernel version: 3.13.0-40-generic)

## Known Behavior

---

The following are known behaviors in this release of Contrail.

- DNS record updates from a controller DNS server to a named server might be missing even after repeated retries. this is because there is no infrastructure currently to sync records across named servers.

- DNS queries from an agent are now sent to both named servers that were learnt using discovery. There is a very low probability of records missing on both named servers.

The first good response from either of the named servers is used to update the DNS client that sent the DNS query request. If there is no good response, the last bad response is sent to the DNS client to inform the client of the error.

Use the following to display a list of named servers to which queries are sent:

**http://x.x.x.x:8085/Snh\_DnsInfoo**

Use the following to trace the queries sent and responses:

**http://x.x.x.x:8085/Snh\_SandeshTraceRequest?x=DnsBind**

## Known Issues

This section lists known limitations with this release. Bug numbers are listed and can be researched in [Launchpad.net](https://bugs.launchpad.net/) at <https://bugs.launchpad.net/>.

Storage:

- 1497047 In Contrail Release 2.20 and earlier, if a Cassandra node is offline for one minute or longer and then brought back online, it might corrupt the database.

In Contrail Release 2.21 and later, a Cassandra node can be offline for up to three hours and then brought back online without corrupting the database.

If the Cassandra node is offline for more than three hours, you need to perform the following procedure:

- After the Cassandra node joins the Cassandra cluster, you must use the **nodetool repair** command
- If the Cassandra node is offline for more than ten days, it should not be brought back online. Instead, you need to remove the Cassandra node using the **nodetool removemode** command and the associated procedure. The procedure can be access at: [http://docs.datastax.com/en/cassandra/1.2/cassandra/operations/ops\\_remove\\_node\\_t.html](http://docs.datastax.com/en/cassandra/1.2/cassandra/operations/ops_remove_node_t.html)
- After the procedure is complete, you can add the node back as a new node.

Contrail Networking:

- 1501951 Upgrading from Contrail Release 2.20 Build 64 to Contrail Release 2.21 build 102 fails due to a Rabbitmq cluster crash. To correct the problem, you must reprovision the whole cluster.
- 1499169 To use the Contrail Release 2.21 package, you must upgrade Server Manager from Release 2.20 to Release 2.21. See the Server Manager upgrade procedure for the Release 2.21 upgrade.
- 1499546 A configuration flag has been added that controls the ability to send flow samples from a vRouter to a collector. When the **disable\_flow\_collection** flag is set, the vRouter does not send flow samples to the collector.

- 1484297 The Cassandra `gc_grace_seconds` value needs to be changed for keyspaces that do deletes to the default value of 10 days to handle node down scenarios and to prevent deleted data from showing up since the tombstones are stored for only the `gc_grace_seconds` time. The `nodetool repair` command needs to be run on the Cassandra nodes periodically to prevent these issues.
- 1496606 You use the `fab install_new_contrail` and `fab join_cluster` commands to add a new control node to a cluster that is already provisioned.

The `fab join_cluster` command succeeds only if the newly added control node is “up” in the `rabbitmqctl cluster_status` command output. Also before purging an existing control node, you need to verify if the control node is displayed in the `rabbitmqctl cluster_status` command output.

For example:

```
root@a12c4s2:/opt/contrail/utils# rabbitmqctl cluster_status
Cluster status of node 'rabbit@a12c4s2-ctrl' ...
[{nodes,[{disc,['rabbit@a12c3s3-ctrl','rabbit@a12c3s4-ctrl',
               'rabbit@a12c4s2-ctrl']}]},
 {running_nodes,['rabbit@a12c3s4-ctrl','rabbit@a12c3s3-ctrl',
                 'rabbit@a12c4s2-ctrl']}],
 {cluster_name,<<"rabbit@a12c3s3">>},
 {partitions,[]}]

root@a12c4s2:/opt/contrail/utils#
mysql -uroot -p$(cat /etc/contrail/mysql.token) -e "show status like 'wsrep%'"
| wsrep_cert_index_size | 41 |
| wsrep_causal_reads | 145146 |
| wsrep_incoming_addresses | 5.5.5.5:3306,5.5.5.6:3306,5.5.5.4:3306 |
| wsrep_cluster_conf_id | 60 |
| wsrep_cluster_size | 3 |
| wsrep_cluster_state_uuid | 3c0286
```

Verify that the hostname of the new control node is listed in the `rabbitmqctl cluster_status` command output and the IP address of the new control node is listed in the `wsrep_incoming_addresses` field.

- 1496605 When adding a new control node using the `fab install_new_contrail` command, the command expects the new control node to be added to the end of each role definition in the `testbed.py` file.

For example, in the following `testbed.py` file example, `host2` is the newly added control node.

```
# Role definition of the hosts.
env.roledefs = {
  'all': [ host3, host4, host5, host1, host2 ],
  'cfgm': [ host3, host5, host1, host2 ],
  'openstack': [ host3, host5, host1, host2 ],
  'control': [ host3, host5, host1, host2 ],
  'compute': [ host4 ],
  'collector': [ host3, host5, host1, host2 ],
  'webui': [ host3, host5, host1, host2 ],
  'database': [ host3, host5, host1, host2 ],
```

```
'build': [host_build],
}
```

This constraint might be removed in a future release.

- 1491644 When bare metal servers are behind an MX Series router, MX redundancy is provisioned in the network, and a bare metal server pings another bare metal server, the ARP cache of the first bare metal server for the second bare metal server is poisoned with the vRouter compute node's MAC address. This leads to connectivity failure between the two bare metal server.

The cause is that when the ARP request from BMS1 is flooded to a compute node by the MX Series router, the vRouter does the source IP address lookup for the bare metal server IP address in the inet (IPv4) route table. This lookup results in the subnet route pointing to the ECMP next hop of two MX Series routers. This makes the vRouter respond with the virtual host's MAC address to force the packets to Layer 3 processing though the ARP request is not meant for any VMs in that compute node.

- 1496609 For a control node to participate in high availability properly, all the control nodes must have a unique priority. When adding a new control node to an already provisioned high availability enabled cluster, the uniqueness in the priority across the control node is not automatic.

You need to adjust the values to ensure uniqueness as follows:

1. Stop the keepalived process using the **service keepalived stop** command
  2. Edit the `/etc/keepalived/keepalived.conf` file in all the control nodes and modify the priority under the `vrrp_instance INTERNAL*` and `vrrp_instance EXTERNAL*` configuration section, so that all the control nodes have unique values.
  3. Start the keepalived process using the **service keepalived start** command.
- 1495697 When you add a new control node using the **fab install\_new\_contrail** command to a cluster that is already provisioned, there is a possibility that the command might fail due to a timing issue. Even though this command reports failure, it actually does everything as expected. You can proceed using the **fab join\_cluster** command as the next step for adding a new control node.
  - 1474258 When the HTTP port is already used by another instance, the TOR Agent crashes at static init.
  - 1462990 When a TOR agent switchover occurs on the Contrail controller, the OVS controller address changes to the TSN. In this scenario, there is a delay of up to two minutes between the OVS being updated in the QFX Series switch and the first multicast packet reaching the TSN node.
  - 1404846 In Juno, VPC VM launch is failing VPC API is not supported with Juno. Its planned to be supported in subsequent release.
  - 1464606 Vtep-ctl is only able to list 7 MAC addresses. The TOR agent is learning fewer MAC addresses than are actually present.
  - 1465744 Contrail/MX interop when a VM is using SNAT to reach a bare metal server floating IP address. This happens only in cases where a SNAT instance and destination Floating IP address are on the same compute node.

- 1466777 There is a need to improve api-server and schema initialization times in a scaled setup. On highly scaled setups it takes up to 40 minutes for an API server and schema transformer to converge.
- 1466731 A QFX Series switch does not handle transient duplicate VxLAN IDs for two different VNs. If a VN is deleted and added quickly the TOR switch may go in to a bad state.
- 1468685 Centos6.5 icehouse - single node setup config processes are killed after a node reboot. A single node Centos installation runs into an API server exception.
- 1469366 In a Device Manager setup which has an MX Series router configured for a virtual network, if the MX Series router is down and then comes back up, the latest configuration is not pushed.
- 1475028 One Contrail controller can get stuck initializing for about 5 minutes. When this happens, another controller is initialized and the first one moves to a backup role.
- 1484600 When a device is moved from one QFX Series switch to another Series switch, the MAC address is not learned on the switch for a period up to 12 minutes.
- 1486387 If you configure compute and config services in the same node, you must use the **fab setup\_nova\_aggregate** command after the node is rebooted. If the command is not used, **setup\_nova\_aggregate** will never get executed.
- 1491202 The TSN stops replicating broadcast packets after broadcast packets with a payload of 1400 bytes are sent.
- 1491791 When the mcast\_controller gets switched as part of a control node switchover there is 6 seconds of multicast traffic loss.
- 1493861 When clearing the setup used for inter-VN communication, the compute node might crash.
- 1414850 Interfaces created for logical routers and other constructs that are not on vRouters, do not get accounted for in the dashboard.
- 1403348 If you attach and then detach a security group, the transparent firewall service interface does not have an internal security group.
- 1447401 On multiple VMs in a Docker cluster, they invariably end up on one compute only.
- 1454813 Setup of a vCenter fails if the same **dv\_port** or **dv\_switch** name is part of multiple data centers.
- 1455944 When creating nova instances in Docker containers, the user-data script is not executed.
- 1457854 If you try to create an analyzer VM with **contrail\_flavor\_small** configured, the VM is not created but multiple instances are respawned and all are in an error state.
- 1458794 DNS configuration in Docker container is wrong. A Docker instance does not learn the DNS address provided by the vVrouter.
- 1459505 An xmpp peer is not deleted until the replicated routes are gone.



- 1460241 If you create twelve virtual routers attached to a single logical router and then clear the router, Neutron experiences an error.
- 1461791 When servers in a cluster are reimaged with an ESX ISO image, only one server is successfully reimaged, all other servers in that cluster will be re-imaging in a loop.
- 1463622 If you create multiple compute nodes and multiple virtual machines, return traffic from server to client converges on a single label. Eventually, all the flows converge on one VM on each compute node.
- 1463786 If you create thousands of logical interfaces and thousands of virtual machine interfaces, deleting all the interfaces using the Web user interface might result in the **Too many pending updates to RabbitMQ: 4096** error.
- 1465372 If a bare metal server and a SNAT instance are attached to a public network and a packet is sent from the network namespace (netns) instance to the bare metal server, it gets Layer 3 lookups rather than a bridge table lookup.
- 1467028 In a three-node Vcenter setup, restarting the plugin service might take up to two minutes for the other node to detect and establish mastership.
- 1467031 A TSN drops ARP packets coming from a backup LBaaS network namespace with the **invalid multicast source** error. As a result, the backup LBaaS network namespace instance continuously sends ARP requests for each of the bare metal server IP addresses that have been configured as members of the load-balance pool.
- 1468420 If you create thousands of virtual machine interfaces and logical interfaces with a thousand virtual networks, and then push the configuration using the device manager, the configuration might get repeatedly added and deleted on the MX Series router.
- 1468474 TOR Agent Switchover: BUM/ARP traffic loss. Currently a control node does not implement the graceful restart feature, so MAC routes are immediately withdrawn on the TOR agent during switchover leading to traffic loss.
- 1468886 Sometimes it takes more than half an hour for cmon to bring up mysql during node failure scenarios.
- 1469296 When an MX Series router is providing NAT service for a bare metal server using floating IP addresses and the bare metal server belongs to overlapping subnets, their respective NAT configurations will collide in the NAT pool section of the config and get rejected.
- 1469312 When HAProxy is stopped on a virtual IP node, one out of three glance requests fail.
- 1399812 The OpenStack high availability control node reboot hangs if the Network File System mount point is unavailable during boot.
- 1441810 After rebooting a compute node, sometimes the nova-compute is started before the virtualization API (libvirt) is initialized. Therefore it fails to come up.
- 1480050 If you assign the same FIP address to two virtual machines, only the VM with an active VRRP address should get the FIP traffic.
- 1489610 If two DNS servers are configured and one is down, the DNS request should only be sent to the server that is up.

- 1490788 When a ToR-Agent's status is busy, it can not reply to HAProxy keepalive messages.
- 1492979 Broadcast routes are always programmed with the EVPN as the next hop. So even if there is no MX Series router to flood the traffic, it is still programmed in the composite next hop.

The Vrouter replicates the traffic for the EVPN next hop and eventually the traffic is discarded. This causes the drop statistics count to increase.

- 1469341 The Vcenter setup does not use the svc-monitor. The **contrail-svc-monitor** status needs to be removed from the **contrail-status** command output.
- 1493687 Fragment packets with partial TCP headers get dropped but the flow still gets created and the next fragment gets forwarded to the receiver.

When a packet fragment has a full TCP header and the next fragments offset is 1, then the Vrouter forwards this fragment.

When a fragment packet head is received after 3 or more fragments, it sometime leads to fragment loss.

- 1485754 When a virtual network is extended to a physical router, the Device Manager allocates an IP address for the IRB interface. If the virtual network to physical router association is broken, the Device Manager tries to free the allocated IP address. This call fails. As a result, the IP address that was previously allocated, is no longer available in the free pool.

## Upgrading Contrail Software from Release 2.00 or Greater to Release 2.20

Use the following procedure to upgrade an installation of Contrail software from one release to a more recent release. This procedure is valid starting from Contrail Release 2.00 and greater.



**NOTE:** If you are installing Contrail for the first time, refer to the full documentation and installation instructions in *Installing the Operating System and Contrail Packages*.

Instructions are given for both CentOS and Ubuntu versions. The only Ubuntu versions supported for upgrading are Ubuntu 12.04 and 14.04.2.

To upgrade Contrail software from Contrail Release 2.00 or greater:

1. Download the file `contrail-install-packages-x.xx-xxx.xxx.noarch.rpm | deb` from <http://www.juniper.net/support/downloads/?p=contrail#sw> and copy it to the `/tmp` directory on the config node, as follows:

*CentOS* : `scp <id@server>:/path/to/contrail-install-packages-x.xx-xxx.xxx.noarch.rpm /tmp`

*Ubuntu* : `scp <id@server>:/path/to/contrail-install-packages-x.xx-xx~havana_all.deb /tmp`



**NOTE:** The variables `xxx-xxx` and so on represent the release and build numbers that are present in the name of the installation packages that you download.

2. Install the `contrail-install-packages`, using the correct command for your operating system:

*CentOS*: `yum localinstall /tmp/contrail-install-packages-x.xx-xxx.xxx..noarch.rpm`

*Ubuntu*: `dpkg -i /tmp/contrail-install-packages_x.xx-xxx~icehouse_all.deb`

3. Set up the local repository by running the `setup.sh`:

`cd /opt/contrail/contrail_packages; ./setup.sh`

4. Ensure that the `testbed.py` file that was used to set up the cluster with Contrail is intact at `/opt/contrail/utils/fabfile/testbeds/`.

- Ensure that `testbed.py` has been set up with a combined `control_data` section (required as of Contrail Release 1.10).
- Ensure that the `do_parallel` flag is set to `True` in the `testbed.py` file, see bug [1426522](#) in Launchpad.net.

See *Populating the Testbed Definitions File*.

5. Upgrade the software, using the correct set of commands to match your operating system and vrouter, as described in the following:

Change directory to the **utils** folder:

```
cd /opt/contrail/utils; \
```

Select the correct upgrade procedure from the following to match your operating system and vrouter. In the following, *<from>* refers to the currently installed release number, such as 2.0, 2.01, 2.1:

*CentOS Upgrade Procedure:*

```
fab upgrade_contrail:<from>,/tmp/contrail-install-packages-x.xx-xxx.xxx.noarch.rpm;
```

*Ubuntu 12.04 Procedure:*

```
fab upgrade_contrail:<from>,/tmp/contrail-install-packages-x.xx-xxx~icehouse_all.deb;
```

*Ubuntu 14.04 Upgrade, Two Procedures:*

There are two different upgrade procedures for Ubuntu 14.04 upgrade to Contrail Release 2.20, depending on which vrouter (**contrail-vrouter-3.13.0-35-generic** or **contrail-vrouter-dkms**) is installed in your current setup.

As of Contrail Release 2.20, the recommended kernel version for an Ubuntu 14.04-based system is 3.13.0-40. Both procedures can use the command **fab upgrade\_kernel\_all** to upgrade the kernel.

#### **Ubuntu 14.04 Upgrade Procedure For a System With contrail-vrouter-3.13.0-35-generic:**

Use the following upgrade procedure for Contrail Release 2.20 systems based on Ubuntu 14.04 with the **contrail-vrouter-3.13.0-35-generic** installed. The command sequence upgrades the kernel version and also reboots the compute nodes when finished.

```
fab install_pkg_all:/tmp/contrail-install-packages-x.xx-xxx~icehouse_all.deb;
```

```
fab migrate_compute_kernel;
```

```
fab
```

```
upgrade_contrail:<from>,/tmp/contrail-install-packages-x.xx-xxx~icehouse_all.deb;
```

```
fab upgrade_kernel_all;
```

```
fab restart_openstack_compute;
```

#### **Ubuntu 14.04 Upgrade Procedure For System with contrail-vrouter-dkms:**

Use the following upgrade procedure for Contrail R2.20 systems based on Ubuntu 14.04 with **contrail-vrouter-dkms** installed. The command sequence upgrades the kernel version and also reboots the compute nodes when finished.

```
fab
```

```
upgrade_contrail:<from>,/tmp/contrail-install-packages-x.xx-xxx~icehouse_all.deb;
```

All nodes in the cluster can be upgraded to kernel version 3.13.0-40 by using the following fab command:

```
fab upgrade_kernel_all
```

6. On the OpenStack node, soft reboot all of the virtual machines.

You can do this in the OpenStack dashboard or log into the node that uses the `openstack` role and issue the following commands:

```
source /etc/contrail/openstackrc ; nova reboot <vm-name>
```

You can also use the following fab command to reboot all virtual machines:

```
fab reboot_vm
```

7. Check to ensure that the `nova-novncproxy` service is still running:

```
service nova-novncproxy status
```

If necessary, restart the service:

```
service nova-novncproxy restart
```

8. *(For Contrail Storage option, only.)*

Contrail Storage has its own packages.

To upgrade Contrail Storage, download the file:

```
contrail-storage-packages_x.x-xx*.deb
```

from <http://www.juniper.net/support/downloads/?p=contrail#sw>

and copy it to the `/tmp` directory on the config node, as follows:

```
Ubuntu: scp <id@server>:/path/to/contrail-storage-packages_x.x-xx*.deb /tmp
```



**NOTE:** Use only Icehouse packages (for example, `contrail-storage-packages_2.0-22~icehouse_all.deb`) because OpenStack Havana is no longer supported.

Use the following statement to upgrade the software:

```
cd /opt/contrail/utils; \
```

```
Ubuntu: fab
```

```
upgrade_storage:<from>,/tmp/contrail-storage-packages_2.0-22~icehouse_all.deb;
```

When upgrading to Contrail Release 2.10, add the following steps if you have live migration configured. Upgrades to Release 2.0 do not require these steps. Select the command that matches your live migration configuration.

```
fab setup_nfs_livem
```

or

```
fab setup_nfs_livem_global
```

- Related Documentation**
- *Contrail Getting Started Guide, Release 2.21*
  - *Contrail Feature Guide, Release 2.21*

## Documentation Updates

---

## 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 at the Juniper Networks Technical Documentation 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](mailto: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: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- 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://tools.juniper.net/SerialNumberEntitlementSearch/>

## 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

---

October 2015—Revision 1, Contrail 2.21

August 2015—Revision 1, Contrail 2.20

April 2014—Revision 1, Contrail 1.05

18 March 2014—Revision 1, Contrail 1.04

January 2014—Revision 1, Contrail 1.03

21 October 2013—Revision 1, Contrail 1.02

16 September 2013—Revision 1, Contrail 1.0

Copyright © 2015, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the 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.