

Contrail[®] Networking

Contrail Networking Installation and Upgrade Guide

Published
2021-03-04

Release
2011

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, 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.

Contrail® Networking Contrail Networking Installation and Upgrade Guide
2011

Copyright © 2021 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About the Documentation | xv

Documentation and Release Notes | xvi

Documentation Conventions | xvi

Documentation Feedback | xix

Requesting Technical Support | xix

Self-Help Online Tools and Resources | xx

Creating a Service Request with JTAC | xx

1

Installing and Upgrading Contrail

Understanding Contrail | 2

Understanding Contrail Networking | 2

Understanding Contrail Networking Components | 4

Understanding Contrail Containers | 5

Contrail Containers | 6

Understanding Contrail Microservices Architecture | 6

What is Contrail Microservices Architecture? | 6

Installing Contrail with Microservices Architecture | 7

Understanding contrail-ansible-deployer used in Contrail Command | 7

What is the contrail-ansible-deployer? | 8

playbooks/provision_instances.yml | 8

playbooks/configure_instances.yml | 8

playbooks/install_contrail.yml | 9

Preparing to Install with Contrail Command | 9

Prerequisites | 9

Supported Providers | 9

Configure a Yaml File for Your Environment | 9

Provider Configuration | 10

Global Services Configuration | 13

Contrail Services Configuration | 13

Kolla Services Configuration | 14

Instances Configuration | 14

Installing a Contrail System | 15

Supported Platforms and Server Requirements | 16

Server Requirements and Supported Platforms | 16

Hardware and Server Requirements | 17

Hardware Requirements for Contrail Networking Release 2011 | 17

Baseline Server Requirement for an All-In-One Setup | 17

Contrail Controller for Object Scale (Number of VLANs) | 18

Baseline Server Requirement for a HA Setup | 18

Contrail Insights Scale | 18

Hardware Requirements for Contrail Networking Release 2008 | 18

Baseline Server Requirement for an All-In-One Setup | 18

Contrail Controller for Object Scale (Number of VLANs) | 19

Baseline Server Requirement for a HA Setup | 19

Contrail Insights Scale | 19

Contrail Command | 20

How to Install Contrail Command and Provision Your Contrail Cluster | 20

When to Use This Document | 21

Server Requirements | 21

Software Requirements | 21

How to Obtain Contrail Images | 22

How to Install Contrail Command | 22

Before You Begin | 24

Preparing Your Contrail Command Server for the Installation | 24

Installing Contrail Command | 29

How to Provision Servers into the Contrail Cluster | 33

Before You Begin | 33

How to Provision the Contrail Cluster | 33

Sample command_servers.yml Files for Installing Contrail Command | 51

Minimal command_servers.yml file | 52

Complete command_servers.yml File | 53

Disaster Recovery and Troubleshooting | 57

| 58

How to Login to Contrail Command | 59

Navigating the Contrail Command UI | 61

Using the Get Started with Contrail Enterprise Multicloud Panel | 62

Navigating to pages using the side panel | 64

Hiding the side panel | 65

Search functionality | 65

Pinning favorite pages | 66

Opening external applications | 72

Using the What's New Panel | 72

Supported Browsers for Installing Contrail Command | 73

Installing a Contrail Cluster using Contrail Command and instances.yml | 74

Importing Contrail Cluster Data using Contrail Command | 79

Adding a New Compute Node to Existing Contrail Cluster Using Contrail Command | 85

How to Deploy Contrail Command and Import a Cluster Using Juju | 89

Overview: Deploying Contrail Command with a Contrail Cluster Using Juju | 90

Preparing the SSL Certificate Authority (CA) for the Deployment | 90

Deploy Contrail Command and Import a Contrail Cluster Using Juju | 93

Example: Config.YML File for Deploying Contrail Command with a Cluster Using Juju | 98

Prerequisites for Contrail Insights and Contrail Insights Flow | 100

Install Contrail Insights on the Juju Cluster after Contrail Command is Installed | 100

Install Contrail Insights Flows on the Juju Cluster after Contrail Insights is Installed | 101

Importing a Canonical Openstack Deployment Into Contrail Command | 104

Overview: Canonical Openstack Deployment into Contrail Command | 105

Importing Canonical Openstack Into Contrail Command | 105

Upgrading Contrail Software | 109

Upgrading Contrail Networking using Contrail Command | 109

Upgrading Contrail Command using Backup Restore Procedure | 112

How to Perform a Zero Impact Contrail Networking Upgrade using the Ansible Deployer | **114**

Updating Contrail Networking using the Zero Impact Upgrade Process in an Environment using Red Hat Openstack | **121**

When to Use This Procedure | **121**

Prerequisites | **122**

Before You Begin | **122**

Updating Contrail Networking in an Environment using Red Hat Openstack | **123**

Updating Contrail Networking using the Zero Impact Upgrade Procedure in a Canonical Openstack Deployment with Juju Charms | **130**

Prerequisites | **130**

When to Use This Procedure | **131**

Recommendations | **131**

Updating Contrail Networking in a Canonical Openstack Deployment Using Juju Charms | **132**

Upgrading Contrail Networking using contrail-ansible Deployer | **136**

Upgrading Contrail Networking using In-Place Upgrade Procedure | **138**

Upgrading Contrail Networking Release 19xx with RHOSP13 to Contrail Networking Release 2011 with RHOSP16.1 | **139**

Backup and Restore Contrail Software | 142

How to Backup and Restore Contrail Databases in JSON Format | **142**

Before You Begin | **142**

Simple Database Backup in JSON Format | **143**

Examples: Simple Database Backups in JSON Format | **147**

Restore Database from the Backup in JSON Format | **150**

Example: How to Restore a Database Using the JSON Backup (Ansible Deployer Environment) | **156**

Example: How to Restore a Database Using the JSON Backup (Red Hat Openstack Deployer Environment) | **160**

Using Contrail Networking with VMware vCenter | 164

VMware vCenter with Containerized Contrail Networking | **164**

Integrating vCenter for Contrail | **164**

Prerequisites | **165**

ESX Agent Manager | **165**

Set Up vCenter Server | **165**

Configure Contrail Parameters	170
Install Contrail	170
Monitor and Manage ContrailVM from ESX Agent Manager	170
Installing and Provisioning Contrail VMware vRealize Orchestrator Plugin	172
Accessing vRO Control Center	173
Installing vRO Plugin	175
Accessing vRO Desktop Client	178
Connecting to vRO using the Desktop Client	178
Connecting to Contrail Controller	179
Deploying Contrail vRO Plugin	182
Integrating Contrail Release 5.0.X with VMware vRealize Orchestrator	182
Components of vRealize Orchestrator	183
Contrail Workflows	183
Managing Networks From Contrail Command and VMware vCenter User Interfaces	187
User Interfaces for Contrail Integration with VMware vCenter	187
Feature Configuration for Contrail vCenter	188
Configuring Underlay Network for ContrailVM	198
Standard Switch Setup	199
Distributed Switch Setup	200
PCI Pass-Through Setup	202
SR-IOV Setup	205
Using Contrail Networking with OpenStack 	210
Setting Up Contrail with Red Hat OpenStack 16.1	210
Understanding Red Hat OpenStack Platform Director	211
Red Hat OpenStack Platform Director	211
Contrail Networking Roles	212
Undercloud Requirements	213
Overcloud Requirements	213
Networking Requirements	214
Compatibility Matrix	215

Installation Summary	215
Setting Up the Infrastructure	216
Target Configuration (Example)	216
Configure the External Physical Switch	218
Configure KVM Hosts	219
Create the Overcloud VM Definitions on the Overcloud KVM Hosts	221
Create the Undercloud VM Definition on the Undercloud KVM Host	223
Setting Up the Undercloud	225
Install the Undercloud	225
Perform Post-Install Configuration	226
Setting Up the Overcloud	227
Configuring the Overcloud	227
Customizing the Contrail Service with Templates (contrail-services.yaml)	232
Customizing the Contrail Network with Templates	233
Installing Overcloud	269
Setting Up Contrail with Red Hat OpenStack 13	270
Understanding Red Hat OpenStack Platform Director	271
Red Hat OpenStack Platform Director	271
Contrail Roles	272
Undercloud Requirements	273
Overcloud Requirements	273
Networking Requirements	273
Compatibility Matrix	275
Installation Summary	275
Setting Up the Infrastructure	275
Target Configuration (Example)	275
Configure the External Physical Switch	277
Configure KVM Hosts	278
Create the Overcloud VM Definitions on the Overcloud KVM Hosts	280
Create the Undercloud VM Definition on the Undercloud KVM Host	282
Setting Up the Undercloud	284
Install the Undercloud	284
Perform Post-Install Configuration	286

Setting Up the Overcloud	286
Configuring the Overcloud	287
Customizing the Contrail Service with Templates (contrail-services.yaml)	293
Customizing the Contrail Network with Templates	294
Installing Overcloud	329
Using Netronome SmartNIC vRouter with Contrail Networking	331
Installing OpenStack Octavia LBaaS with RHOSP in Contrail Networking	334
Configuring Virtual Networks	339
Creating Projects in OpenStack for Configuring Tenants in Contrail	340
Creating a Virtual Network with OpenStack Contrail	342
Creating an Image for a Project in OpenStack Contrail	345
Using Security Groups with Virtual Machines (Instances)	349
Security Groups Overview	349
Creating Security Groups and Adding Rules	349
Using Contrail Resources in Heat Templates	353
Using the Contrail Heat Template	353
Introduction to Heat	354
Heat Architecture	354
Support for Heat Version 2 Resources	354
Heat Version 2 with Service Templates and Port Tuple Sample Workflow	355
Example: Creating a Service Template Using Heat	355
QoS Support in Contrail Networking	358
Quality of Service in Contrail	358
Overview: Quality of Service	358
Contrail QoS Model	358
Features of Fabric Interfaces	359
QoS Configuration Parameters for Provisioning	359
Configuring QoS in Contrail Networking Release 5.0 and Later	359
Queuing Implementation	360
Contrail QoS Configuration Objects	360
Example: Mapping Traffic to Forwarding Classes	361
QoS Configuration Object Marking on the Packet	362

Queuing	363
Configuring Network QoS Parameters	365
Overview	365
QoS Configuration Examples	366
Limitations	367
Load Balancers	367
Using Load Balancers in Contrail	368
Invoking LBaaS Drivers	368
Benefits of Creating Configuration Objects	368
Using a Service Appliance Set as the LBaaS Provider	369
Understanding the Load Balancer Agent	371
F5 Networks Load Balancer Integration in Contrail	372
Example: Creating a Load Balancer	375
Using the Avi Networks Load Balancer for Contrail	376
Support for OpenStack LBaaS	381
OpenStack Neutron LBaaS Version 2.0	381
OpenStack Octavia LBaaS	383
Configuring Load Balancing as a Service in Contrail	385
Overview: Load Balancing as a Service	385
Contrail LBaaS Implementation	386
Configuring LBaaS Using CLI	387
Configuring LBaaS using the Contrail Command UI	390
Optimizing Contrail Networking	399
Multiqueue Virtio Interfaces in Virtual Machines	399
Multiqueue Virtio Overview	399
Requirements and Setup for Multiqueue Virtio Interfaces	399
Contrail Networking OpenStack Analytics	400
Ceilometer Support in Contrail	400
Overview	401
Ceilometer Details	401
Verification of Ceilometer Operation	401
Contrail Ceilometer Plugin	404
Ceilometer Installation and Provisioning	407

Contrail OpenStack APIs | 407

Working with Neutron | 408

Data Structure | 408

Network Sharing in Neutron | 409

Commands for Neutron Network Sharing | 409

Support for Neutron APIs | 410

Contrail Neutron Plugin | 410

DHCP Options | 411

Incompatibilities | 411

Using Contrail with Juju Charms | 412

Installing Contrail with OpenStack by Using Juju Charms | 412

Preparing to Deploy Contrail by Using Juju Charms | 413

Deploying Contrail Charms | 415

Deploy Contrail Charms in a Bundle | 415

Deploying Juju Charms with OpenStack Manually | 423

Options for Juju Charms | 429

Installing Contrail with Kubernetes by Using Juju Charms | 437

Understanding Juju Charms with Kubernetes | 438

Preparing to Deploy Contrail with Kubernetes by Using Juju Charms | 438

Deploying Contrail Charms with Kubernetes | 440

Deploying Contrail Charms in a Bundle | 440

Deploying Juju Charms with Kubernetes Manually | 446

Installing Contrail with Kubernetes in Nested Mode by Using Juju Charms | 451

Installing OpenStack Octavia LBaaS with Juju Charms in Contrail Networking | 455

Using Netronome SmartNIC vRouter with Contrail Networking and Juju Charms | 464

Prepare to Install Contrail Networking by Using Juju Charms | 465

Deploy Contrail Charms in a Bundle | 466

Using Contrail and Contrail Insights with Kolla/Ocata OpenStack | 475

Contrail, Contrail Insights, and OpenStack Kolla/Ocata Deployment Requirements | 475

Software Requirements | 475

Hardware Requirements | 476

Preparing for the Installation | 476

Preparing the Targets | 476

Preparing the Base Host using Ansible Installer | 476

TCP/IP Port Conflicts Between Contrail and Contrail Insights | 477

Plugins to Enable for Contrail and Contrail Insights Deployment | 477

Configuring Contrail Monitoring in Contrail Insights | 477

Compute Monitoring: Listing IP Addresses to Monitor | 478

Configuring Openstack_Controller Hosts for Contrail Insights | 478

Other Contrail Insights group_vars That Must be Enabled in instances.yaml | 478

Contrail Insights License | 479

Run the Playbooks | 479

Accessing Contrail in Contrail Insights Management Infrastructure in UI | 480

Notes and Caveats | 481

Example Instances.yml for Contrail and Contrail Insights OpenStack Deployment | 481

Contrail Insights Installation and Configuration for OpenStack | 485

Getting Started with OpenStack | 486

Requirements | 487

Workflow in Three Steps | 487

Configure OpenStack | 488

Create Ansible Inventory | 488

Ansible Configuration Variables | 489

Configure an HTTP Proxy for Fetching Prerequisites | 489

Install Contrail Insights | 490

Additional Ansible Configuration | 493

Upgrade Notices | 496

Remove a Node from Contrail Insights | 498

Upgrade Contrail Insights for an OpenStack Cluster | 498

Uninstall Contrail Insights from an OpenStack Cluster | 498

Contrail Insights Installation for OpenStack in HA | 499**HA Design Overview | 499****Requirements | 499****Connectivity | 500****Install Contrail Insights for High Availability | 500****Post Installation Tasks | 504****Configuring Role and Resource-Based Access Control | 504****Contrail Role and Resource-Based Access (RBAC) Overview | 504****API-Level Access Control | 505****Rule Sets and ACL Objects | 506****Object Level Access Control | 506****Configuration | 507****Parameter: aaa-mode | 507****Parameter: cloud_admin_role | 508****Global Read-Only Role | 508****Parameter Changes in /etc/neutron/api-paste.ini | 509****Upgrading from Previous Releases | 509****Configuring RBAC Using the Contrail User Interface | 509****Configuring RBAC at the Global Level | 510****Configuring RBAC at the Domain Level | 510****Configuring RBAC at the Project Level | 510****Configuring RBAC Details | 511****RBAC Resources | 513****Configuring Role-Based Access Control for Analytics | 513****Configuring the Control Node with BGP | 514****Configuring the Control Node from Contrail Web UI | 515****Configuring the Control Node with BGP from Contrail Command | 521****Configuring MD5 Authentication for BGP Sessions | 524**

Configuring Transport Layer Security-Based XMPP in Contrail | 526

Overview: TLS-Based XMPP | 526

TLS XMPP in Contrail | 526

Configuring XMPP Client and Server in Contrail | 526

Configuring Control Node for XMPP Server | 526

Configuring DNS Server for XMPP Server | 527

Configuring Control Node for XMPP Client | 527

Configuring Graceful Restart and Long-lived Graceful Restart | 528

Application of Graceful Restart and Long-lived Graceful Restart | 528

BGP Graceful Restart Helper Mode | 529

Feature Highlights | 529

XMPP Helper Mode | 529

Configuration Parameters | 530

Cautions for Graceful Restart | 531

Configuring Graceful Restart using Contrail Command | 532

Scaling Up Contrail Networking Configuration API Server Instances | 534

Scaling Up Contrail Networking Configuration API | 538

From Contrail Ansible Deployer | 538

From Contrail Command UI | 538

From RHOSP Deployer | 539

From JUJU Deployer | 539

About the Documentation

IN THIS SECTION

- [Documentation and Release Notes | xvi](#)
- [Documentation Conventions | xvi](#)
- [Documentation Feedback | xix](#)
- [Requesting Technical Support | xix](#)

Use this guide to install and upgrade Contrail Networking solution. This guide covers various installation scenarios including:

- Contrail Command.
- Contrail with Contrail Insights.
- Contrail with VMware vCenter.
- Contrail with Red Hat.
- Contrail and Contrail Insights with Kolla/Ocata OpenStack.
- Contrail with Juju Charms.

For information on Contrail Networking in cloud-native environments using Kubernetes orchestration, see [Contrail Networking Cloud-Native User Guide](#).

Contrail Networking product documentation is organized into multiple guides as shown in [Table 1 on page xv](#), according to the task you want to perform or the deployment scenario.

Table 1: Contrail Networking Guides

Guide Name	Description
Contrail Networking Cloud-Native User Guide	Provides information about installing and using Contrail Networking in cloud-native environments using Kubernetes orchestration.
Contrail Networking Fabric Lifecycle Management Guide	Provides information about Contrail underlay management and data center automation.
Contrail Networking and Security User Guide	Provides information about creating and orchestrating highly secure virtual networks.

Table 1: Contrail Networking Guides (*continued*)

Guide Name	Description
Contrail Networking Service Provider Focused Features Guide	Provides information about the features that are used by service providers.
Contrail Networking Monitoring and Troubleshooting Guide	Provides information about Contrail Insights and Contrail analytics.

Documentation and Release Notes

To obtain the most current version of all Juniper Networks[®] technical documentation, see the product documentation page on the Juniper Networks website at <https://www.juniper.net/documentation/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <https://www.juniper.net/books>.

Documentation Conventions

[Table 2 on page xvii](#) defines notice icons used in this guide.

Table 2: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 3 on page xvii defines the text and syntax conventions used in this guide.

Table 3: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	- Introduces or emphasizes important new terms. - Identifies guide names. - Identifies RFC and Internet draft titles.	- A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> - RFC 1997, <i>BGP Communities Attribute</i>

Table 3: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. • The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i>>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		


Table 3: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	- In the Logical Interfaces box, select All Interfaces. - To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use 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 Juniper Care or Partner Support Services 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 <https://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/>
- Create a service request online: <https://myjuniper.juniper.net>

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

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- 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://support.juniper.net/support/requesting-support/>.

1

PART

Installing and Upgrading Contrail

Understanding Contrail | **2**

Supported Platforms and Server Requirements | **16**

Contrail Command | **20**

Upgrading Contrail Software | **109**

Backup and Restore Contrail Software | **142**

Using Contrail Networking with VMware vCenter | **164**

Using Contrail Networking with OpenStack | **210**

Using Contrail with Juju Charms | **412**

Using Contrail and Contrail Insights with Kolla/Ocata OpenStack | **475**

Post Installation Tasks | **504**

Understanding Contrail

IN THIS CHAPTER

- [Understanding Contrail Networking | 2](#)
- [Understanding Contrail Networking Components | 4](#)
- [Understanding Contrail Containers | 5](#)
- [Understanding Contrail Microservices Architecture | 6](#)
- [Understanding contrail-ansible-deployer used in Contrail Command | 7](#)

Understanding Contrail Networking

Contrail Networking provides dynamic end-to-end networking policy and control for any cloud, any workload, and any deployment, from a single user interface. It translates abstract workflows into specific policies, simplifying the orchestration of virtual overlay connectivity across all environments.

It unifies policy for network automation with seamless integrations for systems such as: Kubernetes, OpenShift, Mesos, OpenStack, VMware, a variety of popular DevOps tools like Ansible, and a variety of Linux operating systems with or without virtualization like KVM and Docker containers.

Contrail Networking is a fundamental building block of Contrail Enterprise Multicloud for enterprises. It manages your data center networking devices, such as QFX Series Switches, Data Center Interconnect (DCI) infrastructures, as well as public cloud gateways, extending the continuous connectivity from your on-premises to private and public clouds.

Contrail Networking reduces the friction of migrating to cloud by providing a virtual networking overlay layer that delivers virtual routing, bridging, and networking services (IPAM, NAT, security, load balancing, VPNs, etc.) over any existing physical or cloud IP network. It also provides multitenant structure and API compatibility with multitenant public clouds like Amazon Web Services (AWS) virtual private clouds (VPCs) for truly unifying policy semantics for hybrid cloud environments.


For service providers, Contrail Networking automates network resource provisioning and orchestration to dynamically create highly scalable virtual networks and to chain a rich set of Juniper Networks or third-party virtualized network functions (VNFs) and physical network functions (PNFs) to form differentiated service chains on demand.

Contrail Networking is also integrated with Contrail Cloud for service providers. It enables you to run high-performance Network Functions Virtualization (NFV) with always-on reliability so that you can deliver innovative services with greater agility.

Contrail Networking is equipped with always-on advanced analytics capabilities to provide deep insights into application and infrastructure performance for better visualization, easier diagnostics, rich reporting, custom application development, and machine automation. It also supports integration with other analytics platforms like Juniper Networks Contrail Insights and streaming analytics through technologies like Apache Kafka and its API.

Contrail Networking also provides a Graphical User Interface (GUI). This GUI is built entirely using the REST APIs.

Figure 1: Contrail Networking Architecture



RELATED DOCUMENTATION


[Understanding Contrail Containers](#) | 5

Understanding Contrail Networking Components

Contrail Networking is comprised of the following key components:

- *Contrail Networking management Web GUI and plug-ins* integrate with orchestration platforms such as Kubernetes, OpenShift, Mesos, OpenStack, VMware vSphere, and with service provider operations support systems/business support systems (OSS/BSS). Many of these integrations are built, certified, and tested with technology alliances like Red Hat, Mirantis, Canonical, NEC, and more. Contrail Networking sits under such orchestration systems and integrates northbound via published REST APIs. It can be automatically driven through the APIs and integrations, or managed directly using the Web GUI, called Contrail Command GUI.
- *Contrail Networking control and management systems*, commonly called the controller, have several functions. Few of the major functions are:
 - *Configuration Nodes*—This function accepts requests from the API to provision workflows like adding new virtual networks, new endpoints, and much more. It converts these abstract high-level requests, with optional detail, into low-level directions that map to the internal data model.
 - *Control Nodes*—This function maintains a scalable, highly available network model and state by federating with other peer instances of itself. It directs network provisioning for the Contrail Networking vRouters using Extensible Messaging and Presence Protocol (XMPP). It can also exchange network connectivity and state with peer physical routers using open industry-standard MP-BGP which is useful for routing the overlay networks and north-south traffic through a high-performance cloud gateway router.
 - *Analytics Nodes*—This function collects, stores, correlates, and analyzes data across network elements. This information, which includes statistics, logs, events, and errors, can be consumed by end-user or network applications through the northbound REST API or Apache Kafka. Through the Web GUI, the data can be analyzed with SQL style queries.
- *Contrail Networking vRouter* runs on the compute nodes of the cloud or NFV infrastructure. It gets network tenancy, VPN, and reachability information from the control function nodes and ensures native Layer 3 services for the Linux host on which it runs or for the containers or virtual machines of that host. Each vRouter is connected to at least two control nodes to optimize system resiliency. The vRouters run in one of two high performance implementations: as a Linux kernel module or as an Intel Data Plane Development Kit (DPDK)-based process.

Figure 2: Contrail Networking Overview



RELATED DOCUMENTATION

[Understanding Contrail Networking | 2](#)

Understanding Contrail Containers

IN THIS SECTION

- [Contrail Containers | 6](#)

Some subsystems of Contrail Networking solution are delivered as Docker containers.

Contrail Containers

The following are key features of the new architecture of Contrail containers:

- All of the Contrail containers are multiprocess Docker containers.
- Each container has an INI-based configuration file that has the configurations for all of the applications running in that container.
- Each container is self-contained, with minimal external orchestration needs.
- A single tool, *Ansible*, is used for all levels of building, deploying, and provisioning the containers. The *Ansible* code for the Contrail system is named **contrail-ansible** and kept in a separate repository. The *Contrail Ansible* code is responsible for all aspects of Contrail container build, deployment, and basic container orchestration.

Understanding Contrail Microservices Architecture

IN THIS SECTION

- [What is Contrail Microservices Architecture? | 6](#)
- [Installing Contrail with Microservices Architecture | 7](#)

What is Contrail Microservices Architecture?

Employing microservices provides a number of benefits which includes:

- Deploying patches without updating the entire Contrail deployment.
- Better ways to manage the lifecycles of containers.
- Improved user experiences with Contrail provisioning and upgrading.
- Provisioning with minimum information provided.
- Configuring every feature.
- Simplify application complexity by implementing small, independent processes.

The containers and their processes are grouped as services and microservices, and are similar to pods in the Kubernetes open-source software used to manage containers on a server cluster.

- [Configure a Yaml File for Your Environment | 9](#)
- [Installing a Contrail System | 15](#)

This topic provides an overview of **contrail-ansible-deployer** used by *Contrail Command* tool. It is used for installing Contrail Networking with microservices architecture.

To understand Contrail microservices, refer to [“Understanding Contrail Microservices Architecture” on page 6](#). For step by step procedure on how to install Contrail using Contrail Command deployer, refer to [“Installing a Contrail Cluster using Contrail Command and instances.yml” on page 74](#).

What is the contrail-ansible-deployer?

IN THIS SECTION

- [playbooks/provision_instances.yml | 8](#)
- [playbooks/configure_instances.yml | 8](#)
- [playbooks/install_contrail.yml | 9](#)

The **contrail-ansible-deployer** is a set of Ansible playbooks designed to deploy Contrail Networking with microservices architecture.

The **contrail-ansible-deployer** contains three plays:

playbooks/provision_instances.yml

This play provisions the operating system instances for hosting the containers. It supports the following infrastructure providers:

- kvm.
- gce.
- aws.

playbooks/configure_instances.yml

This play configures the provisioned instances. The playbook installs software and configures the operating system to meet the required prerequisite standards. This is applicable to all providers.

playbooks/install_contrail.yml

This play pulls, configures, and starts the Contrail containers.

Preparing to Install with Contrail Command

This section helps you prepare your system before installing Contrail Networking using **contrail-command-deployer**.

Prerequisites

Make sure your system meets the following requirements before running **contrail-command-deployer**.

- Confirm that you are running compatible versions of CentOS, Ansible, Docker, and any other software component for your system in your environment. See [Contrail Networking Supported Platforms List: Release 20xx](#).
- Name resolution is operational for long and short host names of the cluster nodes, through either DNS or the host file.
- For high availability (HA), confirm that the time is in sync between the cluster nodes.
- The time must be synchronized between the cluster nodes using Network Time Protocol (ntp).

Supported Providers

The playbooks support installing Contrail Networking on the following providers:

- bms—bare metal server.
- kvm—kernel-based virtual machine (KVM)-hosted virtual machines.
- gce—Google compute engine (GCE)-hosted virtual machines.
- aws—Amazon Web Services (AWS)-hosted virtual machines.

Configure a Yaml File for Your Environment

IN THIS SECTION

- [Provider Configuration | 10](#)
- [Global Services Configuration | 13](#)
- [Contrail Services Configuration | 13](#)
- [Kolla Services Configuration | 14](#)
- [Instances Configuration | 14](#)

The configuration for all three plays is contained in a single file, **config/instances.yaml**.

The configuration has multiple main sections, including:

The main sections of the **config/instances.yaml** file are described in this section. Using the sections that are appropriate for your system, configure each with parameters specific to your environment.

Provider Configuration

The section **provider_config** configures provider-specific settings.

KVM Provider Example

Use this example if you are in a kernel-based virtual machine (kvm) hosted environment.

NOTE: Passwords are provided in this output for illustrative purposes only. We suggest using unique passwords in accordance with your organization's security guidelines in your environment.

```
provider_config:                                # the provider section contains
  all provider relevant configuration
  kvm:                                           # Mandatory.
    image: CentOS-7-x86_64-GenericCloud-1710.qcow2.xz # Mandatory for provision
    play. Image to be deployed.
    image_url: https://cloud.centos.org/centos/7/images/ # Mandatory for provision
    play. Path/url to image.
    ssh_pwd: contrail123                        # Mandatory for
    provision/configuration/install play. Ssh password set/used.
    ssh_user: centos                           # Mandatory for
    provision/configuration/install play. Ssh user set/used.
    ssh_public_key: /home/centos/.ssh/id_rsa.pub # Optional for
    provision/configuration/install play.
    ssh_private_key: /home/centos/.ssh/id_rsa   # Optional for
    provision/configuration/install play.
    vcpu: 12                                    # Mandatory for provision
    play.
    vram: 64000                                # Mandatory for provision
    play.
    vdisk: 100G                                # Mandatory for provision
    play.
    subnet_prefix: ip-address                   # Mandatory for provision
    play.
    subnet_netmask: subnet-mask                 # Mandatory for provision
    play.
    gateway: gateway-ip-address                # Mandatory for
```

```

provision play.
    nameserver: dns-ip-address                                # Mandatory for
provision play.
    ntpserver: ntp-server-ip-address                            # Mandatory
for provision/configuration play.
    domainsuffix: local                                         # Mandatory for provision
play.

```

BMS Provider Example

Use this example if you are in a bare metal server (bms) environment.

NOTE: Passwords are provided in this output for illustrative purposes only. We suggest using unique passwords in accordance with your organization's security guidelines in your environment.

```

provider_config:
  bms:                                                         # Mandatory.
    ssh_pwd: contrail123                                       # Optional. Not needed if ssh
keys are used.
    ssh_user: centos                                           # Mandatory.
    ssh_public_key: /home/centos/.ssh/id_rsa.pub              # Optional. Not needed if ssh
password is used.
    ssh_private_key: /home/centos/.ssh/id_rsa                 # Optional. Not needed if ssh
password is used.
    ntpserver: ntp-server-ip-address                           # Optional. Needed if
ntp server should be configured.
    domainsuffix: local                                         # Optional. Needed if
configuration play should configure /etc/hosts

```



CAUTION: *SSH Host Identity Keys* must be accepted or installed on the Deployer node before proceeding with Contrail installation.

To do so:

- Make SSH connection to each target machine from the Deployer VM using Deployer user credentials and click **Yes** to accept the *SSH Host Key*.

or

- Set the environmental variable `ANSIBLE_HOST_KEY_CHECKING` value to **False**.

```
ANSIBLE_HOST_KEY_CHECKING=false
```

or

- Set `[defaults] host_key_checking` value to **False** in `ansible.cfg` file.

```
[defaults] host_key_checking=false
```

AWS Provider Example

Use this example if you are in an Amazon Web Services (AWS) environment.

```
provider_config:
  aws:
    ec2_access_key: THIS_IS_YOUR_ACCESS_KEY      # Mandatory.
    ec2_secret_key: THIS_IS_YOUR_SECRET_KEY      # Mandatory.
    ssh_public_key: /home/centos/.ssh/id_rsa.pub # Optional.
    ssh_private_key: /home/centos/.ssh/id_rsa    # Optional.
    ssh_user: centos                             # Mandatory.
    instance_type: t2.xlarge                     # Mandatory.
    image: ami-337be65c                         # Mandatory.
    region: eu-central-1                       # Mandatory.
    security_group: SECURITY_GROUP_ID           # Mandatory.
    vpc_subnet_id: VPC_SUBNET_ID               # Mandatory.
    assign_public_ip: yes                       # Mandatory.
    volume_size: 50                             # Mandatory.
    key_pair: KEYPAIR_NAME                     # Mandatory.
```

GCE Provider Example

Use this example if you are in a Google Cloud environment.

```
provider_config:
  gce:
    # Mandatory.
    service_account_email: # Mandatory. GCE service account email address.
    credentials_file:      # Mandatory. Path to GCE account json file.
    project_id:            # Mandatory. GCE project name.
    ssh_user:              # Mandatory. Ssh user for GCE instances.
    ssh_pwd:               # Optional. Ssh password used by ssh user, not
    needed when public is used
    ssh_private_key:       # Optional. Path to private SSH key, used by by
    ssh user, not needed when ssh-agent loaded private key
    machine_type: n1-standard-4 # Mandatory. Default is too small
    image: centos-7           # Mandatory. For provisioning and configuration
    only centos-7 is currently supported.
    network: microservice-vn   # Optional. Defaults to default
    subnetwork: microservice-sn # Optional. Defaults to default
    zone: us-west1-aA         # Optional. Defaults to ?
    disk_size: 50             # Mandatory. Default is too small
```

Global Services Configuration

This section sets global service parameters. All parameters are optional.

```
global_configuration:
  CONTAINER_REGISTRY: hub.juniper.net/contrail
  REGISTRY_PRIVATE_INSECURE: True
  CONTAINER_REGISTRY_USERNAME: YourRegistryUser
  CONTAINER_REGISTRY_PASSWORD: YourRegistryPassword
```

Contrail Services Configuration

This section sets global Contrail service parameters. All parameters are optional.

```
contrail_configuration: # Contrail service configuration section
  CONTRAIL_VERSION: latest
  UPGRADE_KERNEL: true
```

For a complete list of parameters available for `contrail_configuration.md`, see [Contrail Configuration Parameters for Ansible Deployer](#).

AWS Default Three Node HA Instance

The following example uses three AWS EC2 instances to deploy a three node high availability setup with all roles and default parameters.

```
instances:
  aws1:
    provider: aws
  aws2:
    provider: aws
  aws3:
    provider: aws
```

More Examples

Refer to the following for more configuration examples for instances.

- [GCE Kubernetes \(k8s\) HA with separate control and data plane instances](#)
- [AWS Kolla HA with separate control and data plane instances](#)

Installing a Contrail System

To perform a full installation of a Contrail system, refer to the installation instructions in: [“Installing a Contrail Cluster using Contrail Command and instances.yml”](#) on page 74.

RELATED DOCUMENTATION

| [Installing a Contrail Cluster using Contrail Command and instances.yml](#) | 74

Supported Platforms and Server Requirements

IN THIS CHAPTER

- Server Requirements and Supported Platforms | 16
- Hardware and Server Requirements | 17

Server Requirements and Supported Platforms

This topic discusses server requirements in a Contrail Networking cluster.

Each server must have a minimum of:

- 64 GB memory.
- 300 GB hard drive.
- 4 CPU cores.
- At least one Ethernet port.

A server can either be a physical device or a virtual machine. For scalability and availability reasons, it is highly recommended to use physical servers in most use cases whenever possible.

Server role assignments vary by environment. All non-compute roles can be configured in each controller node if desired in your topology.

All installation images are available in repositories and can also be downloaded from [Contrail Downloads page](#).

The Contrail image includes the following software:

- All dependent software packages needed to support installation and operation of OpenStack and Contrail.
- Contrail Controller software – all components.
- OpenStack release currently in use for Contrail.

All components required for installing the Contrail Controller are available for each Contrail release, for the supported Linux operating systems and versions, and for the supported versions of OpenStack.

For a list of supported platforms for all Contrail Networking releases, see [Contrail Networking Supported Platforms List](#).

Access *Container Tags* are located at [README Access to Contrail Registry 20XX](#).

If you need access to Contrail docker private secure registry, e-mail **contrail-registry@juniper.net** for Contrail container registry credentials.

RELATED DOCUMENTATION

| [Hardware and Server Requirements](#) | 17

Hardware and Server Requirements

IN THIS SECTION

- [Hardware Requirements for Contrail Networking Release 2011](#) | 17
- [Hardware Requirements for Contrail Networking Release 2008](#) | 18

The following tables list the minimum and total memory and disk requirements per x86 server or per virtual machine for installing Contrail Networking.

Hardware Requirements for Contrail Networking Release 2011

The following tables list the minimum and total memory and disk requirements for installing Contrail Networking Release 2011.

Baseline Server Requirement for an All-In-One Setup

Baseline for 32 switches.

Role	vCPUs (hyper threaded)	Memory (GB)	Disk (GB)	Networking
Contrail Command	4	32	100	1 Ethernet port
Control Host	20	48	350	2 Ethernet ports

Contrail Insights + Flow Node	10	50	300	2 Ethernet ports

Contrail Controller for Object Scale (Number of VLANs)

Baseline for 256 switches.

Number of Devices	Number of VMIs (2500 VPGs each with 102 VLANs)	vCPUs (hyper threaded)	RAM (GB)	Disk (GB)
256	256K	38	135	350

Baseline Server Requirement for a HA Setup

Role	vCPUs (hyper threaded)/ Node	Memory (GB)/ Node	Disk (GB)/ Node	Networking
Contrail Command	4	32	100	1 Ethernet port
Control Host (x3)	20	48	350	2 Ethernet ports
Contrail Insights + Flow Node (x3) Note: Number of sampled flows: 500K	10	50	300	2 Ethernet ports

Contrail Insights Scale

Number of Contrail Insight Nodes	Number of VMs	Number of Compute Nodes	vCPUs	RAM (GB)	Disk (TB)
3	10K	300	64	64	2
5	100K	1000	64	128	4

Hardware Requirements for Contrail Networking Release 2008

The following tables list the minimum and total memory and disk requirements for installing Contrail Networking Release 2008.

Baseline Server Requirement for an All-In-One Setup

Baseline for 32 switches.

Role	vCPUs (hyper threaded)	Memory (GB)	Disk (GB)	Networking

Contrail Command	4	32	100	1 Ethernet port
Control Host	20	48	350	2 Ethernet ports
Contrail Insights + Flow Node	10	50	300	2 Ethernet ports

Contrail Controller for Object Scale (Number of VLANs)

Baseline for 128 switches.

Number of Devices	Number of VMIs (2500 VPGs each with 102 VLANs)	vCPUs (hyper threaded)	RAM (GB)	Disk (GB)
128	256K	35	133	350

Baseline Server Requirement for a HA Setup

Role	vCPUs (hyper threaded)/ Node	Memory(GB)/ Node	Disk (GB)/ Node	Networking
Contrail Command	4	32	100	1 Ethernet port
Control Host (x3)	20	48	350	2 Ethernet ports
Contrail Insights + Flow Node (x3) Note: Number of sampled flows: 500K	10	50	300	2 Ethernet ports

Contrail Insights Scale

Number of Contrail Insight Nodes	Number of VMs	Number of Compute Nodes	vCPUs	RAM (GB)	Disk (TB)
3	10K	300	64	64	2
5	100K	1000	64	128	4

Contrail Command

IN THIS CHAPTER

- [How to Install Contrail Command and Provision Your Contrail Cluster | 20](#)
- [How to Login to Contrail Command | 59](#)
- [Navigating the Contrail Command UI | 61](#)
- [Installing a Contrail Cluster using Contrail Command and instances.yml | 74](#)
- [Importing Contrail Cluster Data using Contrail Command | 79](#)
- [Adding a New Compute Node to Existing Contrail Cluster Using Contrail Command | 85](#)
- [How to Deploy Contrail Command and Import a Cluster Using Juju | 89](#)
- [Importing a Canonical Openstack Deployment Into Contrail Command | 104](#)

How to Install Contrail Command and Provision Your Contrail Cluster

IN THIS SECTION

- [When to Use This Document | 21](#)
- [Server Requirements | 21](#)
- [Software Requirements | 21](#)
- [How to Obtain Contrail Images | 22](#)
- [How to Install Contrail Command | 22](#)
- [How to Provision Servers into the Contrail Cluster | 33](#)
- [Sample command_servers.yml Files for Installing Contrail Command | 51](#)

Use this document to install Contrail Command—the graphical user interface for Contrail Networking—and provision your servers or VMs as nodes in a Contrail cluster. Servers or VMs are provisioned into compute nodes, control nodes, orchestrator nodes, Contrail Insights nodes, Contrail Insights Flows nodes, or service nodes to create your Contrail cluster using this procedure.

NOTE: Contrail Insights and Contrail Insights Flows were previously named Appformix and Appformix Flows.

When to Use This Document

We strongly recommend Contrail Command as the primary interface for configuring and maintaining Contrail Networking.

You should, therefore, complete the procedures in this document as an initial configuration task in your Contrail Networking environment.

Server Requirements

A Contrail Networking environment can include physical servers or VMs providing server functions, although we highly recommended using physical servers for scalability and availability reasons whenever possible.

Each server in a Contrail environment must have a minimum of:

- 64 GB memory.
- 300 GB hard drive.
- 4 CPU cores.
- At least one Ethernet port.

For additional information on server requirements for Contrail Networking, see [“Server Requirements and Supported Platforms” on page 16](#).

Software Requirements

- *Contrail Command and Contrail Networking*

Contrail Command and Contrail Networking are updated simultaneously and always run the same version of Contrail Networking software.

Each Contrail Networking release has software compatibility requirements based on the orchestration platform version, the deployer used to deploy the orchestration platform, the supported server operating system version, and other software requirements.

For a list of supported platforms for all Contrail Networking releases and additional environment-specific software requirements, see [Contrail Networking Supported Platforms List](#).

- *Contrail Insights and Contrail Insights Flows*

For Contrail Release 2003, the supported Contrail Insights version is 3.1.15 and the supported Contrail Insights Flows version is 1.0.7. Separate Contrail Insights and Contrail Insights Flows images—often called Appformix or Appformix Flows in earlier releases since Appformix was renamed Contrail Insights—need to be downloaded in Contrail Release 2003 and earlier installations. The step-by-step process for downloading the images is included in [“How to Install Contrail Command” on page 22](#).

Starting in Contrail Release 2005, the Contrail Insights and Contrail Insights Flows images that support a Contrail Networking release are automatically provisioned within Contrail Command. When you download your version of Contrail Command, Contrail Command pulls the Contrail Insights and Contrail Insights Flows images for your Contrail Networking version automatically from within the Juniper Contrail registry. You do not, therefore, need to separately download any individual Contrail Insights software or have awareness of Contrail Insights or Contrail Insights version numbers for your installation.

How to Obtain Contrail Images

The procedures used in this document download the Contrail Command, Contrail Insights, and Contrail Insights Flows software from the Juniper Networks Contrail docker private secure registry at hub.juniper.net. Email contrail-registry@juniper.net to obtain access credentials to this registry.

You will need to know the *Container Tags* for your Contrail image to retrieve Contrail images from the Contrail registry. See [README Access to Contrail Registry 20XX](#).

Contrail Networking images are also available at the [Contrail Downloads page](#). Enter *Contrail Networking* as the product name.

Contrail Insights and Contrail Insights Flows images are also available at the [Contrail Insights Download page](#). Enter *Contrail Insights* as the product name.

How to Install Contrail Command

IN THIS SECTION

- [Before You Begin | 24](#)
- [Preparing Your Contrail Command Server for the Installation | 24](#)
- [Installing Contrail Command | 29](#)

Contrail Command is a single pane-of-glass GUI interface for Contrail Networking. For an optimized Contrail Networking experience, we strongly recommend installing Contrail Command before creating your Contrail clusters. Contrail Command is installed using these instructions.

For additional information on Contrail Command, see [“Understanding Contrail Networking Components”](#) on page 4.

Before You Begin

Ensure your Contrail Command server—the server that will host Contrail Command—is a virtual machine (VM) or a physical x86 server that meets these minimum system requirements:

- 4 vCPUs
- 32 GB RAM
- 100 GB disk storage with all user storage in the “/” partition.

If the “/home” partition exists, remove it and increase the “/” partition by the amount of freed storage.

- Meets the specifications listed in [“Server Requirements” on page 21](#).
- Runs a version of CentOS that supports your version of Contrail Networking.

For a list of CentOS versions that are supported with Contrail Networking and orchestration platform combinations, see [Contrail Networking Supported Platforms List](#).

You can install CentOS with updated packages using the **yum update** command.

- Has access to the Contrail Container registry at hub.juniper.net. This access is needed because the Contrail Command deployer, which includes the Contrail Command docker images, is retrieved from this registry during this installation procedure.

If you do not have access to the Contrail Container registry, email contrail-registry@juniper.net to obtain access credentials. See [README Access to Contrail Registry 20XX](#) for additional information about accessing this registry.

- Has an active connection to the Internet.
- Includes at least one active IP interface attached to the management network. Contrail Command manages Contrail and orchestrator clusters over a management IP interface.

Obtain the container tag for the release that you are installing. A container tag is necessary to identify the Contrail Command container files in the hub.juniper.net repository that are installed during this procedure.

The container tag for any Contrail Release 20xx image can be found in [README Access to Contrail Registry 20XX](#).

Preparing Your Contrail Command Server for the Installation

To prepare your servers or VMs for the installation:

1. Log onto the server that will host Contrail Command and all servers in your Contrail cluster. The servers in your Contrail cluster are the devices that will be provisioned into compute, control, orchestrator, Contrail Insights, Contrail Insights Flows, or service node roles.
2. Verify the hosts in the hosts file, and add the name and IP address of each host that you are adding to the file.

In this example, the hosts file is edited using VI to include the name and IP address of the three other servers—contrail-cluster, insights, and insights-flows—that will be provisioned into the contrail cluster during this procedure.

```
[root@ix-cn-ccmd-01 ~]# cat /etc/hosts
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6

[root@ix-cn-ccmd-01 ~]# vi /etc/hosts
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
10.1.1.2    contrail-cluster
10.1.1.3    insights
10.1.1.4    insights-flows
```

NOTE: The hosts file is typically overwritten during the provisioning process. This step can be skipped in most Contrail cluster provisioning scenarios, but is recommended as a precaution.

3. Verify the hostname file listing on the Contrail Command server.

```
[root@ix-cn-ccmd-01 ~]# cat /etc/hostname
ix-cn-ccmd-01
```

If needed, update the Contrail Command hostname accordingly to match the hostname that you will use in the Contrail Command cluster.

NOTE: The hostname file is typically overwritten during the provisioning process. This step can be skipped in most Contrail cluster provisioning scenarios, but is recommended as a precaution.

4. If you haven't already generated a shared RSA key for the servers in the cluster, generate and share the RSA key.

```
[root@ix-cn-ccmd-01 ~]# ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
```

```

Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:$ABC123
root@ix-cn-ccmd-01
The key's randomart image is:
+---[RSA 2048]---+
|           .o=o*. |
|           . =o+ o |
|           o. o= + |
|           .o.+ +B |
|           ..S.+++ |
|           ...O..O.. |
|           .. oo=oo|
|           . ..=o*=|
|           E...+=|
+-----[SHA256]-----+

[root@ix-cn-ccmd-01 ~]# ssh-copy-id -i /root/.ssh/id_rsa.pub 10.1.1.2
/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed:
"/root/.ssh/id_rsa.pub"
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter
out any that are already installed
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted
now it is to install the new keys
root@10.1.1.2's password:

Number of key(s) added: 1

Now try logging into the machine, with:  "ssh '10.1.1.2'"
and check to make sure that only the key(s) you wanted were added.

[root@ix-cn-ccmd-01 ~]# ssh-copy-id -i /root/.ssh/id_rsa.pub 10.1.1.3
...

[root@ix-cn-ccmd-01 ~]# ssh-copy-id -i /root/.ssh/id_rsa.pub 10.1.1.4
...

```

5. SSH into each server that will be provisioned into the Contrail cluster to confirm reachability and accessibility:

```
[root@ix-cn-ccmd-01 ~]# ssh 10.1.1.2
[root@contrail-cluster ~]# exit
logout
Connection to 10.1.1.2 closed.

[root@ix-cn-ccmd-01 ~]# ssh 10.1.1.3
[root@insights ~]# exit
logout
Connection to 10.1.1.3 closed.

[root@ix-cn-ccmd-01 ~]# ssh 10.1.1.4
[root@insights-flows ~]# exit
logout
Connection to 10.1.1.4 closed.
```

6. Verify that routes to each server are established on your server.

NOTE: The routes connecting the servers are created outside the Contrail Networking environment and the process to create the routes varies by environment. This procedure, therefore, does not provide the instructions for creating these routes.

In this example, the routes are verified on the Contrail Command server.

```
[root@ix-cn-ccmd-01 ~]# ip route
default via 10.102.70.254 dev ens192 proto static metric 100
10.1.1.0/24 dev ens224 proto kernel scope link src 10.1.1.1 metric 101
10.102.70.0/24 dev ens192 proto kernel scope link src 10.102.70.216 metric 100
```

Perform this step on the Contrail Command server and all servers in your Contrail cluster.

7. Ping each server to verify connectivity.

In this example, each node in the Contrail cluster is pinged from the Contrail Command server.

```
[root@ix-cn-ccmd-01 ~]# ping 10.1.1.2
PING 10.1.1.2 (10.1.1.2) 56(84) bytes of data.
64 bytes from 10.1.1.2: icmp_seq=1 ttl=64 time=0.602 ms
64 bytes from 10.1.1.2: icmp_seq=2 ttl=64 time=0.220 ms
```

```

^C
--- 10.1.1.2 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 0.220/0.411/0.602/0.191 ms

[root@ix-cn-ccmd-01 ~]# ping 10.1.1.3
PING 10.1.1.3 (10.1.1.3) 56(84) bytes of data.
64 bytes from 10.1.1.3: icmp_seq=1 ttl=64 time=0.435 ms
64 bytes from 10.1.1.3: icmp_seq=2 ttl=64 time=0.299 ms
64 bytes from 10.1.1.3: icmp_seq=3 ttl=64 time=0.260 ms
^C
--- 10.1.1.3 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2000ms
rtt min/avg/max/mdev = 0.260/0.331/0.435/0.076 ms

[root@ix-cn-ccmd-01 ~]# ping 10.1.1.4
PING 10.1.1.4 (10.1.1.4) 56(84) bytes of data.
64 bytes from 10.1.1.4: icmp_seq=1 ttl=64 time=0.248 ms
64 bytes from 10.1.1.4: icmp_seq=2 ttl=64 time=0.266 ms
64 bytes from 10.1.1.4: icmp_seq=3 ttl=64 time=0.232 ms
^C
--- 10.1.1.4 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 1999ms
rtt min/avg/max/mdev = 0.232/0.248/0.266/0.022 ms
[root@ix-cn-ccmd-01 ~]#

```

Perform this step on the Contrail Command server and all servers in your Contrail cluster.

8. Check the Linux kernel version and, if needed, update the Linux kernel version. If a kernel version update is performed, reboot the server to complete the update.

NOTE: Obtaining the Linux kernel is not shown in this document.

In this example, the Linux kernel is verified on the Contrail Command server.

```

[root@ix-cn-ccmd-01 ~]# uname -a
Linux ix-cn-ccmd-01 3.10.0-1062.el7.x86_64 #1 SMP Wed Aug 7 18:08:02 UTC 2019
x86_64 x86_64 x86_64 GNU/Linux

[root@ix-cn-ccmd-01 ~]# ls

```

```

anaconda-ks.cfg  kernel-3.10.0-1062.12.1.el7.x86_64.rpm

[root@ix-cn-ccmd-01 ~]# rpm -ihv kernel-3.10.0-1062.12.1.el7.x86_64.rpm
warning: kernel-3.10.0-1062.12.1.el7.x86_64.rpm: Header V3 RSA/SHA256 Signature,
key ID f4a80eb5: NOKEY
Preparing...                               ##### [100%]
Updating / installing...
 1:kernel-3.10.0-1062.12.1.el7             ##### [100%]

[root@ix-cn-ccmd-01 ~]# shutdown -r now

```

After the server reboots, confirm that the kernel is updated.

```

[root@ix-cn-ccmd-01 ~]# uname -a
Linux ix-cn-ccmd-01 3.10.0-1062.12.1.el7.x86_64 #1 SMP Tue Feb 4 23:02:59 UTC
2020 x86_64 x86_64 x86_64 GNU/Linux

```

Perform this step on the Contrail Command server and all servers in your Contrail cluster.

Installing Contrail Command

To install Contrail Command onto a server:

1. Log into the server that will host the Contrail Command containers. This server will be called the Contrail Command server for the remainder of this procedure.

```

$ ssh root@10.12.70.192
root@10.12.70.192's password: password

```

2. Remove all installed Python Docker libraries—**docker** and **docker-py**—from the Contrail Command server:

```

pip uninstall docker docker-py

```

The Python Docker libraries will not exist on the server if a new version of CentOS 7-based software was recently installed. Entering this command when no Python Docker libraries are installed does not harm any system functionality.

The Contrail Command Deployer, which is deployed later in this procedure, installs all necessary libraries, including the Python Docker libraries.

3. Install and start the Docker Engine.

There are multiple ways to perform this step. In this example, Docker Community Edition version 18.03 is installed using **yum install** and **yum-config-manager** commands and started using the **systemctl start docker** command.

NOTE: The Docker version supported with Contrail Networking changes between Contrail releases and orchestration platforms. See [Contrail Networking Supported Platforms List](#). The **yum install -y docker-ce-18.03.1.ce** command is used to illustrate the command for one version of Docker.

```
yum install -y yum-utils device-mapper-persistent-data lvm2
yum-config-manager --add-repo
https://download.docker.com/linux/centos/docker-ce.repo
yum install -y docker-ce-18.03.1.ce
systemctl start docker
```

4. Retrieve the **contrail-command-deployer** Docker image by logging into *hub.juniper.net* and entering the **docker pull** command.

```
docker login hub.juniper.net --username <container_registry_username> --password
<container_registry_password>
docker pull hub.juniper.net/contrail/contrail-command-deployer:<container_tag>
```

Variables:

- **<container_registry_username>** and **<container_registry_password>**—Registry access credentials. You can email contrail-registry@juniper.net to obtain your username and password credentials to access the Contrail Container registry.
- **<container_tag>**—container tag for the Contrail Command (UI) container deployment for the release that you are installing. The **<container_tag>** for any Contrail Release 20xx image can be found in [README Access to Contrail Registry 20XX](#).

5. Create and save the **command_servers.yml** configuration file on the Contrail Command server.

The configuration of the **command_servers.yml** file is unique to your environment and complete documentation describing all **command_servers.yml** configuration options is beyond the scope of this document. Two sample **command_servers.yml** files for a Contrail environment are provided with this document in [“Sample command_servers.yml Files for Installing Contrail Command”](#) on page 51 to provide configuration assistance.

Be aware of the following key configuration parameters when configuring the **command_servers.yml** file for Contrail Command:

- The **contrail_config** hierarchy defines the Contrail Command login credentials:

```
contrail_config:
  database:
    type: postgres
    dialect: postgres
    password: contrail123
  keystone:
    assignment:
      data:
        users:
          admin:
            password: contrail123
```



CAUTION: For security purposes, we strongly recommend creating unique username and password combinations in your environment.

- (Contrail Networking Release 2003 or earlier) The following configuration lines must be entered if you want to deploy Contrail Insights and Contrail Insights Flows:

NOTE: Appformix and Appformix Flows were renamed Contrail Insights and Contrail Insights Flows. The Appformix naming conventions still appear during product usage, including within these directory names.

```
---
user_command_volumes:
- /opt/software/appformix:/opt/software/appformix
- /opt/software/xflow:/opt/software/xflow
```

The configuration lines must be entered outside of the “**command_servers**” hierarchy, either immediately after the “---” at the very top of the file or as the last two lines at the very bottom of the file. See [“Complete command_servers.yml File” on page 53](#) for an example of these lines added at the beginning of the **command_servers.yml** file.

This step is not required to install Contrail Insights and Contrail Insights Flows starting in Contrail Networking Release 2005.

6. Run the **contrail-command-deployer** container to deploy Contrail Command.

```
docker run -td --net host -v
<ABSOLUTE_PATH_TO_command_servers.yml_FILE>:/command_servers.yml --privileged
--name contrail_command_deployer
hub.juniper.net/contrail/contrail-command-deployer:<container_tag>
```

where `<ABSOLUTE_PATH_TO_command_servers.yml_FILE>` is the absolute path to the `command_servers.yml` file that you created in step 5, and `<container_tag>` is the container tag for the Contrail Command (UI) container deployment for the release that you want to install. The `<container_tag>` for any Contrail Release 20xx image can be found in [README Access to Contrail Registry 20XX](#).

7. (Optional) Track the progress of step 6.

```
docker logs -f contrail_command_deployer
```

8. Verify that the Contrail Command containers are running:

```
[root@centos254 ~]# docker ps -a
```

CONTAINER ID	IMAGE	<trimmed>	STATUS	<trimmed>	NAMES
2e62e778aa91	hub.juniper.net/...		Up	<trimmed>	contrail_command
c8442860e462	circleci/postgre...		Up	<trimmed>	contrail_psycopg2
57a666e93d1a	hub.juniper.net/...		Exited	<trimmed>	contrail_command_deployer

The `contrail_command` container is the GUI and the `contrail_psycopg2` container is the database. Both containers should have a STATUS of **Up**.

The `contrail-command-deployer` container should have a STATUS of **Exited** because it exits when the installation is complete.

9. Open a web browser and enter `https://<Contrail-Command-Server-IP-Address>:9091` as the URL. The Contrail Command home screen appears.

Enter the username and password combination specified in the `command_servers.yml` file in step 5. If you use the sample `command_servers.yml` files in ["Sample command_servers.yml Files for Installing Contrail Command" on page 51](#), the username is `admin` and the password is `contrail123`.



CAUTION: For security purposes, we strongly recommend creating unique username and password combinations in your environment.

For additional information on logging into Contrail Command, see ["How to Login to Contrail Command" on page 59](#).

How to Provision Servers into the Contrail Cluster

IN THIS SECTION

- [Before You Begin | 33](#)
- [How to Provision the Contrail Cluster | 33](#)

Use this procedure to provision servers into your Contrail cluster. A Contrail cluster is a collection of interconnected servers that have been provisioned as compute nodes, control nodes, orchestrator nodes, Contrail Insights nodes, Contrail Insights Flows nodes, or service nodes in a cloud networking environment.

Before You Begin

Before you begin:

- Plan your topology.
- Ensure an out-of-band management network is established.
- Ensure Contrail Command is installed. See [“How to Install Contrail Command” on page 22](#).
- Ensure all servers hosting Contrail cluster functions meet the specifications listed in [“Server Requirements” on page 21](#).

How to Provision the Contrail Cluster

To provision the Contrail cluster:

1. (Contrail Networking Release 2003 target release installations using Appformix only) Download the Appformix and—if you are also using Appformix Flows—the Appformix Flows images from the [Contrail Appformix Download page](#).

NOTE: Appformix and Appformix Flows were renamed Contrail Insights and Contrail Insights Flows. The Appformix filename conventions are used to name these files for use with Contrail Networking Release 2003.

For Contrail Release 2003, the supported AppFormix version is 3.1.15 and the supported AppFormix Flows version is 1.0.7.

```
appformix-<version>.tar.gz  
appformix-platform-images-<version>.tar.gz
```

```

appformix-dependencies-images-<version>.tar.gz
appformix-network_device-images-<version>.tar.gz
appformix-openstack-images-<version>.tar.gz
appformix-flows-<version>.tar.gz
appformix-flows-ansible-<version>.tar.gz

```

- Copy the **tar.gz** files to the **/opt/software/appformix/** directory on the Contrail Command server.
- Copy your AppFormix license to the **/opt/software/appformix/** directory.
- (Appformix Flows environments only) Copy the two **appformix-flows** files to the **/opt/software/xflow** directory.

You can ignore this step if you are not using Appformix Flows.

You can skip this step if you are using Contrail Networking Release 2005 or later or are not using Appformix or Appformix Flows in your environment.


2. Login to Contrail Command at **https://<Contrail-Command-Server-IP-Address>:9091** in most scenarios. See [“How to Login to Contrail Command” on page 59](#) if you are not seeing the Contrail Command login screen at this URL.

Leave the **Select Cluster** field blank to enter Contrail Command in a wizard that guides you through the cluster provisioning process. If Contrail Command is not currently managing a cluster, this is your only Contrail Command login option.


Your Contrail Command access credentials were specified in the *command_servers.yml* files in step 5 when you installed Contrail Command. If you used the sample *command_servers.yml* file to enable Contrail Command, your username is **admin** and your password is **contrail123**.

NOTE: Username and password combinations are provided in this document for illustrative purposes only. We suggest using unique username and password combinations to maximize security in accordance with your organization's security guidelines.

3. You are placed into the **Infrastructure > Clusters** menu upon login. Click the **Add Cluster** button to start the cluster provisioning process.



4. Click the **Credentials** tab to move to the **Credentials** box, then the **Add** button to add access credentials for a device that will be added to the cluster.




5. In the **Add** box, add the access credentials for a device in your cluster. Click the **Add** button to complete the process and add the access credentials.

The screenshot shows the Contrail Command web interface. On the left, a sidebar lists steps: STEP 1 Inventory, STEP 2 Provisioning Options, STEP 3 Infrastructure Networks, STEP 4 (optional) Overcloud, STEP 5 (optional) Undercloud Nodes, STEP 6 (optional) Jumpstart Nodes, STEP 7 Control Nodes, and STEP 8 Orchestrator Nodes. The main area displays the 'Add' dialog box with the following fields: Name (compute1-root), SSH User (root), SSH Password (masked with dots), Confirm SSH Password (masked with dots), and Key Pair (a dropdown menu). At the bottom of the dialog are 'Cancel' and 'Add' buttons. The background shows a blurred view of the 'Available servers' box with an 'Add' button.

Repeat steps 4 and 5 to add the access credentials for each server or VM in your cluster.

6. After clicking the **Add** button to add the credentials of your last server or VM, click the **Server** tab to return to the **Available servers** box.

- Click the **Add** button in the **Available Servers** box.



The **Create Server** dialog box appears.

- Complete the fields in the **Create Server** dialog box for each physical server or VM in your Contrail cluster. Each physical server or VM that will function as a compute node, control node, orchestrator

node, Contrail Insights node, Contrail Insights Flows node, or service node in your cluster must be added as a server at this stage of the provisioning process.

Field descriptions:

- **Choose Mode**—Options include: *Express*, *Detailed*, or *Bulk Import (CSV)*. We recommend using the *Detailed* or *Bulk Import (CSV)* modes in most environments to ensure all server field data is entered and to avoid performing manual configuration tasks later in the procedure.
 - *Express*—includes a limited number of required fields to enter for each server or VM.
 - *Detailed*—provides all fields to enter for each server or VM.
 - *Bulk Import (CSV)*—Import the physical server or VM fields from a CSV file.
- **Select workload type this server will be used for**
 - *Physical/Virtual Node*—A virtualized physical server or a VM. This is the option used for most servers or VMs in Contrail Networking environments.
 - *Baremetal*—A non-virtualized server.
- **Hostname**—the name of the physical server or VM.
- **Management IP**—the management IP address of the physical server or VM.
- **Management Interface**—the name of the management-network facing interface on the physical server or VM.
- **Credentials**—Select any credentials that appear in the drop-down menu.

- **Disk Partition(s)**—(Optional) Specify the disk partitions that you want to use.

This field is often left blank.

- **Name** (Network interfaces)—the name of a network-facing interface on the physical server or VM.
- **IP Address** (Network interfaces)—the IP address of the network-facing interface on the physical server or VM.

Click **Add** in the *Network Interfaces* box to add additional network interfaces for the server or VM.

Click the **Create** button after completing all fields to add the server or VM.

Repeat this step for each physical server or VM that will function as a compute node, control node, orchestrator node, Contrail Insights node, Contrail Insights Flows node, or service node in the Contrail cluster.

9. You are returned to the **Infrastructure > Clusters > Servers** menu after adding the final server. Click the **Next** button to proceed to the **Provisioning Options** page.

The screenshot shows the Contrail Command web interface. The top navigation bar includes the Contrail Command logo, the 'INFRASTRUCTURE' menu, and the 'Clusters' sub-menu. The main content area is titled 'Servers' and contains a table of available servers. The table has columns for 'NAME', 'TYPE', 'IP ADDRESS', and 'SERVER PROFILE'. Three servers are listed: 'ix-afx-03', 'ix-afx-flows-03', and 'ix-aio-03', all with a 'physical/virtual' type and IP addresses 10.1.1.3, 10.1.1.4, and 10.1.1.2 respectively. A 'Next' button is located at the bottom right of the table, highlighted with a red box. A sidebar on the left shows the progress of the setup process, with 'STEP 1 Inventory' selected.

NAME	TYPE	IP ADDRESS	SERVER PROFILE
ix-afx-03	physical/virtual	10.1.1.3	...
ix-afx-flows-03	physical/virtual	10.1.1.4	...
ix-aio-03	physical/virtual	10.1.1.2	...

No items selected

Next

10. Complete the fields on the **Provisioning Options** page.

The screenshot shows the 'Provisioning Options' page in the Contrail Command interface. The sidebar on the left lists steps from 'Inventory' to 'Orchestrator Nodes', with 'Provisioning Options' being the active step. The main content area contains several configuration fields:

- Choose Provisioning Manager:** Two radio buttons are present; 'Contrail Enterprise Multicloud' is selected.
- Cluster Name:** A text field containing 'contrail-cluster-1'.
- Container Registry:** A text field containing 'hub.juniper.net/contrail'.
- Insecure:** An unchecked checkbox.
- Container Registry Username:** A text field containing 'admin'.
- Container Registry Password:** A masked password field with a toggle to show/hide.
- Contrail Version:** A text field containing '2003.1.40'.
- Domain Suffix:** A text field containing 'local'.
- NTP Server:** A text field containing '172.21.100.10'.
- Default Vrouter Gateway:** A text field containing '10.1.11.2'.
- Encapsulation Priority:** A dropdown menu showing 'VXLAN,MPLSoUDP,MPL...'.
- Fabric Management:** A checked checkbox.

At the bottom of the form are 'Previous' and 'Next' buttons.

Field Descriptions:

- **Choose Provisioning Manager**
 - **Contrail Cloud**—Contrail Cloud Provisioning Manager. Do not use this provisioning manager option.
 - **Contrail Enterprise Multicloud**—(Default) Contrail Enterprise Multicloud Provisioning Manager. Select Contrail Enterprise Multicloud as your provisioning manager.

The remaining steps of this procedure assume Contrail Enterprise Multicloud is selected as the provisioning manager.

- **Cluster Name**—Name the Contrail cluster.
- **Container Registry**—Path to the container registry to obtain the Contrail Networking image. The path to the Juniper container registry is `hub.juniper.net/contrail` and is set as the default container registry path. Enter this path or the path to the repository used by your organization.
- **Insecure** checkbox—This option should only be selected if you want to connect to an insecure registry using a non-secure protocol like HTTP.

This box is unchecked by default. Leave this box unchecked to connect to the Juniper container registry at `hub.juniper.net/contrail` or to access any other securely-accessible registry.

- **Container Registry Username**—Username to access the container registry.

The Juniper container registry is often used in this field to obtain the Contrail Networking image. Email contrail-registry@juniper.net to receive a registry username and password combination to access the Juniper container registry.

- *Container Registry Password*—Password to access the container registry.

The Juniper container registry is often used in this field to obtain the Contrail Networking image. Email contrail-registry@juniper.net to receive a registry username and password combination to access the Juniper container registry.

- *Contrail Version*—Specify the version of the Contrail Networking image to use for the upgrade that is in the repository.

You can use the *latest* tag to retrieve the most recent image in the repository, which is the default setting. You can also specify a specific release in this field using the version's release tag.

See [README Access to Contrail Registry 20XX](#) to obtain the release tag for any Contrail Networking Release 20xx release tag.

- *Domain Suffix*—(Optional) Domain name for the cluster.
- *NTP Server*—The IP address of the NTP server.
- *Default vRouter Gateway*—The IP address of the default vRouter gateway.

This address is typically the IP address of the interface on the leaf device in the fabric that connects to the server's network-facing interface.

- *Encapsulation Priority*—Select the Encapsulation priority order from the drop down menu.

Select *VXLAN*, *MPLSoUDP*, *MPLSoGRE* in most Contrail Networking environments.

- *Fabric Management* checkbox—Select this option if your deploying in an environment using Openstack for orchestration.
- Click **Add** and enter the following **Key/Value** pairs.

Key	Value
CONTROL_NODES	List of comma-separated user data interface IP addresses for the controller(s)
PHYSICAL_INTERFACE	The user data interface name
TSN_NODES	List of comma-separated user data interface IP addresses for the Contrail Service Node(s)
CONTRAIL_CONTAINER_TAG	The container tag for the desired Contrail and OpenStack release combination as specified in README Access to Contrail Registry 20XX

Key	Value
API__DEFAULTS__enable_latency_stats_log	(Optional. Available starting in Contrail Networking Release 2008) Enable logging and storing of latency statistics in Contrail Networking and Contrail Insights for calls to Cassandra, Zookeeper, and Keystone from the API server.
API__DEFAULTS__enable_api_stats_log	(Optional Available starting in Contrail Networking Release 2008) Enable logging and storing of latency statistics and call time statistics in Contrail Networking and Contrail Insights for Rest API calls.

The screenshot shows the 'Contrail Configuration' section of the provisioning process. It includes a table with the following entries:


Key	Value
CONTROL_NODES	10.1.11.1
PHYSICAL_INTERFACE	ens256
TSN_NODES	10.1.11.1
CONTRAIL_CONTAINER_TAG	2003.1.40

At the bottom of the configuration area, there is a '+ Add' button. Navigation buttons 'Previous' and 'Next' are located at the bottom of the page.

Click the **Next** button to proceed to the **Control Nodes** provisioning page.

- From the **Control Nodes** provisioning page, assign any server that you created in step 8 as a control node by clicking the > icon next to the server to move it into the **Assigned Control Nodes** box.

You have the option to remove roles from a control node within the **Assigned Control Nodes**. There is no need to remove control node roles in most deployments and you should only remove roles if you are an expert user familiar with the consequences.



(Installations using VMWare vCenter only) Complete the following steps to install a control node that is integrated with VMware vCenter. For additional information on vCenter integration with Contrail Networking, see *Understanding VMware-Contrail Networking Fabric Integration*.

Prerequisites:

- Installed vCenter version 6.5 or later.
- Installed ESX version 6.5 or later.
- A vCenter license with Distributed Virtual Switch (DVS) support.
- Login credentials for vCenter.

To perform the integration:

- a. Select the **Manage vCenter** check box.

The vCenter Credentials section is displayed.

- b. Enter the following information:

- Enter the vCenter IP address in the **vCenter IP Address** field.
- In the **Data Center Name** field, enter the name of the data center under vCenter that CVFM will work on.

- Enter the vCenter username in the **Username** field.
 - Enter the vCenter password in the **Password** field.
- c. Click **>**, next to the name of the server, to assign a server from the Available Servers table as a control node. The server is then added to the Assigned Control Nodes table.

Note that the **contrail_vcenter_fabric_manager_node** is added to the list of roles.

- d. Click **Next**.

After assigning all control nodes, click the **Next** button to move to the **Orchestrator Nodes** provisioning page.

12. Select your orchestration platform from the **Orchestrator Type** drop-down menu.

Assign any one of the servers that you created in step 8 as an orchestrator node by clicking the **>** icon next to the server to move it into the **Assigned nodes** box.

The screenshot shows the 'Orchestrator Nodes' configuration page in the Contrail Command interface. On the left, a sidebar lists steps from 'Inventory' to 'Compute Nodes', with 'Orchestrator Nodes' (Step 8) highlighted. The main panel features a dropdown menu for 'Orchestrator type' set to 'Openstack'. Below this, there are two tables: 'Available servers' and 'Assigned Openstack nodes'. The 'Available servers' table lists 'ix-cn-afx-flows-01' (IP 10.1.1.4) and 'ix-cn-afx-01' (IP 10.1.1.3). The 'Assigned Openstack nodes' table shows 'ix-cn-cluster-01' (IP 10.1.1.2) has been assigned. Below the assigned node table, a 'Roles' section lists 'openstack_control_node', 'openstack_network_node', 'openstack_storage_node', and 'openstack_monitoring_node'. At the bottom, 'Previous' and 'Next' buttons are visible.

The remaining processes for this step depend on your orchestration platform:

- **Openstack**

Click the **Show Advanced** box then scroll to **Kolla Globals** and click **+Add**.

Add the following Kolla global **Key** and **Value** pairs in most environments:

Key	Value
enable_haproxy	no

Key	Value
enable_ironic	no
enable_swift	yes
swift_disk_partition_size	20GB

After assigning all orchestrator nodes and Kolla global keys and values, click the **Next** button to progress to the **Compute Nodes** provisioning page.

- **Kubernetes**

Select the Kubernetes nodes from the list of available servers and assign corresponding roles to the servers.

By default, the Kubernetes nodes are assigned the `kubernetes_master_node`, `kubernetes_kubemanager_node`, and `kubernetes_node` roles.

After assigning roles to all nodes, click the **Next** button to progress to the **Compute Nodes** provisioning page.

13. Assign any server that you created in step 8 as a compute node by clicking the > icon next to the server to move it into the **Assigned Compute nodes** box.

Enter the default vRouter gateway IP Address in the **Default Vrouter Gateway** box after moving the server into the **Assigned Compute nodes** box.

The screenshot shows the 'Compute Nodes' provisioning page in the Contrail Command interface. The left sidebar displays a progress bar with steps 1 through 9, where Step 9 'Compute Nodes' is the current step. The main area is titled 'Clusters' and contains a 'Datapath encryption' checkbox (unchecked). Below this, there are two panels: 'Available servers' and 'Assigned Compute nodes'. The 'Available servers' panel has a search bar and a table with columns 'HOSTNAME', 'IP ADDRESS', and 'DISK PARTITION'. It lists two servers: 'ix-cn-afx-flows-01' (10.1.1.4) and 'ix-cn-afx-01' (10.1.1.3), each with a right-pointing arrow icon. The 'Assigned Compute nodes' panel also has a search bar and a table with the same columns. It lists one server: 'ix-cn-cluster-01' (10.1.1.2) with a trash icon. Below the table, there is a 'Default Vrouter Gateway' field with the value '172.16.1.10', a 'Type' dropdown menu set to 'Kernel', and a 'Next' button at the bottom right.

After assigning all compute nodes, click the **Next** button to progress to the **Contrail Service Nodes** provisioning page.

14. Assign any server that you created in step 8 as a Contrail Services node by clicking the > icon next to the server to move it into the **Assigned Service Nodes** box.

Contrail service nodes are only used in environments with bare metal servers. If you are not using Contrail Service nodes in your environment, click the **Next** button without assigning any servers into the **Assigned Service Nodes** box.

The default vRouter gateway IP Address might be autocompleted in the **Default Vrouter Gateway** box. This default vRouter gateway is typically the IP address of a leaf device in the fabric that is directly connected to the server fulfilling the service node role.

After assigning all Contrail Service nodes, click the **Next** button to progress to the **Insights Nodes** provisioning page.

NOTE: The **Insights Nodes** provisioning workflow is called the **Appformix Nodes** workflow in Contrail Networking Release 2005 and earlier releases.

15. Contrail Insights is an optional product that isn't used in all environments. If your are not using Contrail Insights in your environment, simply click the **Next** button without assigning a server as an Appformix node in this step.

NOTE: Appformix was renamed Contrail Insights. The Appformix naming is still used in some Contrail Command screens.

- *Contrail Insights*

If you are using Contrail Insights in your environment, click the > icon next to the server or VM in the **Available servers** box to move it into the **Assigned Insights Nodes** box.

NOTE: The **Assigned Insights Nodes** box is called **Assigned Appformix Nodes** in Contrail Networking Release 2005 and earlier releases.

By default, the server is assigned the *appformix_platform_node* role. You can maintain this default setting in most environments. If the role needs to be changed, click within the **Roles** drop-down menu and select from the available roles.

- *Contrail Insights Flows*

If you are also using Contrail Insights Flows in your environment, click the > icon next to the server or VM in the **Available servers** box to move it into the **Assigned Insights Nodes** box.

NOTE: The **Assigned Insights Nodes** box is called **Assigned Appformix Nodes** in Contrail Networking Release 2005 and earlier releases.

Click within the **Roles** drop-down menu and uncheck the default *appformix_platform_node* role selection. Select *appformix_bare_host_node* from within the **Roles** drop-down menu to set it as the role.

Click the **Next** button to progress to the **Appformix Flows** provisioning page.

16. Contrail Insights Flows is an optional product that isn't used in all environments. If you are not using Contrail Insights Flows in your environment, simply click the **Next** button without assigning a server as an Appformix Flows node in this step.

NOTE: Appformix Flows was renamed Contrail Insights Flows. The Appformix Flows naming is still used on this Contrail Command page.

If you are using Contrail Insights Flows in your environment, make the following configuration selections:

- *Appformix Flows Node Provisioning Type:*
 - *Out-of-Band*—(Default) The Appformix Flows node is managed from an out-of-band management network.
 - *In-Band*—The Appformix Flows node is managed from an in-band connection.
- *Virtual IP Address*—The virtual IP address management address on the Appformix Flows node that connects the node to the management network.

(Contrail Insights and Contrail Insights Flows on same server only) Starting in Contrail Networking Release 2008, you can enable Contrail Insights and Contrail Insights Flows on the same server node.

Perform these steps if you are enabling Contrail Insights and Contrail Insights Flows on the same node:

- a. Click the **Show Advanced** box. The advanced configuration options appear.
- b. From the **AppFormix Flows Configuration Parameters** box, click the **+Add** option to open the **Key** and **Value** configuration options.

Add the following key value pairs:

- **Key:** health_port
Value: 8205
- **Key:** kafka_broker_port
Value: 9195
- **Key:** zookeeper_client_port
Value: 3281
- **Key:** redis_port
Value: 6479


The screenshot shows the Contrail Command web interface. The top navigation bar includes the Contrail Command logo, the 'INFRASTRUCTURE' section, and the 'Clusters' sub-section. On the right of the navigation bar are icons for notifications, user 'admin', and a help icon. The main content area is divided into a left sidebar and a main panel. The sidebar shows a vertical list of steps: STEP 7 Control Nodes, STEP 8 Orchestrator Nodes, STEP 9 (optional) Compute Nodes, STEP 10 (optional) Contrail Service Nodes, STEP 11 (optional) Insights Nodes, STEP 12 (optional) Contrail Insight Flows (which is the current step and highlighted with a blue circle), STEP 13 Summary, and STEP 14 Provisioning. The main panel is titled 'AppFormix Flows Configuration Parameters'. It contains a table with four rows, each representing a configuration parameter. The first row is 'health_port' with a value of '8205'. The second row is 'kafka_broker_port' with a value of '9195'. The third row is 'zookeeper_client_port' with a value of '3281'. The fourth row is 'redis_port' with a value of '6479'. Each row has a 'Key' column, a 'Value' column, and a trash icon on the right. Below the table is a '+ Add' button. At the bottom of the panel are 'Previous' and 'Next' buttons.

Key*	Value*	
health_port	8205	▼ ▲ 🗑️
kafka_broker_port	9195	▼ ▲ 🗑️
zookeeper_client_port	3281	▼ ▲ 🗑️
redis_port	6479	▼ ▲ 🗑️

+ Add


Previous Next

Click the > icon next to the server or VM in the **Available servers** box to move it into the **Assigned AppFormix Flows Nodes** box.



Click the **Next** button to progress to the **Summary** page.

17. Review your settings in the **Cluster overview** screen.



Click any tab in the **Nodes Overview** box to review any configuration.

The screenshot shows the Contrail Command web interface. On the left is a sidebar with a vertical list of steps: STEP 3 (optional) Infrastructure Networks, STEP 4 (optional) Overcloud, STEP 5 (optional) Undercloud Nodes, STEP 6 (optional) Juniphost Nodes, STEP 7 Control Nodes, STEP 8 Orchestrator Nodes, STEP 9 (optional) Compute Nodes, STEP 10 (optional) Contrail Service Nodes, STEP 11 (optional) AppFormix Nodes, STEP 12 Summary, and STEP 13 Provisioning. The main area is titled 'Clusters' and contains a configuration tree on the left and a 'Nodes overview' table on the right.

The configuration tree includes sections like Fabric Management, Contrail configuration, and various node settings. The 'Nodes overview' table has columns for NAME, TYPE, IP ADDRESS, SERVER PROFILE, and ROLES. It lists three nodes: ix-cn-afx-01, ix-cn-afx-flows-01, and ix-cn-cluster-01.

NAME	TYPE	IP ADDRESS	SERVER PROFILE	ROLES
ix-cn-afx-01	physical/virtual	10.1.1.3		AppFormix node
ix-cn-afx-flows-01	physical/virtual	10.1.1.4		AppFormix node
ix-cn-cluster-01	physical/virtual	10.1.1.2		Control node; Compute node; ...

At the bottom of the main area, there are 'Previous' and 'Provision' buttons.

Click the **Provision** button after verifying your settings to provision the cluster.

The cluster provisioning process begins. This provisioning process time varies by environment and deployment. It has routinely taken 90 minutes or more in our testing environments.

18. (Optional) Monitor the provisioning process by logging onto the Contrail Command server and entering the `docker exec contrail-command tail /var/log/contrail/deploy.log` command.

19. When the provisioning process completes, click the **Proceed to Login** option.

You are taken to the Contrail Command login screen.

20. Login to Contrail Command from the web browser.

Enter the following values:

- **Select Cluster:** Select a Contrail Cluster from the dropdown menu. The cluster is presented in the `<cluster-name>-<string>` format. The `<cluster-name>` options should include the cluster that you just created and should match the cluster name assigned in step 10 of this procedure.
- **Username:** Enter the username credential to access Contrail Command. This username was set in the `command_servers.yml` file configured in step 5 of the [“How to Install Contrail Command” on page 22](#) procedure.

- **Password:** Enter the password credential to access Contrail Command. This password was set in the `command_servers.yml` file configured in step 5 of the [“How to Install Contrail Command” on page 22](#) procedure.
- **Domain:** You can often leave this field blank. Contrail Command logs into the **default_domain**—the default domain for all orchestration platforms supported by Contrail Command except Canonical Openstack—when the **Domain** field is empty.

If you are logging into a cluster that includes Canonical Openstack as its orchestration platform, you can enter `admin_domain`—the default domain name for Canonical Openstack—in the **Domain** field if your default domain name was not manually changed.

You can enter the personalized domain name of your cloud network’s orchestration platform in the **Domain** field if you’ve changed the default domain name.

See [“How to Login to Contrail Command” on page 59](#) for additional information on logging into Contrail Command.

21. (Optional. Contrail Insights only) Click the **Contrail Insights** icon on the bottom-left hand corner of the Contrail Command page to open Contrail Insights.

NOTE: This is an **Appformix** icon in Contrail Networking Release 2005 and earlier releases.

If you are not accessing Contrail Command through the fabric network, you might also have to configure an External IP address to access Contrail Insights externally. Navigate to **Infrastructure > Advanced Options > Endpoints** and locate **insights** in the **Prefixes** list. Click the **Edit** button—the pencil icon—and change the **Public URL** field to a usable external IP address.

Contrail Insights Flows is integrated into Contrail Command. See *Contrail Insights Flows in Contrail Command*.

Sample `command_servers.yml` Files for Installing Contrail Command

IN THIS SECTION

- [Minimal `command\_servers.yml` file | 52](#)
- [Complete `command\_servers.yml` File | 53](#)
- [Disaster Recovery and Troubleshooting | 57](#)
- [\[xref target has no title\]](#)

Minimal command_servers.yml file

The following sample file has the minimum configuration that you need when you install Contrail Command.



CAUTION: For security purposes, we strongly recommend creating unique username and password combinations in your environment. Username and password combinations are provided in this example for illustrative purposes only.

```
---
# Required for Appformix and Appformix Flows installations in Release 2003 and
earlier
user_command_volumes:
- /opt/software/appformix:/opt/software/appformix
- /opt/software/xflow:/opt/software/xflow
command_servers:
  server1:
    ip: <IP Address> # IP address of server where you want to install Contrail
    Command
    connection: ssh
    ssh_user: root
    ssh_pass: <contrail command server password>
    sudo_pass: <contrail command server root password>
    ntpserver: <NTP Server address>

    registry_insecure: false
    container_registry: hub.juniper.net/contrail
    container_tag: <container_tag>
    container_registry_username: <registry username>
    container_registry_password: <registry password>
    config_dir: /etc/contrail

    contrail_config:
      database:
        type: postgres
        dialect: postgres
        password: contrail123
      keystone:
        assignment:
          data:
            users:
              admin:
                password: contrail123
```

```

insecure: true
client:
  password: contrail123

```

Complete `command_servers.yml` File

The following sample file has an exhaustive list of configurations and supporting parameters that you can use when you install Contrail Command.



CAUTION: For security purposes, we strongly recommend creating unique username and password combinations in your environment. Username and password combinations are provided in this example for illustrative purposes only.

```

---
# Required for Appformix and Appformix Flows installations in Release 2003 and
earlier
user_command_volumes:
- /opt/software/appformix:/opt/software/appformix
- /opt/software/xflow:/opt/software/xflow

# User defined volumes
#user_command_volumes:
# - /var/tmp/contrail:/var/tmp/contrail

command_servers:
  server1:
    ip: <IP Address>
    connection: ssh
    ssh_user: root
    ssh_pass: <contrail command server password>
    sudo_pass: <contrail command server root password>
    ntpserver: <NTP Server address>

# Specify either container_path
#container_path: /root/contrail-command-051618.tar
# or registry details and container_name
registry_insecure: false
container_registry: hub.juniper.net/contrail
container_name: contrail-command
container_tag: <container_tag>
container_registry_username: <registry username>

```

```

container_registry_password: <registry password>
config_dir: /etc/contrail

# contrail command container configurations given here go to
/etc/contrail/contrail.yml
contrail_config:
  # Database configuration. PostgreSQL supported
  database:
    type: postgres
    dialect: postgres
    host: localhost
    user: root
    password: contrail123
    name: contrail_test
    # Max Open Connections for DB Server
    max_open_conn: 100
    connection_retries: 10
    retry_period: 3s

# Log Level
log_level: debug

# Cache configuration
cache:
  enabled: true
  timeout: 10s
  max_history: 100000
  rdbms:
    enabled: true

# Server configuration
server:
  enabled: true
  read_timeout: 10
  write_timeout: 5
  log_api: true
  address: ":9091"

# TLS Configuration
tls:
  enabled: true
  key_file: /usr/share/contrail/ssl/cs-key.pem
  cert_file: /usr/share/contrail/ssl/cs-cert.pem

```

```

# Enable GRPC or not
enable_grpc: false

# Static file config
# key: URL path
# value: file path. (absolute path recommended in production)
static_files:
    /: /usr/share/contrail/public

# API Proxy configuration
# key: URL path
# value: String list of backend host
#proxy:
#    /contrail:
#    - http://localhost:8082

notify_etcd: false

# VNC Replication
enable_vnc_replication: true

# Keystone configuration
keystone:
    local: true
    assignment:
        type: static
    data:
        domains:
            default: &default
            id: default
            name: default
        projects:
            admin: &admin
            id: admin
            name: admin
            domain: *default
            demo: &demo
            id: demo
            name: demo
            domain: *default
        users:
            admin:
                id: admin
                name: Admin

```

```

        domain: *default
        password: contrail123
        email: admin@juniper.nets
        roles:
        - id: admin
          name: admin
          project: *admin
    bob:
        id: bob
        name: Bob
        domain: *default
        password: bob_password
        email: bob@juniper.net
        roles:
        - id: Member
          name: Member
          project: *demo
    store:
        type: memory
        expire: 36000
    insecure: true
    authurl: https://localhost:9091/keystone/v3

# disable authentication with no_auth true and comment out keystone
configuraion.
    #no_auth: true
    insecure: true

    etcd:
        endpoints:
        - localhost:2379
        username: ""
        password: ""
        path: contrail

    watcher:
        enabled: false
        storage: json

    client:
        id: admin
        password: contrail123
        project_name: admin
        domain_id: default

```



```

    schema_root: /
    endpoint: https://localhost:9091

    compilation:
        enabled: false
        # Global configuration
        plugin_directory: 'etc/plugins/'
        number_of_workers: 4
        max_job_queue_len: 5
        msg_queue_lock_time: 30
        msg_index_string: 'MsgIndex'
        read_lock_string: "MsgReadLock"
        master_election: true

        # Plugin configuration
        plugin:
            handlers:
                create_handler: 'HandleCreate'
                update_handler: 'HandleUpdate'
                delete_handler: 'HandleDelete'

    agent:
        enabled: true
        backend: file
        watcher: polling
        log_level: debug

    # The following are optional parameters used to patch/cherrypick
    # revisions into the contrail-ansible-deployer sandbox. These configs
    # go into the /etc/contrail/contrail-deploy-config.tmpl file
    #
    # cluster_config:
    #     ansible_fetch_url:
    # "https://review.opencontrail.org/Juniper/contrail-ansible-deployer
    # refs/changes/80/40780/20"
    #     ansible_cherry_pick_revision: FETCH_HEAD
    #     ansible_revision: GIT_COMMIT_HASH

```

Disaster Recovery and Troubleshooting

SUMMARY

This section lists commonly seen errors and failure scenarios and procedures to fix them.

Problem

Description: Recovering the Galera Cluster Upon Server Shutdown—In an OpenStack HA setup provisioned using Kolla and OpenStack Rocky, if you shut down all the servers at the same time and bring them up later, the Galera cluster fails.

Solution

To recover the Galera cluster, follow these steps:

1. Edit the `/etc/kolla/mariadb/galera.cnf` file to remove the `wsrep` address on one of the controllers as shown here.

```
wsrep_cluster_address = gcomm://  
#wsrep_cluster_address = gcomm://10.x.x.8:4567,10.x.x.10:4567,10.x.x.11:4567
```

NOTE: If all the controllers are shut down in the managed scenario at the same time, you must select the controller that was shut down last.

2. Docker start mariadb on the controller on which you edited the file.
3. Wait for a couple of minutes, ensure that the mariadb container is not restarting, and then Docker start mariadb on the remaining controllers.
4. Restore the `/etc/kolla/mariadb/galera.cnf` file changes and restart the mariadb container on the previously selected controller.

Problem

Description: Containers from Private Registry Not Accessible—You might have a situation in which containers that are pulled from a private registry named `CONTAINER_REGISTRY` are not accessible.

Solution

To resolve, check to ensure that `REGISTRY_PRIVATE_INSECURE` is set to **True**.

RELATED DOCUMENTATION

[Video: Using Contrail Command to Install Contrail Networking 2005 and Contrail Insights](#)

[Server Requirements and Supported Platforms | 16](#)

Contrail Networking Installation and Upgrade Guide

How to Login to Contrail Command

Contrail Command is a single pane-of-glass GUI that configures and monitors Contrail Networking. You can login to Contrail Command using these instructions.

To login to Contrail Command:

Before you begin:

- Install Contrail Command. See [“How to Install Contrail Command and Provision Your Contrail Cluster” on page 20](#).

1. Open a web browser and enter **`https://<Contrail-Command-Server-IP-Address>:<port-number>`** as the URL.

The default port number is **9091** in most environments. The port number can be reset using the **address:** field in the **server:** hierarchy within the **command_servers.yml** file.

The Contrail Command home screen appears.

2. Select your Contrail cluster from the **Select Cluster** drop-down menu. Leave this field blank if you are logging into Contrail Command to create a cluster.

Your cluster is presented within the drop-down menu in the `<cluster-name>-<string>` format. Your *cluster-name* was defined during the cluster creation process and the *string* is a randomly-generated character string.

In [Figure 4 on page 60](#), the *cluster-name* is **clstr1**.

Figure 4: Contrail Command Login Homepage—Select Cluster

Log in

Select Cluster ⓘ

Select cluster ^

clstr1-1d52fcb0-c87c-4699-af8a-280d980339b6

Username

Password

Domain ⓘ

default

Log in

3. Enter the username and password credentials in the **Username** and **Password** fields. The username and password credentials are set in the **command_servers.yml** file during the Contrail Command installation. See [“How to Install Contrail Command and Provision Your Contrail Cluster” on page 20](#).

For security purposes, we strongly recommend creating unique username and password combinations in your environment. If you didn't change your password in the **command_servers.yml** file, however, the default username to access Contrail Command in most deployments is **admin** and the password is **contrail123**.


4. Enter the domain name of the cluster in the **Domain** field.

You can often leave this field blank. Contrail Command logs into the **default_domain**—the default domain for all orchestration platforms supported by Contrail Command except Canonical Openstack—when the **Domain** field is empty.

You may have to enter a domain in the following use cases:

- Canonical Openstack orchestration. If you are logging into a cluster that includes Canonical Openstack as its orchestration platform, you can enter *admin_domain*—the default domain name for Canonical Openstack—in the **Domain** field if your default domain name was not manually changed.
- You have manually changed the domain name in your cloud network's orchestration platform. You can enter the personalized domain name of your cloud network's orchestration platform in the **Domain** field if you've changed the default domain name.

Figure 5: Contrail Command Login Page Example



The image shows the Contrail Command login page. On the left is a dark sidebar with the Contrail Command logo. The main area is white and contains the following fields:

- Log in** (Section Header)
- Select Cluster** (Dropdown menu): clstr1-1d52fcb0-c87c-4699-af8a-280d980339b6
- Username** (Text input): admin
- Password** (Text input): masked with dots and an eye icon for toggling visibility
- Domain** (Text input): default
- Log in** (Blue button)

5. Press the **Log in** button to enter Contrail Command.

RELATED DOCUMENTATION

[How to Install Contrail Command and Provision Your Contrail Cluster | 20](#)

[Video: Using Contrail Command to Install Contrail Networking 2005 and Contrail Insights](#)

[Contrail Networking Installation and Upgrade Guide](#)

Navigating the Contrail Command UI

IN THIS SECTION

- [Using the Get Started with Contrail Enterprise Multicloud Panel | 62](#)
- [Navigating to pages using the side panel | 64](#)
- [Hiding the side panel | 65](#)
- [Search functionality | 65](#)

- Pinning favorite pages | 66
- Opening external applications | 72
- Using the What's New Panel | 72
- Supported Browsers for Installing Contrail Command | 73

Contrail Networking Release 2003 introduces a redesigned Contrail Command UI. This topic describes how to navigate the UI, some of the new and improved features that the UI offers, and the supported browsers to install Contrail Command.

Using the Get Started with Contrail Enterprise Multicloud Panel

Starting with Contrail Networking Release 2003, you can use the *Get Started with Contrail Enterprise Multicloud* panel in Contrail Command. This *Get Started* panel provides a user-friendly walkthrough of initial Contrail Command configuration tasks. The panel includes **Begin** buttons that allow for quick task initiation and a dynamic tracking mechanism that tracks task progress.

The *Get Started* panel appears automatically when Contrail Command is initially accessed and can always be opened by selecting the **Get Started with Contrail** option in the ? help menu. If you choose to close the panel, it remains closed within Contrail Command—including across login sessions—unless you choose to open the panel by selecting the **Get Started with Contrail** option.

[Figure 6 on page 63](#) shows the *Get Started with Contrail Enterprise Multicloud* panel home screen.

Figure 6: Get Started with Contrail Enterprise Multicloud Home Screen





Figure 7 on page 63 shows how to open the *Get Started with Contrail Enterprise Multicloud* panel within Contrail Command.

Figure 7: Get Started with Contrail Option



The *Get Started* panel has the following limitations:


- if you switch between web browsers or login using incognito mode, the task status within the panel is lost.
- If you clear your web browser cache or delete your web browser's cookies, the task status within the panel is lost.
- If you access Contrail Command from a different web browser or login using incognito mode, the panel opens by default even if it had previously been closed.
- If you clear your web browser cache or delete your web browser's cookies, the panel opens by default even if it had previously been closed.

Navigating to pages using the side panel

All menu options are available in a side panel on the left of the UI. The side panel displays the main menu categories available in Contrail Command. Mouse over each category to view the corresponding second tier menu options. These second tier menu options correspond to pages and you can click the menu option to open the page.

The category as well as the page name can be viewed in the bread crumbs visible on the top banner. You can click on page names in the bread crumbs to navigate to the page, but the category name is not clickable.

Figure 8: Contrail Command Side Panel




Hiding the side panel

You can also hide the side panel from view. To hide the side panel, click the toggle button (≡) on the top left of the banner on any page in Contrail Command. To view the side panel, click the toggle button (≡) again.

Search functionality


If you are unaware of the navigation path to any page, you can use the search functionality to search for the page. The updated Contrail Command UI has a search text box at the top of the left side panel. When you enter a string in the search text box, all the categories containing matching terms are highlighted in a bold font. All other categories are greyed out.

Figure 9: Contrail Command Search Function–Relevant Categories



Mouse over each highlighted category to view the corresponding second tier menu options. While all the menu options of the category are displayed, the pages matching the search input are highlighted in yellow.

Figure 10: Contrail Command Search Function–Relevant Pages




Pinning favorite pages

You can pin frequently visited pages to the favorites category. The left side panel has a **Favorites** category under the search text box and you can add second tier menu options to this category. Initially this category is empty and is grayed out.


- Adding to **Favorites** – To add a page to favorites, mouse over a category in the side panel to display the second tier menu options. Mouse over the menu option and click the pin icon in-line with the page name to add to favorites.

Figure 11: Contrail Command Favorites Function–Add Page



Once pages are added, the **Favorites** category gets enabled and the pinned page is displayed underneath it.


Figure 12: Contrail Command Favorites Function



NOTE: Pinned favorite pages are stored in the Web browser cache in local storage. The existing favorite-page list disappears if you switch between Web browsers, or if you log in under the incognito mode, or if you clear Web browser cache and cookies.


- **Deleting from Favorites** - You can delete pages from the favorites category in any one of the following two ways. You can click the enabled pin icon available in-line with the pinned page in the favorites category to remove the page.

Figure 13: Contrail Command Favorites Function–Remove Page In-line



Alternatively, mouse over the corresponding category of the pinned page to display all the second tier menu options. Mouse over the pinned page and click the enabled pin icon in-line with the page name to remove it from the favorites category.

Figure 14: Contrail Command Favorites Function–Remove Page in Panel




You can also collapse and expand the **Favorites** category by clicking the arrow icon (^ or v) next to the category name.

Opening external applications

To open integrated external applications such as Contrail Insights, click the application name in the footer of the side panel on the left.

NOTE: If no external applications are available, the footer in the side panel is not visible.

Figure 15: Contrail Command Accessing External Applications




Using the What's New Panel

Starting with Contrail Networking Release 2005, you can use the *What's New* panel within Contrail Command to gather a summary list of the new Contrail Networking features in your Contrail Networking release. The *What's New* panel provides a high-level description of each new feature and a *See Release Notes* option that takes you to the Contrail Networking Release Notes for additional feature information.

You can access the *What's New* panel by selecting the **What's New** option in the ? help menu.

Figure 16: What's New Panel



Supported Browsers for Installing Contrail Command

Use the following supported browsers to install Contrail Command.

Contrail Networking Release	Supported Browsers and Versions

2008	• Chrome 80, 81, 83, 84 • Firefox 68, 75, 76, 77, 78 • Microsoft Edge 80, 81, 83, 84 • Opera 66, 67, 68, 69 • Safari 12, 12.1, 13, 13.1
2005	• Chrome 78, 79, 80, 81 • Firefox 68, 72, 73, 74, 75 • Microsoft Edge 18, 79, 80, 81 • Opera 65, 66, 67, 68 • Safari 12, 12.1, 13, 13.1
2003	• Chrome 77, 78, 79, 80 • Firefox 68, 71, 72, 73, 74 • Microsoft Edge 17, 18, 79, 80 • Opera 63, 64, 65, 66 • Safari 12, 12.1, 13, 13.1

Release History Table

Release	Description
2003	Contrail Networking Release 2003 introduces a redesigned Contrail Command UI.

Installing a Contrail Cluster using Contrail Command and instances.yml

Contrail Networking supports deploying Contrail cluster using Contrail Command and the **instances.yml** file. A YAML file provides a concise format for specifying the instance settings.

We recommend installing Contrail Command and deploying your Contrail cluster from Contrail Command in most Contrail Networking deployments. See [“How to Install Contrail Command and Provision Your Contrail Cluster” on page 20](#). You should only use the procedure in this document if you have a strong reason to not use the recommended procedure.

System Requirements

- A VM or physical server with:
 - 4 vCPUs
 - 32 GB RAM
 - 100 GB disk

- Internet access to and from the physical server, hereafter referred to as the Contrail Command server
- (Recommended) x86 server with CentOS 7.6 as the base OS to install Contrail Command

For a list of supported platforms for all Contrail Networking releases, see [Contrail Networking Supported Platforms List](#).

NOTE: Contrail Release 5.1 does not support Contrail Insights deployment from command line with Contrail Cluster instances.yml file.

Before you begin

docker-py Python module is superseded by **docker** Python module. You must remove **docker-py** and **docker** Python packages from all the nodes where you want to install the Contrail Command UI.

```
pip uninstall docker-py docker
```

Configuration

Perform the following steps to deploy a Contrail Cluster using Contrail Command and the **instances.yml** file.

1. Install Docker to pull *contrail-command-deployer* container. This package is necessary to automate the deployment of Contrail Command software.

```
yum install -y yum-utils device-mapper-persistent-data lvm2
```

```
yum-config-manager --add-repo https://download.docker.com/linux/centos/docker-ce.repo
```

```
yum install -y docker-ce-18.03.1.ce
```

```
systemctl start docker
```

2. Download the contrail-command-deployer Docker container image from hub.juniper.net. To download these containers and for access to hub.juniper.net, refer to the *Access to Contrail Registry* topic on the [Contrail software download](#) page. Allow Docker to connect to the private secure registry.

```
docker login hub.juniper.net --username <container_registry_username> --password  
<container_registry_password>
```

Pull contrail-command-deployer container from the private secure registry.

```
docker pull hub.juniper.net/contrail/contrail-command-deployer:<container_tag>
```

Example, for container_tag: 5.1.0-0.38, use the following command:

```
docker pull hub.juniper.net/contrail/contrail-command-deployer:5.1.0-0.38
```

3. Edit the input configuration **instances.yml** file. See [Sample instances.yml File on page 77](#) for a sample **instances.yml** file.
4. Start the **contrail_command_deployer** container to deploy the Contrail Command (UI) server and provision Contrail Cluster using the **instances.yml** file provided.

```
docker run -td --net host -e action=provision_cluster -v  
<ABSOLUTE_PATH_TO_COMMAND_SERVERS_FILE>:/command_servers.yml -v <  
ABSOLUTE_PATH_TO_INSTANCES_FILE>:/instances.yml --privileged --name  
contrail_command_deployer hub.juniper.net/contrail/contrail-command-deployer:<container_tag>
```

The **contrail_command** and **contrail_psql** Contrail Command containers will be deployed. Contrail Cluster is also provisioned using the given **instances.yml** file.

5. (Optional) Track the progress of [6](#).

```
docker logs -f contrail_command_deployer
```

6. Once the playbook execution completes, log in to Contrail Command using <https://Contrail-Command-Server-IP-Address:9091>. Use the same user name and password that was entered in [5](#). Default username is admin and password is contrail123.

NOTE: We strongly recommend creating a unique username and password for Contrail Command. See *Installing Contrail Command* for additional information on creating username and password combinations.

NOTE: Enable subscription on all the RedHat nodes.

```
sudo subscription-manager register --username <USERNAME> --password <PASSWORD>  
sudo subscription-manager attach --pool pool_id  
  
sudo subscription-manager repos --enable=rhel-7-server-rpms  
--enable=rhel-7-server-rh-common-rpms --enable=rhel-ha-for-rhel-7-server-rpms  
--enable=rhel-7-server-extras-rpms
```

Sample instances.yml File

```

global_configuration:
  CONTAINER_REGISTRY: hub.juniper.net/contrail
  CONTAINER_REGISTRY_USERNAME: < container_registry_username >
  CONTAINER_REGISTRY_PASSWORD: < container_registry_password >
provider_config:
  bms:
    ssh_pwd: <Pwd>
    ssh_user: root
    ntpserver: <NTP Server>
    domainsuffix: local
instances:
  bms1:
    provider: bms
    ip: <BMS IP>
    roles:
      config_database:
        config:
      control:
      analytics_database:
      analytics:
      webui:
      vrouter:
      openstack:
      openstack_compute:
  bms2:
    provider: bms
    ip: <BMS2 IP>
    roles:
      openstack:
  bms3:
    provider: bms
    ip: <BMS3 IP>
    roles:
      openstack:
  bms4:
    provider: bms
    ip: <BMS4 IP>
    roles:
      config_database:
        config:
      control:
      analytics_database:
      analytics:

```

```

    webui:
bms5:
  provider: bms
  ip: <BMS5 IP>
  roles:
    config_database:
    config:
    control:
    analytics_database:
    analytics:
    webui:
bms6:
  provider: bms
  ip: <BMS6 IP>
  roles:
    config_database:
    config:
    control:
    analytics_database:
    analytics:
    webui:
bms7:
  provider: bms
  ip: <BMS7 IP>
  roles:
    vrouter:
      PHYSICAL_INTERFACE: <Interface name>
      VROUTER_GATEWAY: <Gateway IP>
    openstack_compute:
bms8:
  provider: bms
  ip: <BMS8 IP>
  roles:
    vrouter:
      # Add following line for TSN Compute Node
      TSN_EVPN_MODE: True
    openstack_compute:
contrail_configuration:
  CLOUD_ORCHESTRATOR: openstack
  CONTRAIL_VERSION: latest or <contrail_container_tag>

RABBITMQ_NODE_PORT: 5673
VROUTER_GATEWAY: <Gateway IP>
ENCAP_PRIORITY: VXLAN, MPLSoUDP, MPLSoGRE

```

```

AUTH_MODE: keystone
KEYSTONE_AUTH_HOST: <Internal VIP>
KEYSTONE_AUTH_URL_VERSION: /v3
CONTROLLER_NODES: < list of mgmt. ip of control nodes >
CONTROL_NODES: <list of control-data ip of control nodes>
OPENSTACK_VERSION: queens
kolla_config:
  kolla_globals:
    openstack_release: queens
    kolla_internal_vip_address: <Internal VIP>
    kolla_external_vip_address: <External VIP>
    openstack_release: queens
    enable_haproxy: "no"      ("no" by default, set "yes" to enable)
    enable_ironic: "no"       ("no" by default, set "yes" to enable)
    enable_swift: "no"        ("no" by default, set "yes" to enable)
    swift_disk_partition_size = 10GB
    keepalived_virtual_router_id: <Value between 0-255>
  kolla_passwords:
    keystone_admin_password: <Keystone Admin Password>

```

RELATED DOCUMENTATION

Installing Contrail Command

Installing a Contrail Cluster Using Contrail Command

[Importing Contrail Cluster Data using Contrail Command | 79](#)

Importing Contrail Cluster Data using Contrail Command

Contrail Networking supports importing of Contrail Cluster data to Contrail Command provisioned using one of the following applications - OpenStack, Kubernetes, VMware vCenter, and TripleO.

System Requirements

- A VM or physical server with:
 - 4 vCPUs
 - 32 GB RAM
 - 100 GB storage

- Internet access to and from the physical server, which is the Contrail Command server.
- (Recommended) x86 server with CentOS 7.6 as the base OS to install Contrail Command.

For a list of supported platforms for all Contrail Networking releases, see [Contrail Networking Supported Platforms List](#).

Access *Container Tags* are located at [README Access to Contrail Registry 20XX](#).

If you need access to Contrail docker private secure registry, e-mail **contrail-registry@juniper.net** for Contrail container registry credentials.

Before you begin

docker-py Python module is superseded by **docker** Python module. You must remove **docker-py** and **docker** Python packages from all the nodes where you want to install the Contrail Command UI.

```
pip uninstall docker-py docker
```

Configuration

Perform the following steps to import Contrail Cluster data.

1. Install Docker to pull *contrail-command-deployer* container. This package is necessary to automate the deployment of Contrail Command software.

```
yum install -y yum-utils device-mapper-persistent-data lvm2
```

```
yum-config-manager --add-repo https://download.docker.com/linux/centos/docker-ce.repo
```

```
yum install -y docker-ce-18.03.1.ce
```

```
systemctl start docker
```

2. Download the *contrail-command-deployer* Docker container image to deploy *contrail-command* (*contrail_command*, *contrail_psql* containers) from *hub.juniper.net*. Allow Docker to connect to the private secure registry.

```
docker login hub.juniper.net --username <container_registry_username> --password  
<container_registry_password>
```

Pull *contrail-command-deployer* container from the private secure registry.

```
docker pull hub.juniper.net/contrail/contrail-command-deployer:<container_tag>
```

Example, for *container_tag*:5.1.0-0.38, use the following command:

```
docker pull hub.juniper.net/contrail/contrail-command-deployer:5.1.0-0.38
```

3. Get the **command_servers.yml** file that was used to bring the Contrail Command server up and the configuration file that was used to provision the Contrail Cluster.

NOTE: "For OpenShift orchestrator use the ose-install file instead of instances.yml file.

4. Start the **contrail-command-deployer** container to deploy the Contrail Command (UI) server and import Contrail Cluster data to Contrail Command (UI) server using the Cluster configuration file provided.

- Import Contrail-Cluster provisioned using a supported orchestrator (OpenStack/Kubernetes/OpenShift/vCenter/Mesos).

```
docker run -td --net host -e orchestrator=<YOUR_ORCHESTRATOR> -e action=import_cluster -v < ABSOLUTE_PATH_TO_COMMAND_SERVERS_FILE>:/command_servers.yml -v < ABSOLUTE_PATH_TO_CLUSTER_CONFIG_FILE>:/instances.yml --privileged --name contrail_command_deployer hub.juniper.net/contrail/contrail-command-deployer:<container_tag>
```

To use the following supported orchestrators, replace **<YOUR_ORCHESTRATOR>** in the command with the options given below.

- For OpenStack, use **openstack**.
- For Kubernetes, use **kubernetes**.
- For Red Hat OpenShift, use **openshift**.

NOTE: You must use **ose-install** file instead of **instances.yml** file.

- For VMware vCenter, use **vcenter**.
- For Mesos, use **mesos**.
- Import Contrail-Cluster provisioned using OSPDirector/TripleO Life Cycle Manager for RedHat OpenStack Orchestration.

Prerequisites:

- **IP_ADDRESS_OF_UNDERCLOUD_NODE** is an Undercloud node IP that must be reachable from the *contrail-command-deployer* node. You must be able to SSH to Undercloud node from the *contrail-command-deployer* node.
- **External VIP** is an Overcloud VIP where OpenStack and Contrail public endpoints are available. **External VIP** must be reachable from Contrail Command node.
- DNS host name for Overcloud external VIP must be resolvable on Contrail Command node. Add the entry in the **/etc/hosts** file.

```
docker run -td --net host -e orchestrator=tripleo -e action=import_cluster -e undercloud=<IP_ADDRESS_OF_UNDERCLOUD_NODE> -e undercloud_user=<user-name> -e undercloud_password=<STACK_USER_PASSWORD_FOR_SSH_TO_UNDERCLOUD> -v <
```

ABSOLUTE_PATH_TO_COMMAND_SERVERS_FILE>:/command_servers.yml --privileged --name contrail_command_deployer hub.juniper.net/contrail/contrail-command-deployer:<container_tag>

NOTE: *undercloud_user* is root by default. Otherwise, the undercloud username is stack in most of the cases.

- Contrail command server must have access to External VIP network to communicate with the configured endpoints.

Run the following commands:

```
ovs-vsctl add-port br0 vlan<externalNetworkVlanID>
tag=<externalNetworkVlanID> -- set interface vlan<externalNetworkVlanID>
type=internal
ip link set dev vlan<externalNetworkVlanID> up
ip addr add <externalNetworkGatewayIP>/<subnetMask> dev
vlan<externalNetworkVlanID>
```

- If you have used domain name for the external VIP, add the entry in the **/etc/hosts** file.

Run the following commands:

```
docker exec -it contrail_command bash
vi /etc/hosts
<externalVIP> <externalVIP'sDomainName>
```

Sample instances.yml file

```
global_configuration:
  CONTAINER_REGISTRY: hub.juniper.net/contrail
  CONTAINER_REGISTRY_USERNAME: < container_registry_username >
  CONTAINER_REGISTRY_PASSWORD: < container_registry_password >
provider_config:
  bms:
    ssh_pwd: <Pwd>
    ssh_user: root
    ntpserver: <NTP Server>
    domainsuffix: local
instances:
  bms1:
    provider: bms
    ip: <BMS1 IP>
    roles:
```

```

    openstack:
bms2:
  provider: bms
  ip: <BMS2 IP>
  roles:
    openstack:
bms3:
  provider: bms
  ip: <BMS3 IP>
  roles:
    openstack:
bms4:
  provider: bms
  ip: <BMS4 IP>
  roles:
    config_database:
    config:
    control:
    analytics_database:
    analytics:
    webui:
bms5:
  provider: bms
  ip: <BMS5 IP>
  roles:
    config_database:
    config:
    control:
    analytics_database:
    analytics:
    webui:
bms6:
  provider: bms
  ip: <BMS6 IP>
  roles:
    config_database:
    config:
    control:
    analytics_database:
    analytics:
    webui:
bms7:
  provider: bms
  ip: <BMS7 IP>

```

```

roles:
  vrouter:
    PHYSICAL_INTERFACE: <Interface name>
    VROUTER_GATEWAY: <Gateway IP>
  openstack_compute:
bms8:
  provider: bms
  ip: <BMS8 IP>
  roles:
    vrouter:
      # Add following line for TSN Compute Node
      TSN_EVPN_MODE: True
    openstack_compute:
contrail_configuration:
  CLOUD_ORCHESTRATOR: openstack
  CONTRAIL_VERSION: latest or <contrail_container_tag>
  CONTRAIL_CONTAINER_TAG: <contrail_container_tag>-queens
  RABBITMQ_NODE_PORT: 5673
  VROUTER_GATEWAY: <Gateway IP>
  ENCAP_PRIORITY: VXLAN,MPLSoUDP,MPLSoGRE
  AUTH_MODE: keystone
  KEYSTONE_AUTH_HOST: <Internal VIP>
  KEYSTONE_AUTH_URL_VERSION: /v3
  CONTROLLER_NODES: < list of mgmt. ip of control nodes >
  CONTROL_NODES: <list of control-data ip of control nodes>
  OPENSTACK_VERSION: queens
kolla_config:
  kolla_globals:
    openstack_release: queens
    kolla_internal_vip_address: <Internal VIP>
    kolla_external_vip_address: <External VIP>
    openstack_release: queens
    enable_haproxy: "no"      ("no" by default, set "yes" to enable)
    enable_ironic: "no"       ("no" by default, set "yes" to enable)
    enable_swift: "no"        ("no" by default, set "yes" to enable)
    keepalived_virtual_router_id: <Value between 0-255>
  kolla_passwords:
    keystone_admin_password: <Keystone Admin Password>

```

RELATED DOCUMENTATION

Installing Contrail Command

Adding a New Compute Node to Existing Contrail Cluster Using Contrail Command


You can add or remove a new node from an existing containerized Contrail cluster.

To add a new compute node to an existing Contrail OpenStack cluster:

1. Login to Contrail Command UI as a super user using credentials *admin* for username and *contrail123* for password.

The default credentials for Contrail Command are *admin* for username and *contrail123* for password. We strongly recommend creating a unique username and password combination. See *Installing Contrail Command*.

2. Click **Servers**.
 - a. Click **Create**.




- b. Enter the required details.

c. Click **Create**.

3. Click **Cluster**.

a. Click **Add** under **Compute Nodes**.



b. Select the required server from **Available Servers** list.

CONTRAIL COMMAND | INFRASTRUCTURE | Cluster

Overview | Cluster Nodes | Subclusters | Cluster Import | Advanced Options >

Assign Compute Nodes

Available servers

Search servers Add all <

HOSTNAME	IP ADDRESS	DISK PARTITION
centos3	10. . .73	>

Assigned Compute nodes

Search servers Remove all


Assign one or more servers

Cancel Assign Nodes

8.79 MB 4.39 MB 0.00 MB 1:48:00 PM 1:56:00 PM 2:03:00 PM 2:08:00 PM 2:15:00 PM Response s... Response ti...

120 60 0 1:47:00 PM 1:51:00 PM 1:55:00 PM 1:59:00 PM 2:03:00 PM 2:07:00 PM 2:11:00 PM 2:15:00 PM Response ti...


c. Click **Assign Nodes**.



Perform the following steps to remove a compute node from an existing Contrail OpenStack cluster.

NOTE: Workloads on the deleted computes must be removed before removing the compute node from the cluster.

1. Login to Contrail Command UI as a super user using credentials *admin* for username and *contrail123* for password.
The default credentials for Contrail Command are *admin* for username and *contrail123* for password. We strongly recommend creating a unique username and password combination for security purposes. See *Installing Contrail Command*.
2. Click **Cluster**.
3. Click **Compute Nodes**.
4. Remove the required compute node.



You can also add a compute node to existing Contrail cluster using **instances.yaml** file. For details, refer to *How to Add a New Compute Node to an Existing Contrail Cluster Using the instances.yaml File*.

How to Deploy Contrail Command and Import a Cluster Using Juju

IN THIS SECTION

- Overview: Deploying Contrail Command with a Contrail Cluster Using Juju | 90
- Preparing the SSL Certificate Authority (CA) for the Deployment | 90
- Deploy Contrail Command and Import a Contrail Cluster Using Juju | 93
- Example: Config.YML File for Deploying Contrail Command with a Cluster Using Juju | 98
- Prerequisites for Contrail Insights and Contrail Insights Flow | 100
- Install Contrail Insights on the Juju Cluster after Contrail Command is Installed | 100
- Install Contrail Insights Flows on the Juju Cluster after Contrail Insights is Installed | 101

You can use this document to deploy Contrail Command and import an existing cluster into Contrail Command using Juju with a single procedure. This procedure can be applied in environments using Canonical Openstack or environments that are running Juju and using Kubernetes for orchestration.

If you are already running Contrail Command in a Canonical Openstack environment and want to import a cluster, see [“Importing a Canonical Openstack Deployment Into Contrail Command” on page 104](#).

Overview: Deploying Contrail Command with a Contrail Cluster Using Juju

Starting in Contrail Release 2005, you can deploy Contrail Command and import a cluster using Juju in a Canonical Openstack environment.

Starting in Contrail Release 2008, you can deploy Contrail Command and import a cluster using Juju in an environment using Kubernetes orchestration.

This document makes the following assumptions about your initial environment:

- Juju is already running in your environment, and your environment is either a Canonical Openstack deployment or a deployment using Kubernetes orchestration.
- Contrail Networking Release 2005 or later is running if you are operating a Canonical Openstack deployment.

Contrail Networking Release 2008 or later is running if you are operating an environment using Kubernetes orchestration.

See [Contrail Networking Supported Platforms](#) for information on the supported software components for any Contrail Networking release.

- A Juju controller is configured and reachable.
- Contrail Command is not running.

Preparing the SSL Certificate Authority (CA) for the Deployment

A base64-encoded SSL Certificate Authority (CA) for the Juju controller is required to deploy Contrail Command with an existing cluster in a Canonical Openstack or Kubernetes environment.

There are multiple ways to generate a base64-encoded SSL CA. You can use this procedure or a more familiar procedure to generate your base64-encoded SSL CA.

To create a base64-encoded SSL CA:

1. From the Juju jumphost, enter the **juju show-controller** command and locate the certificate output in the **ca-cert:** hierarchy.

```
$ juju show-controller
jc5-cloud:
  details:
    ...<output removed for readability>...
    ca-cert: |
      -----BEGIN CERTIFICATE-----
      MIIErTCCAxWgAwIBAgIVAKRPIub8Q7imJ2+T2U8AK4th0ss7MA0GCSqGSIb3DQEB
      CwUAMG4xDALBgNVBAoTBGp1anUxLjAsBgNVBAMMJWp1anUtZ2VuZXJhdGVkIENB
      IGZvciBtb2RlbCAianVqdS1jYSIxLTArBgNVBAUTJDl0ZDJjODg0LTllYWYtNDU2
      Ni04NTA0LWJkZGYxZWJiYTgzYjAeFw0yMDA0MTUwMzE2MzdaFw0zMDA0MjIwMzE2
      MzVaMG4xDALBgNVBAoTBGp1anUxLjAsBgNVBAMMJWp1anUtZ2VuZXJhdGVkIENB
      IGZvciBtb2RlbCAianVqdS1jYSIxLTArBgNVBAUTJDl0ZDJjODg0LTllYWYtNDU2
      Ni04NTA0LWJkZGYxZWJiYTgzYjCCAAIwDQYJKoZIhvcNAQEBBQADggGPADCCAYoC
      ggGBAL/7d3JtNcHW6ue6y0eKv0lSDhXgGs4vYLD00QzliMyw39+BytB4XY+05EBg
      A5JKfYV+u8xXL0meLvH+4yE87cwR0bsT1WYFCDFVTiGSeSN3w+2UJxHWwuAubDl7
      zfAKnGgIzq/KZJJimxa6Yuqw5isCxffu3fQz+H5ULSpLCpFxxAq38VjrW7FnjEm1
      c4fFlBf07LU0QbXSiS0gxar01DQE2IQv4mfIAFvJgT/5UKJYuGEX3NH9DerYqjJa
      NchyGMkXgyBj3YVec8bFE4+erDMISBvJHBMwyx74PTDQys+KlfnXptup5FH/FwBb
      9ZRBAD99c0f0VW6moNxoAKKhrGVZt1w7CxxwvGRZnWUezthwoHI8yFqBvkT+Lq6Nd
      jvLEv1DQ+3zmMfhz/emRD1D0QQfn3mQhSk40Nd03kw/B8bH0IXmgIgNbv48g0Ac7
      /hQ002moDxrLkCZNN0fVg0KvonDjbSo5YNCH/7fleacmQN3Mug3wXp9kYh7rKDHw
      6pkQQwIDAQAB0IwQDA0BgNVHQ8BAf8EBAMCAQwDwYDVR0TAAQ/BAUwAwEB/zAd
      BgNVHQ4EFgQUcGE6bMiGsQQyiDYKBL+txAfeFAkwDQYJKoZIhvcNAQELBQADggGB
      AG7pivQhJVNSCbg+9jVy3owhg/P0np2sewD1t8BM0kPTPgAa/37vrp4KSPdNXKZz
      hnFzBXkL8jBUP0qg2Vfy9iqlgXNdVadb4Ijk440hLwWNGUiZwL2nNbnUL7NnTeh
      jqZaIb60e2y1ByNrQweVM085qdrYJCelf9Wh9fYdtOfx4Ty0Mg+ZqPqmvTR08yTx
      KOupywxmezbjhEaaILXo9kouU4UV2gAIdYiHfvsbTaLkwbYeNgvVE5WAan8HuQqb
      YVnvxggIN45UgEgqGUHEgcj9tHgssfbnX3f2sCb0JkXL2cv7D+wK7hvUCS5tKS6H
      6070oXxfimFBdSZQuuqhqiMYafnRo48Q2oCyQn1Q+g/qG+GYxmujiigoiYS1srV
      mIUaJQUGHtgXvyZGJFivQiAzImQCylq1iyz77Da3myDRX0i0dauu5MACn5i9cgu9
      w7/MD2xR3kKMAY3b4y+pP7CKbEJ6UDswLyAQUkwPyELi1r82vGh6CasinnGaUhk+
      zg==
      -----END CERTIFICATE-----
    ...<additional output removed for readability>...
```

2. Copy the contents of the SSL CA into the cert.pem file.

Copy and paste options vary by user interface. The SSL CA content—all highlighted text from step 1 starting at the beginning of the -----BEGIN CERTIFICATE----- line and ending at the end of the -----END CERTIFICATE----- line—should be the only content in the cert.pem file.

Confirm that leading white spaces are not added to the SSL CA after copying the SSL CA into the cert.pem file. These leading white spaces are introduced by some user interfaces—often at the start of new lines—and will cause the SSL CA certification to be unusable. If leading whitespaces are added to

the SSL CA after it is copied into the `cert.pem` file, manually delete the whitespaces before proceeding to the next step.

3. Generate the `cert.pem` file into base64-encoded output.

You can generate the `cert.pem` file into base64-encoded output without saving the file contents by entering the following command:

```
cat cert.pem | base64
```

You can also generate the base-64 encoded output and save the SSL CA contents into a separate file.

In this example, the base64-encoded output is generated and a new file containing the output—**`cert.pem.b64`**—is saved.

```
cat cert.pem | base64 > "cert.pem.b64"
```

The SSL CA in the **`cert.pem.b64`** file is now a base64-encoded SSL CA.

The base64-encoded SSL CA will be entered as the *juju-CA-certificate* variable in [“Deploy Contrail Command and Import a Contrail Cluster Using Juju” on page 93](#).

Deploy Contrail Command and Import a Contrail Cluster Using Juju

To deploy Contrail Command and import a Contrail cluster into Contrail Command:

1. From the Juju jumphost, deploy Contrail Command using one of the following command strings:

```
juju deploy cs:~juniper-os-software/contrail-command --constraints
tags=<machine-tag> --config docker-registry=<registry-directory> --config
image-tag=<image-tag>
```

```
juju deploy cs:~juniper-os-software/contrail-command --to <machine-name> --config
docker-registry=<registry-directory> --config image-tag=<image-tag>
```

where:

- *machine-name*—the name of the machine instance in Juju that will host Contrail Command.

The IP address of this machine—which can be obtained by entering the **juju status** command—is used to access Contrail Command from a web browser after the installation is complete.

- *registry-directory*—the directory path to the Contrail Networking registry.

This *registry-directory* path can be obtained from Juniper Networks. Contact contrail-registry@juniper.net for information on accessing the Juniper registry.

- *image-tag*—the image tag for your target Contrail release.

The image tag is used to identify your Contrail Networking image within the registry. You can retrieve the image tag for any Contrail Release 20xx image from [README Access to Contrail Networking Registry 20XX](#).

2. Create a juju relation between the Contrail Command charm and the Contrail Controller charm:

```
juju add-relation contrail-command contrail-controller
```

3. Import the Contrail cluster into Contrail command:

- a. Create a **config.yaml** file with the following parameters:

```
$ cat config.yaml
juju-controller: juju-controller-ip
juju-controller-password: password
juju-ca-cert: |
    juju-CA-certificate
juju-model-id: juju-model-id
juju-controller-user: juju-controller-user
```

The command variables:

- *juju-controller-ip*—The IP address of the Juju controller.

You can retrieve the *juju-controller-ip* from the **juju show-controller** command output:

```
username@contrail-ci:~$ juju show-controller
jc5-cloud:
  details:
    ...<output removed for readability>...
    api-endpoints: [10.102.72.40:17070]
    ...<output removed for readability>...
```

- *password*—The password for Juju controller access.

You can set the password for Juju controller access using the **juju change-user-password** command.

- *juju-CA-certificate*—The base64-encoded SSL Certificate Authority (CA) for the Juju controller.

The *juju-CA-certificate* is the base64-encoded SSL CA created in [“Preparing the SSL Certificate Authority \(CA\) for the Deployment” on page 90](#).

See [“Example: Config.YML File for Deploying Contrail Command with a Cluster Using Juju” on page 98](#) for a sample *juju-CA-certificate* entry.

- *juju-model-id*—The universally unique identifier (UUID) assigned to the model environment that includes the Contrail Networking cluster..

You can retrieve the *juju-model-id* from the **juju show-controller** command output:

```
$ juju show-controller
jc5-cloud:
  ...<output removed for readability>...
  models:
    default:
```

```
model-uuid: 4a62e0b0-bcfe-4b35-8db7-48e55f439217
...<output removed for readability>...
```

- **juju_controller_user**—(Optional) The username of the user with Juju controller access.

The **admin** username is used by default if no user with Juju controller access is configured.

See “[Example: Config.YML File for Deploying Contrail Command with a Cluster Using Juju](#)” on [page 98](#) for a sample **config.yaml** configuration for this deployment.

- Save the **config.yaml** file.
- Import the Contrail cluster with the parameters defined in the **config.yaml** file:

```
juju run-action contrail-command/0 import-cluster --params config.yaml
Action queued with id: 1
```

- Check the cluster import status.

You can check the import status by entering the **juju show-action-status *action-ID*** and **juju show-action-output *action-ID* | grep result** commands.

The *action-ID* is assigned immediately after entering the **juju run-action** command in the previous step.

The cluster import is complete when the **status** field output in the **juju show-action-status *action-ID*** command shows **completed**, or when the **result** field in the **juju show-action-output *action-ID* | grep result** indicates **Success**.

Examples:

```
juju show-action-status 1
actions:
- action: import-cluster
completed at: "2020-04-03 12:49:55"
id: "60"
status: completed
unit: contrail-command/19
```



```
juju show-action-output 1 | grep result
results:
  result: Success
```

4. Login to Contrail Command by opening a web browser and entering `https://<juju-machine-ip-address>:<port-number>` as the URL.

The `<juju-machine-ip-address>` is the IP address of the machine hosting Contrail command that was specified in 1. You can retrieve the IP address using the `juju status` command:

NOTE: Some `juju status` output removed for readability.

```
juju status
Unit                                Workload Agent Machine Public address
contrail-command/0*                active   idle   3         10.0.12.40
```

The `port-number` typically defaults to 9091 or 8079. You can, however, configure a unique port number for your environment using the `command_servers.yml` file.

Enter the following values after the Contrail Command homescreen appears:

- **Select Cluster:** Select a Contrail Cluster from the dropdown menu. The cluster is presented in the `<cluster-name>-<string>` format.
- **Username:** Enter the username of the Juju keystone user.
- **Password:** Enter the password of the Juju keystone user.
- **Domain:** If you are running Juju in a Canonical Openstack environment, enter `admin_domain`—the default domain name for Canonical Openstack— if you haven't established a unique domain in Canonical Openstack. Enter the name of your domain if you have created a unique domain.

If you are running Juju in a Kubernetes environment, you can leave this field blank unless you've established a unique domain name in Kubernetes. Enter the name of your domain if you have created a unique domain.

Figure 17 on page 98 illustrates an example Contrail Command login to complete this procedure.

Figure 17: Contrail Command Login Example—Cluster in Environment using Canonical Openstack

See [“How to Login to Contrail Command” on page 59](#) for additional information on logging into Contrail Command.

Example: Config.YML File for Deploying Contrail Command with a Cluster Using Juju

This sample **config.yml** file provides a representative example of a configuration that could be used to deploy Contrail Command with Contrail clusters in an environment running Juju.

See [“Deploy Contrail Command and Import a Contrail Cluster Using Juju” on page 93](#) for step-by-step procedures to create this **config.yml** file and [“Preparing the SSL Certificate Authority \(CA\) for the Deployment” on page 90](#) for instructions on generating the *juju-ca-cert* in the required base64-encoded format.

This sample **config.yml** file does not contain the **juju-controller-user:** field to specify a user with Juju controller access, so the default **admin** username is used.



CAUTION: The password **password** is used in this example for illustrative purposes only.

We strongly recommend creating a unique password that meets your organization's security requirements for your environment.

```
$ cat config.yaml
juju-controller: 10.102.72.40
juju-ca-cert: |
  LS0tLS9CRUdJTiBDRVJUSUZJQ0FURS0tLS0tck1JSUVyVENDQXhXZ0F3SUJBZ0lwQUtSUEl1Yjhr
  N2ltSjIrVDJVOEFLNHRoT3NzN01BMEdDU3FHU0liM0RRRUlKQ3dVQU1HNHhEVEFMQmdOVkJBb1RC
  R3AxyW5VeExqQXNCZ05WQkFNTUpXcDFhbV0wJjJWdVpYSmhkR1ZrSUV0QgpJR1p2Y2lCdGIyUmxi
  Q0FpYW5WcWRTMWpZU0l7TFRBckJnTlZCQVVUSkRJMFPESmpPRGcwTFRsbFlXWXRORFUyCk5pMDRO
  VEEwTFdKa1pHWXhaV0ppWVRnellqQwVGdzB5TURBME1UVXdNekUyTXpkYUZZ3MHPNREEwTwpJd016
  RTIKTXpWYU1HNHhEVEFMQmdOVkJBb1RCR3AxyW5VeExqQXNCZ05WQkFNTUpXcDFhbV0wJjJWdVpY
  SmhkR1ZrSUV0QgpJR1p2Y2lCdGIyUmxiQ0FpYW5WcWRTMWpZU0l4TFRBckJnTlZCQVVUSkRJMFPES
  mpPRGcwTFRsaFlXWXRORFUyCk5pMDROVEEwTFdKa1pHWXhaV0ppWVRnellqQ0NBYU13RFFZSktv
  WklodmNOQVFFQkJRQURnZ0dQURDQ0FZb0MKZ2dHQkFMLzdkM0p0TmNIVzZ1ZTZ5T2VLdk9sU0Ro
  eGdHczR2WUXETzBREmxJTXlXMzkrQnl0QjRYWSSwNUVCZwpBNUpLZlLWk3U4eFhMMG1lTHZoKzR5
  RTg3Y3dST2JwVDFXWUZDREZWVGLHU2VTTjN3KzJVSnhIV3d1QXViRGw3CnpmQUtuR2dJenEvS1pK
  SmlteGE2WXVxdzVpc0N4ZmZ1M2ZReitINVVsU3BMQ3BGeHZBcTM4VmpyVzdGbmpFbTEKYzRmRmxC
  ZjA3TFVPCUJ4U0lTMGd4YXJPMURRRTJJUXY0bWZJQUZ2SmdULzVVS0pZdUdFWDN0SDLEZXJZCwP
  YQp0Y2h5R01rWGD5QmozWVZlYzhIRkU0K2VyRE1JU0J2SkhCTXd5eDc0UFREUXlZK0tsZk5YCHR1
  cDVGSC9Gd0JiCjlaUkJBdK5YzBmMFZXNm1vTnhvQWtLaHJHVlp0MXc3Q3h3dmdSWM5XVW6dGh3
  b0hJ0HLGcUJ2a1QrbHE2TmQKanZMRXYxRFErM3ptTWZoei9lbVJEMURPUVFmbjNtUWhTazQwTmRA
  M2t3L0I4YkhPSVhtZ0lnTmJ2NDhnMEFjNwovAFFPMDJtb0R4ckxrQ1p0TjBmVmdPS3ZvbkrQYlNv
  NVl0Q0gVn2ZsZWfjbVFOM011ZzN3WHA5a1lon3JLREh3CjZwa1FRd0LEQVFBQm8wSXdrREFPQmd0
  VkhR0EJBZjhFQkFNQ0FfUXdEd1LEVlIiwVEFRSC9CQV3QXdfQ096QWQKQmd0VkhRNEVGZ1FVY0dF
  NmJNaUdzUVF5aURZS0JpK3R4QWZlRkFrD0RRWUpLb1pJaHJjTkFRRUxCUUFEZ2dHQgpBRzdwaXZR
  aEpWTlNDYkcrOWpweTNad2hnL1BPbnAyc2V3RDF00EJNT2tQVFBnQWEvMzd2cnA0S1NQZE5YS1p6
  CmhuRnpCWGtMOGpCVVAwcWcyVmZ5OWlxbGdYTMwRQWRiNElqazQ0T2hsd1d0R1VpWndsMm50YnZu
  VUw3Tm5UZWgKanFaYUliNk9lMnkxQnl0clF3ZVZNTZg1cWRyWUpDZWxmOVdoOWZZZHRvZng0VHlP
  TwcrWnFQcw12VFJPOHlUeApLT3VweXd4bWV6YmpoRWFhSUxYbzlr3VVNFVWmmdBSWRZaUhmndNi
  VGFMa1diWWV0Z3Z2RTVXQWfu0Eh1UXFiClLWbnZ4Z2dJTjQ1VWdFZ3FHVUhfZ2NqOXRIZ3NzZmJu
  WDNmNmNDYk9Ka1hMMmN2N0Qrd0s3aHZVQ1M1dEtTNkgKNk83T29YeGZpbUZCZFNauXV1cWxheWlN
  WWFmbLjVNDhRMM9DeRFuMVERZy9xRyTHWxhtdWpJawdvaVlTMXNyVgptSVVhSlFVR0h0Z1h2eVpH
  SkZJdlFpQXpJbVFDewxMWl5ejc3RGEzbXlEULgwaTBkYXV1NU1BQ241aTljZ3U5Clc3L0lEMnhs
  M2tLTUFZM2I0eStwUddDS2JFSjZVRHN3THlBUVVRd1B5ZUxpMXI4MnZHaDZDYXNpbm5HYVVoaysK
  eac9PQotLS0tLUVORCBDRVJUSUZJQ0FURS0tLS0tCg==
juju-model-id: 4a62e0b0-bcfe-4b35-8db7-48e55f439217
juju-controller-password: password
```

Prerequisites for Contrail Insights and Contrail Insights Flow

Contrail Networking Release 2011 supports installing Contrail Insights and Contrail Insights Flows on a Juju cluster after Contrail Networking and Contrail Command are installed. The following prerequisites apply.

docker, **python2.7**, **python-pip** must be installed on the Contrail Insights node and Contrail Insights Flows node.

To install the Docker engine, you need the 64-bit version of one of these Ubuntu versions:

- Ubuntu Groovy 20.10
- Ubuntu Focal 20.04 (LTS)
- Ubuntu Bionic 18.04 (LTS)
- Ubuntu Xenial 16.04 (LTS)

Docker Engine is supported on x86_64 (or amd64), armhf, and arm64 architectures. For more information, see <https://docs.docker.com/engine/install/ubuntu/>.

To install **python 2.7** and **python-pip** run the following commands:

```
sudo apt install python2.7
sudo apt install python-pip
```

If you are running the playbooks as **root** user then this step can be skipped. As a non-root user (for example, “ubuntu”), the user “ubuntu” needs access to the docker user group. The following command adds the user to the docker group:

```
sudo usermod -aG docker ubuntu
```

For more information, see “[Contrail Insights Installation for OpenStack in HA](#)” on page 499.

Install Contrail Insights on the Juju Cluster after Contrail Command is Installed

NOTE: Appformix and Appformix Flows were renamed Contrail Insights and Contrail Insights Flows. The Appformix naming conventions still appear during product usage, including within these directory names.

To install Contrail Insights on the Juju Cluster:

1. Copy the Contrail Insights and Contrail Insights Flows installation directories to the **/opt/software/appformix/** and **/opt/software/xflow** directories inside the Contrail Command container, if not already present.

```
docker run -v /opt/software/appformix:/opt/software/appformix
svl-artifactory.juniper.net/contrail-nightly/appformix/appformix/contrail-insights-ansible:<Contrail
Insights Version>

docker run -v /opt/software/flow:/opt/software/flow
svl-artifactory.juniper.net/contrail-nightly/appformix/flows/contrail-insights-flows-ansible:<Contrail
Insights Flow Version>
```

For example **<Contrail Insights Version> = 3.3.0-a8**.

2. Create the following two inventory files:

```
docker exec -it contrail_command bash
vi /opt/software/appformix/inventory/group_vars/all
vi /opt/software/appformix/inventory/hosts
```

3. Run the following commands to install Contrail Insights in HA mode:

```
cd /usr/share/contrail/appformix-ansible-deployer/appformix/
. venv/bin/activate

cd /opt/software/appformix/
ansible-playbook -i inventory --skip-tags=install_docker
contrail-insights-ansible/appformix_openstack_ha.yml -v
```

Install Contrail Insights Flows on the Juju Cluster after Contrail Insights is Installed

Disclaimer: Official installation method for installation is using the Contrail-Command UI.

contrail-ansible-deployer installs all packages needed for Contrail Insights and Contrail Insights Flows. **appformix-ansible-deployer** creates inventory files for the installation. There are many variables set in the inventory files for specific releases, so setting them manually is prone to errors.

To install Contrail Insights Flows on the Juju Cluster:

1. Log in to the **contrail-command** container:

```
docker exec -it contrail_command bash
```

2. Run the following two commands:

```
cd /usr/share/contrail/appformix-ansible-deployer/xflow
source venv/bin/activate
```

3. Run one of the following commands dependent on your Contrail Networking Release version.

If you are running a Contrail Networking Release later than 2005:

```
bash deploy_contrail_insights_flows.sh <path-to-instances-yml>/instances.yml
--cluster-id <cluster_id>
```

If you are running a Contrail Networking Release earlier than 2005:

```
bash deploy_xflow.sh <path-to-instances-yml>/instances.yml
```

If you are running a Contrail Networking Release earlier than 2005, add the following snippet to the end of the existing **instances.yml** before running the **deploy_contrail_insights_flows.sh** or **deploy_xflow.sh**.

Example **instances.yml** snippet for in-band configuration:

```
global_configuration:
  CONTAINER_REGISTRY: hub.juniper.net/contrail
  CONTAINER_REGISTRY_USERNAME: < container_registry_username >
  CONTAINER_REGISTRY_PASSWORD: < container_registry_password >
provider_config:
  bms:
    ssh_pwd: <Root Pwd>
    ssh_user: root
    ntpserver: <NTP Server>
    domainsuffix: local
instances: < under existing hierarchy >
  a7s33:
    ip: 10.84.30.201
    provider: bms
    roles:
      appformix_flows:
        telemetry_in_band_interface_name: enp4s0f0
xflow_configuration:
```

```
clickhouse_retention_period_secs: 7200
loadbalancer_collector_vip: 30.1.1.3
telemetry_in_band_cidr: 30.1.1.0/24
loadbalancer_management_vip: 10.84.30.195
telemetry_in_band_vlan_id: 11
```

Example **instances.yml** snippet for out-of-band configuration:

```
global_configuration:
  CONTAINER_REGISTRY: hub.juniper.net/contrail
  CONTAINER_REGISTRY_USERNAME: < container_registry_username >
  CONTAINER_REGISTRY_PASSWORD: < container_registry_password >
provider_config:
  bms:
    ssh_pwd: <Root Pwd>
    ssh_user: root
    ntpserver: <NTP Server>
    domainsuffix: local
instances: < under existing hierarchy >
  a7s33:
    ip: 10.84.30.201
    provider: bms
    roles:
      appformix_flows:
xflow_configuration:
  clickhouse_retention_period_secs: 7200
  loadbalancer_collector_vip: 10.84.30.195
```

4. Add the collector nodes:

```
bash deploy_insights_flows.sh <instance_file> --skip-provision --cluster-id
<cluster_id>
```

Release History Table

Release	Description
2011	Contrail Networking Release 2011 supports installing Contrail Insights and Contrail Insights Flows on a Juju cluster after Contrail Networking and Contrail Command are installed.
2008	Starting in Contrail Release 2008, you can deploy Contrail Command and import a cluster using Juju in an environment using Kubernetes orchestration.
2005	Starting in Contrail Release 2005, you can deploy Contrail Command and import a cluster using Juju in a Canonical Openstack environment.

RELATED DOCUMENTATION

[How to Login to Contrail Command | 59](#)

[Importing a Canonical Openstack Deployment Into Contrail Command | 104](#)

Importing a Canonical Openstack Deployment Into Contrail Command

IN THIS SECTION

- [Overview: Canonical Openstack Deployment into Contrail Command | 105](#)
- [Importing Canonical Openstack Into Contrail Command | 105](#)

This document provides the steps needed to import a Canonical Openstack deployment into Contrail Command.

This procedure assumes that Contrail Command is already running in your Contrail Networking environment that is using Canonical Openstack as its orchestration platform. See [“How to Deploy Contrail Command and Import a Cluster Using Juju” on page 89](#) if you’d like to deploy Contrail Command and import the Contrail cluster into Contrail Command in an environment using Contrail Networking and Canonical Openstack.

Overview: Canonical Openstack Deployment into Contrail Command

Starting in Contrail Networking Release 2003, Canonical Openstack deployments can be managed using Contrail Command.

This document provides the steps needed to import a Canonical Openstack deployment into Contrail Command. Contrail Command can be used to manage the Canonical Openstack deployment after this procedure is complete.

This document makes the following assumptions about your environment:

- A Canonical Openstack deployment managed by Contrail Networking is already operational.
- Contrail Command is running in your environment. See *Installing Contrail Command*.
- Contrail Command has access to the Juju jumphost and the Juju cluster.

Canonical Openstack is imported into Contrail Command using Juju in this procedure.

Importing Canonical Openstack Into Contrail Command

To import Canonical Openstack into Contrail Cloud:

1. Install and start the Docker Engine.

There are multiple ways to perform this step. In this example, Docker Community Edition version 18.03 is installed using **yum install** and **yum-config-manager** commands and started using the **systemctl start docker** command.

```
yum install -y yum-utils device-mapper-persistent-data lvm2
yum-config-manager --add-repo
https://download.docker.com/linux/centos/docker-ce.repo
yum install -y docker-ce-18.03.1.ce
systemctl start docker
```

2. Retrieve the **contrail-command-deployer** Docker image by logging into *hub.juniper.net* and entering the **docker pull** command.

```
docker login hub.juniper.net --username <container_registry_username> --password
<container_registry_password>
docker pull hub.juniper.net/contrail/contrail-command-deployer:<container_tag>
```

where *<container_tag>* is the container tag for the Contrail Command (UI) container deployment for the release that you are installing.

The `<container_tag>` for any Contrail Release 20xx image can be found in [README Access to Contrail Registry 20XX](#).

3. Update the **config.yml** configuration file on the Contrail Command server.

The configuration of the **config.yml** file is unique to your environment and complete documentation describing all **config.yml** configuration options is beyond the scope of this document.

The following configuration parameters must be present in the **config.yml** file to support Canonical Openstack in Contrail Command:

- `ntpserver: <NTP_IP>`.

The `NTP_IP` variable is the IP address of the NTP server.

- `vrouter_gateway: <VROUTER_GATEWAY_IP>`

The `VROUTER_GATEWAY_IP` variable is the IP address of the vRouter gateway. The `vrouter_gateway` parameter can be left empty, but it must be present.

- `container_registry: <CONTAINER_REGISTRY>`

The `CONTAINER_REGISTRY` variable is the path to the container registry. The `CONTAINER_REGISTRY` is `hub.juniper.net/contrail` in most deployments.

- `container_tag: <COMMAND_BUILD_TAG>`

The `COMMAND_BUILD_TAG` variable is the Contrail Command (UI) container deployment for the release that you are installing. For any Contrail Release 20xx image, you can retrieve this value from [README Access to Contrail Registry 20XX](#).

- `contrail_container_tag: <CONTRAIL_BUILD_TAG>`

The `CONTRAIL_BUILD_TAG` variable is the Contrail build container for the release that you are installing. For any Contrail Release 20xx image, you can retrieve this value from [README Access to Contrail Registry 20XX](#).

4. Run the Contrail Command deployer.

```
docker run -t --net host -e action=import_cluster -e orchestrator=juju -e
juju_controller=<juju_controller>
[-e juju_model=<juju_model_name> ]
[-e juju_controller_user=<juju_controller_user>]
[-e juju_controller_password=<juju_controller_password>]
[-e delete_db=<delete_db>]
[-e persist_rules=<persist-rules>]
-v <config_file>:/cluster_config.yml --privileged --name contrail_command_deployer
<CCD_image>
```

In the following example, Contrail Command is deployed from the Juju jump host at 172.31.40.101.

```
docker run -td --net host --privileged -e action=import_cluster -e
orchestrator=juju -e juju_model=controller -e juju_controller=172.31.40.101 -e
juju_controller_user=ubuntu -e juju_controller_password=password -v
/home/ubuntu/contrail-command-deployer/config.yaml:/cluster_config.yml -v
/root/.ssh:/root/.ssh contrail-command-deployer:latest
```

The command variables:

- **juju_controller**—(Required) The IP address of the Juju jump host. We define the Juju jump host in this context as the device that has installed the Juju CLI and is being used to run Juju commands. The Contrail Command server must have access to the Juju jump host at this IP address.
- **config_file**—(Required) The path to the configuration file. This configuration file was created in the previous step of this procedure.
- **CCD_image**—(Required) The Contrail Command deployer image.
- **delete_db**—(Optional) Specifies whether the PostgreSQL database is deleted during the process. The PostgreSQL database is deleted by default. Enter *no* in this field if you do not want the PostgreSQL database deleted.
- **persist-rules**—(Optional) Specify whether IP rules remain persistent across reboots.
- **juju_model_name**—(Optional) The name of the Juju model. The name can be retrieved by entering the *juju show-models* command.
- **juju_controller_user**—(Optional) The username of the Juju user on the Juju jump server.
- **juju_controller_password**—(Optional) The password for the Juju user on the Juju jumpbox. This password is used if no SSH keys have been installed.

5. (Optional) Track the progress of step 6.

```
docker logs -f contrail_command_deployer
```

6. Verify that the Contrail Command containers are running:

```
[root@centos254 ~]# docker ps -a
```

CONTAINER ID	IMAGE	<trimmed>	STATUS	<trimmed>	NAMES
2e62e778aa91	hub.juniper.net/...		Up	<trimmed>	contrail_command
c8442860e462	circleci/postgre...		Up	<trimmed>	contrail_psql
57a666e93d1a	hub.juniper.net/...		Exited	<trimmed>	contrail_command_deployer

The **contrail_command** container is the GUI and the **contrail_psql** container is the database. Both containers should have a STATUS of **Up**.

The **contrail-command-deployer** container should have a STATUS of **Exited** because it exits when the installation is complete.

- 7. Open a web browser and enter *https://<Contrail-Command-Server-IP-Address>:8079* as a URL. The Contrail Command home screen appears.

Choose token from the drop-down menu. Enter the username and password combination to Juju as the credentials, and use **admin_domain** as the domain.

Release History Table

Release	Description
2003	Starting in Contrail Networking Release 2003, Canonical Openstack deployments can be managed using Contrail Command.

RELATED DOCUMENTATION

| *Installing Contrail Command*

Upgrading Contrail Software

IN THIS CHAPTER

- Upgrading Contrail Networking using Contrail Command | 109
- Upgrading Contrail Command using Backup Restore Procedure | 112
- How to Perform a Zero Impact Contrail Networking Upgrade using the Ansible Deployer | 114
- Updating Contrail Networking using the Zero Impact Upgrade Process in an Environment using Red Hat Openstack | 121
- Updating Contrail Networking using the Zero Impact Upgrade Procedure in a Canonical Openstack Deployment with Juju Charms | 130
- Upgrading Contrail Networking using contrail-ansible Deployer | 136
- Upgrading Contrail Networking using In-Place Upgrade Procedure | 138
- Upgrading Contrail Networking Release 19xx with RHOSP13 to Contrail Networking Release 2011 with RHOSP16.1 | 139

Upgrading Contrail Networking using Contrail Command

Use the following procedure to upgrade Contrail Networking using Contrail Command.

The procedure supports incremental model and you can use it to upgrade from Contrail Networking Release N-1 to N.


1. Create snapshots of your current configurations and upgrade Contrail Command. See ["Upgrading Contrail Command using Backup Restore Procedure" on page 112](#).
2. Login to Contrail Command without logging into a Contrail cluster.

In most environments, Contrail Command is accessed by entering **https://<Contrail-Command-Server-IP-Address>:9091** in your web browser..

Leave the **Select Cluster** field blank. If the field is populated, use the drag-down menu to delete the pre-populated cluster selection.

Enter your username and password credentials and click **Log in**

If you have any issues logging into Contrail Command, see [“How to Login to Contrail Command”](#) on page 59.



The image shows the Contrail Command login interface. On the left is a large dark panel with the Contrail Command logo (a blue cloud icon with a dashed line underneath) and the text "CONTRAIL COMMAND". On the right is a white login form titled "Log in". The form contains the following fields:


- Select Cluster:** A dropdown menu.
- Username:** A text input field containing the text "admin".
- Password:** A text input field with masked characters (dots).
- Log in:** A blue button.

3. Click on **Clusters**.

You will see the list of all the available clusters with the status.

4. Select the cluster you want to upgrade.

Hover your mouse over ellipsis next to the cluster and click on **Upgrade**.



5. Enter **Contrail Version**, **Container Registry**, **Container Registry Username**, **Container Registry Password**.

Contrail Version depicts the current installed Contrail version. You must update the value to the desired version number.

The values for **Container Registry**, **Container Registry Username**, and **Container Registry Password** are pre-populated based on the values used during initial Contrail deployment.

Click on **Contrail Configuration**.

Add **CONTRAIL_CONTAINER_TAG**.

Access **CONTRAIL_CONTAINER_TAG** values for your targeted upgrade from [README Access to Contrail Networking Registry 20xx](#).

6. If you have Contrail Insights and Contrail Insights Flows installed in the cluster:

NOTE: Appformix and Appformix Flows were renamed Contrail Insights and Contrail Insights Flows. The Appformix naming conventions still appear during product usage, including within these directory names.

- *Release 2005 or later:* You do not need to provide appropriate versions of Contrail Insights and Contrail Insights packages in the **/opt/software/appformix** and **/opt/software/xflow** directories. This step is no longer required starting in Release 2005.
- *Release 2003 or earlier:*

Provide appropriate versions of Contrail Insights and Contrail Insights packages in the **/opt/software/appformix** and **/opt/software/xflow** directories on the Contrail Command server. For more details, refer to *Installing Contrail Insights and Contrail Insights Flows using Contrail Command*.

Skip this step if you are not using Contrail Insights or Contrail Insights Flows.

7. Click on **Upgrade**.

[Figure 18 on page 112](#) provides a representative illustration of a user completing the **Upgrade Cluster** workflow to upgrade to Contrail Networking Release 2005.

Figure 18: Upgrade Cluster Workflow Example

Upgrade Cluster

Contrail Version *

2005.62

Container Registry *

hub.juniper.net/contrail

Container Registry Username *

JNPR-FieldUser55

Container Registry Password *

▼ Contrail Configuration

Key *

Value *

CONTRAIL_CONTAINER_TAG

2005.62

+ Add

Cancel

Upgrade

RELATED DOCUMENTATION

| *Installing Contrail Command*

Upgrading Contrail Command using Backup Restore Procedure

You cannot use the SQL data with the new version of Contrail Command container if the database schema changes while upgrading the Contrail Command container.

You can resolve the issue by:

1. Back up SQL database in **yaml format db dump**.

Run the following **docker exec contrail_command** command on the Contrail Command node to backup the DB.

Contrail Networking Release 2005 or later:

```
docker exec contrail_command commandutil convert --intype rdbms --outtype yaml --out
/etc/contrail/db.yml -c /etc/contrail/command-app-server.yml; mkdir ~/backups; mv
/etc/contrail/db.yml ~/backups/
```

Contrail Networking Release 2003 or earlier:

```
docker exec contrail_command contrailutil convert --intype rdbms --outtype yaml --out
/etc/contrail/db.yml -c /etc/contrail/contrail.yml; mkdir ~/backups; mv /etc/contrail/db.yml
~/backups/
```

2. Upgrade the Contrail Command container.

Specify the desired version of Contrail Command container (*container_tag*) in the deployer input file (**command_servers.yml**) and deploy playbook.

You must use *PostgreSQL* in the **command_servers.yml** file.

The step depends on how you have deployed the Contrail Command.

- Contrail Command is deployed using docker installation:

```
docker run -td --net host -v
<ABSOLUTE_PATH_OF_COMMAND_SERVERS_FILE>:/command_servers.yml --privileged --name
contrail_command_deployer_<contrail_container_tag>
hub.juniper.net/contrail/contrail-command-deployer:<<contrail_container_tag>
```

- Contrail Command is deployed through juju-charms:

```
juju config contrail-command image-tag=<contrail_container_tag>
```

The *contrail_container_tag* for any Contrail Release 20 software can be obtained from [README Access to Contrail Networking Registry 20xx](#).

3. This step depends on your Contrail Networking release.

Contrail Networking Release 2005 or later:

Migrate the **yaml formatted db dump** to the new database schema:

```
docker exec contrail_command mkdir /root/backups
docker cp /root/backups/db.yml contrail_command:/root/backups/
```

```
docker exec contrail_command commandutil migrate --in /root/backups/db.yml --out
/root/backups/db_migrated.yml
```

Contrail Networking Release 2003 or earlier:

Modify the yaml-formatted db dump by adding or removing the fields per the new database schema.

4. Restore the modified **yaml formatted db dump** to the SQL database.

Contrail Networking Release 2005 or later:

```
docker exec contrail_command commandutil convert --intype yaml --in
/root/backups/db_migrated.yml --outtype rdbms -c
/etc/contrail/command-app-server.yml
```

Contrail Networking Release 2003 or earlier:

```
docker exec contrail_command mkdir /root/backups docker cp /root/backups/db.yml
contrail_command:/root/backups/
docker exec contrail_command contrailutil convert --intype yaml --in
~/backups/db.yml --outtype rdbms -c /etc/contrail/contrail.yml
```

NOTE: If the restore procedure fails because of schema mismatch, repeat Step 3 and Step 4 with incremental db dump changes.

How to Perform a Zero Impact Contrail Networking Upgrade using the Ansible Deployer

Starting in Contrail Networking Release 2005, you can perform a Zero Impact Upgrade (ZIU) of Contrail Networking using the Contrail Ansible Deployer container. The Contrail Ansible Deployer container image can be loaded from the Juniper Networks Contrail Container Registry hosted at hub.juniper.net/contrail.

Use the procedure in this document to perform a Zero Impact Upgrade (ZIU) of Contrail Networking using the Contrail Ansible Deployer container. This ZIU allows Contrail Networking to upgrade while sustaining minimal network downtime.

Before you begin:

- The target release for this upgrade must be Contrail Release 2005 or later.
- You can use this procedure to incrementally upgrade to the next Contrail Networking release only. For instance, if you are running Contrail Networking Release 2003 and want to upgrade to the next Contrail Release—which is Contrail Networking Release 2005—you can use this procedure to perform the upgrade.

This procedure is not validated for upgrades between releases that are two or more releases apart. For instance, it could not be used to upgrade from Contrail Networking Release 2002 to Contrail Networking Release 2005.

For a list of Contrail Networking releases in a table that illustrates Contrail Networking release order, see [Contrail Networking Supported Platforms](#).

- The Contrail Ansible Deployer container can only be used in CentOS environments.
- Take snapshots of your current configurations before you proceed with the upgrade process. For details, refer to [“How to Backup and Restore Contrail Databases in JSON Format” on page 142](#).

This procedure illustrates how to perform a ZIU using the Ansible deployer container. It includes a representative example of the steps being performed to upgrade from Contrail Networking Release 2005 to Release 2008.

To perform the ZIU using the Ansible deployer:

1. Pull the *contrail-ansible-deployer* file for the target upgrade release. This procedure is typically performed from a Contrail controller running in your environment, but it can also be performed from a separate server which has network connectivity to the deployment that is being upgraded.

This procedure shows you how to load a 2008 image from the Juniper Networks Contrail Container Registry. You can, however, also change the values to load the file from a private registry.

The Juniper Networks Contrail Container Registry is hosted at *hub.juniper.net/contrail*. If you need the credentials to access the registry, email *contrail-registry@juniper.net*.

Enter the following commands to pull the *contrail-ansible-deployer* file from the registry:

```
sudo docker login -u <username> -p <password> hub.juniper.net
sudo docker pull
hub.juniper.net/contrail/contrail-kolla-ansible-deployer:2008.<contrail_container_tag>
```

where:

- *username*—username to access the registry. Email *contrail-registry@juniper.net* if you need to obtain *username* and *password* credentials.
- *password*—password to access the registry. Email *contrail-registry@juniper.net* if you need to obtain *username* and *password* credentials.
- *contrail_container_tag*—the container tag ID for your target Contrail Networking release. The *contrail_container_tag* for any Contrail Release 20 software can be obtained from [README Access to Contrail Networking Registry 20xx](#).

2. Start the Contrail Ansible Deployer:

```
docker run -t --net host -d --privileged --name contrail-kolla-ansible-deployer
hub.juniper.net/contrail/contrail-kolla-ansible-deployer:2008.<contrail_container_tag>
```

3. Navigate to the *instances.yaml* file and open it for editing.

The *instances.yaml* file was used to initially deploy the setup. The *instances.yaml* can be loaded into the Contrail Ansible Deployer and edited to supported the target upgrade version.

Contrail Release 2008 Target Upgrade Example using VI as the editor:

```
docker cp instances.yaml
contrail-kolla-ansible-deployer:/root/contrail-ansible-deployer/config/instances.yaml
docker exec -it contrail-kolla-ansible-deployer bash
cd /root/contrail-ansible-deployer/config/
vi instances.yaml
```

4. Update the `CONTRAIL_CONTAINER_TAG` to the desired version tag in the **instances.yaml** file from the existing deployment. The `CONTRAIL_CONTAINER_TAG` variable is in the `contrail_configuration:` hierarchy within the **instances.yaml** file.

The `CONTRAIL_CONTAINER_TAG` for any Contrail Release 20 software can be obtained from [README Access to Contrail Networking Registry 20xx](#).

Here is an example instances.yml file configuration:

```
contrail_configuration:
  CONTRAIL_CONTAINER_TAG: "2008.121"
  CONFIG_DATABASE_NODEMGR__DEFAULTS__minimum_diskGB: "2"
  DATABASE_NODEMGR__DEFAULTS__minimum_diskGB: "2"
  JVM_EXTRA_OPTS: "-Xms1g -Xmx2g"
  VROUTER_ENCRYPTION: FALSE
  LOG_LEVEL: SYS_DEBUG
  CLOUD_ORCHESTRATOR: kubernetes
```

5. Upgrade the control plane by running the `ziu.yml` playbook file from inside the contrail ansible deployer container.

- For Contrail Networking Release 2005 to Contrail Networking Release 2008:

Upgrade the control plane by running the **ziu.yml** playbook file.

```
sudo -E ansible-playbook -v -e orchestrator=openstack -e config_file=instances.yaml
playbooks/ziu.yml
```

- For Contrail Networking Release 2011 and later:

Upgrade the control plane by running the controller stage of **ziu.yml** playbook file.

```
sudo -E ansible-playbook -v -e stage=controller -e orchestrator=openstack -e
config_file=../instances.yaml playbooks/ziu.yml
```

6. Upgrade the Openstack plugin by running the `install_openstack.yml` playbook file.

- For Contrail Networking Release 2005 to Contrail Networking Release 2008:

```
sudo -E ansible-playbook -v -e orchestrator=openstack -e config_file=instances.yaml
playbooks/install_openstack.yml
```

- For Contrail Networking Release 2011 and later:

```
sudo -E ansible-playbook -v -e stage=openstack -e orchestrator=openstack -e
config_file=../instances.yaml playbooks/ziu.yml
```

7. Enter the **contrail-status** command to monitor upgrade status. Ensure all pods reach the *running* state and all services reach the *active* state.

This **contrail-status** command provides this output after a successful upgrade:

NOTE: Some output fields and data have been removed for readability.

Pod	Service	Original Name	State
	redis	contrail-external-redis	running
	rsyslogd		running
analytics	api	contrail-analytics-api	running
analytics	collector	contrail-analytics-collector	running
analytics	nodemgr	contrail-nodemgr	running
analytics	provisioner	contrail-provisioner	running
analytics-alarm	alarm-gen	contrail-analytics-alarm-gen	running
analytics-alarm	kafka	contrail-external-kafka	running
analytics-alarm	nodemgr	contrail-nodemgr	running
analytics-alarm	provisioner	contrail-provisioner	running
analytics-snmp	nodemgr	contrail-nodemgr	running
analytics-snmp	provisioner	contrail-provisioner	running
analytics-snmp	snmp-collector	contrail-analytics-snmp-collector	running
analytics-snmp	topology	contrail-analytics-snmp-topology	running
config	api	contrail-controller-config-api	running
config	device-manager	contrail-controller-config-devicemgr	running
config	dnsmasq	contrail-controller-config-dnsmasq	running
config	nodemgr	contrail-nodemgr	running
config	provisioner	contrail-provisioner	running
config	schema	contrail-controller-config-schema	running
config	stats	contrail-controller-config-stats	running
config	svc-monitor	contrail-controller-config-svcmonitor	running
config-database	cassandra	contrail-external-cassandra	running

<trimmed>

vrouter kernel module is PRESENT
 == Contrail control ==
 control: active
 nodemgr: active

```
named: active
dns: active

== Contrail analytics-alarm ==
nodemgr: active
kafka: active
alarm-gen: active

== Contrail kubernetes ==
kube-manager: active

== Contrail database ==
nodemgr: active
query-engine: active
cassandra: active

== Contrail analytics ==
nodemgr: active
api: active
collector: active

== Contrail config-database ==
nodemgr: active
zookeeper: active
rabbitmq: active
cassandra: active

== Contrail webui ==
web: active
job: active

== Contrail vrouter ==
nodemgr: active
agent: active

== Contrail analytics-snmp ==
snmp-collector: active
nodemgr: active
topology: active

== Contrail config ==
svc-monitor: active
nodemgr: active
device-manager: active
```

```
api: active
schema: active
```

8. Migrate workloads VM from one group of compute nodes. Leave them uncommented in the `instances.yaml` file. Comment other computes not ready to upgrade in `instances.yaml`.

9. Upgrade compute nodes.

- For Contrail Networking Release 2005 to Contrail Networking Release 2008:

Run the `install_contrail.yml` playbook file to upgrade the compute nodes that were uncommented in the `instances.yaml` file. Only the compute nodes that were left uncommented in 8 are upgraded to the target release in this step.

```
sudo -E ansible-playbook -v -e orchestrator=openstack -e config_file=instances.yaml
playbooks/install_contrail.yml
```

- For Contrail Networking Release 2011 and later:

Run the compute stage of `ziu.yml` playbook file to upgrade the compute nodes that were uncommented in the `instances.yaml` file. Only the compute nodes that were left uncommented in 8 are upgraded to the target release in this step.

```
sudo -E ansible-playbook -v -e stage=compute -e orchestrator=openstack -e
config_file=../instances.yaml playbooks/ziu.yml
```

10. Repeat Steps 8 and 9 until all compute nodes are upgraded.

You can access the Ansible playbook logs of the upgrade at `/var/log/ansible.log`.

Release History Table

Release	Description
2005	Starting in Contrail Networking Release 2005, you can perform a Zero Impact Upgrade (ZIU) of Contrail Networking using the Contrail Ansible Deployer container.

RELATED DOCUMENTATION

[Installing a Contrail Cluster using Contrail Command and `instances.yaml`](#) | 74

[Upgrading Contrail Networking using Contrail Command](#) | 109

Updating Contrail Networking using the Zero Impact Upgrade Process in an Environment using Red Hat Openstack

IN THIS SECTION

- [When to Use This Procedure | 121](#)
- [Prerequisites | 122](#)
- [Before You Begin | 122](#)
- [Updating Contrail Networking in an Environment using Red Hat Openstack | 123](#)

This document provides the steps needed to update a Contrail Networking deployment that is using Red Hat Openstack as it's orchestration platform. The procedure provides a zero impact upgrade (ZIU) with minimal disruption to network operations.

When to Use This Procedure

This procedure is used to upgrade Contrail Networking when it is running in environments using RHOSP13.

The procedure in this document has been validated for the following Contrail Networking upgrade scenarios:

Table 4: Contrail Networking with RHOSP13 Validated Upgrade Scenarios

Starting Contrail Networking Release	Target Upgraded Contrail Networking Release
1912.L0	1912.L1
1912.L0	2003
2003	2005
2005	2008

A different procedure is followed for upgrading to earlier target Contrail Networking releases in environments using RHOSP13 orchestration. See *Upgrading Contrail Networking with Red Hat Openstack 13 using ISSU*.

If you want to use this procedure to upgrade your Contrail Networking release to other releases, you must engage Juniper Networks professional services. Contact your Juniper representative for additional information.

Prerequisites

This document makes the following assumptions about your environment:

- A Contrail Networking deployment using Red Hat Openstack version 13 (RHOSP13) as the orchestration platform is already operational.
- The overcloud nodes in the RHOSP13 environment have an enabled Red Hat Enterprise Linux (RHEL) subscription.
- Your environment is running Contrail Release 1912 and upgrading to Contrail Release 1912-L1 or to Contrail Release 2003 or later.
- If you are updating Red Hat Openstack simultaneously with Contrail Networking, we assume that the undercloud node is updated to the latest minor version and that new overcloud images are prepared for an upgrade if needed for the upgrade. See the [Upgrading the Undercloud](#) section of the [Keeping Red Hat OpenStack Platform Updated](#) guide from Red Hat.

If the undercloud has been updated and a copy of the heat templates are used for the deployment, update the copy of the heat template from the Red Hat's core heat template collection at `/usr/share/openstack-tripleo-heat-templates`. See the [Understanding Heat Templates](#) document from Red Hat for information on this process.

Before You Begin

We recommend performing these procedures before starting the update:

- Backup your Contrail configuration database before starting this procedure. See [“How to Backup and Restore Contrail Databases in JSON Format” on page 142](#).
- Each compute node agent will go down during this procedure, causing some compute node downtime. The estimated downtime for a compute node varies by environment, but typically took between 12 and 15 minutes in our testing environments.

If you have compute nodes with workloads that cannot tolerate this downtime, consider migrating workloads or taking other steps to accommodate this downtime in your environment.

- If you are updating Red Hat Openstack simultaneously with Contrail Networking, update Red Hat Openstack to the latest minor release version and ensure that the new overcloud images are prepared for the upgrade. See the [Upgrading the Undercloud](#) section of the [Keeping Red Hat OpenStack Platform Updated](#) guide from Red Hat for additional information.

If the undercloud has been updated and a copy of the heat templates are used for the deployment, update the Heat templates from Red Hat's core Heat template collection at `/usr/share/openstack-tripleo-heat-templates`. See the [Understanding Heat Templates](#) document from Red Hat for additional information.

Updating Contrail Networking in an Environment using Red Hat Openstack

To update Contrail Networking in an environment that is using Red Hat Openstack as the orchestration platform:

1. Prepare your docker registry. The registry is often included in the undercloud, but it can also be a separate node.

Docker registry setup is environment independent. See [Docker Registry](#) from Docker for additional information on Docker registry setup.

2. Backup the Contrail TripleO Heat Templates. See [Using the Contrail Heat Template](#).

3. Get the Contrail TripleO Heat Templates (Stable/Queens branch) from <https://github.com/Juniper/contrail-tripleo-heat-templates>.

4. (Optional) Update the Contrail TripleO Puppet module to the latest version and prepare Swift Artifacts, as applicable.

Below are sample commands entered in the undercloud:

```
[stack@queensa ~]$ mkdir -p ~/usr/share/openstack-puppet/modules/tripleo
[stack@queensa ~]$ git clone -b stable/queens
https://github.com/Juniper/contrail-tripleo-puppet
usr/share/openstack-puppet/modules/tripleo
[stack@queensa ~]$ tar czvf puppet-modules.tgz usr/
[stack@queensa ~]$ upload-swift-artifacts -c contrail-artifacts -f
puppet-modules.tgz
```

5. Update the parameter *ContrailImageTag* to the new version.

The location of the *ContrailImageTag* variable varies by environment. In the most commonly-used environments, this variable is set in the *contrail-services.yaml* file.

You can obtain the *ContrailImageTag* parameter from the [README Access to Contrail Registry 20XX](#).

6. (Recommended) If you are upgrading to Contrail Networking Release 2005 or later, check and, if needed, enable kernel vRouter huge page support to support future compute node upgrades without rebooting.

You can enable or verify kernel-mode vRouter huge page support in the *contrail-services.yaml* file using either the **ContrailVrouterHugepages1GB:** and **ContrailVrouterHugepages2MB:** parameters:

```
parameter_defaults:
  ...
  ContrailVrouterHugepages1GB: '2'
```

```
parameter_defaults:
  ...
  # ContrailVrouterHugepages2MB: '1024'
```

Notes about kernel-mode vRouter huge page support in Red Hat Openstack environments:

- Kernel-mode vRouter huge page support was introduced in Contrail Networking Release 2005, and is configured to support 2 1GB huge pages by default in Contrail Networking Release 2005 or later.

A compute node has to be rebooted once for a huge page configuration to finalize. After this initial reboot, the compute node can perform future Contrail Networking software upgrades without rebooting.

Notably, a compute node in an environment running Contrail Networking 2005 or later has not enabled huge page support for kernel-mode vRouters until it is rebooted. The 2x1GB huge page support configuration is present by default, but it isn't enabled until the compute node is rebooted.

- We recommend only using 1GB or 2MB kernel-mode vRouter huge pages in most environments. You can, however, simultaneously enable 1GB or 2MB kernel-mode vRouter huge pages in Red Hat Openstack environments if your environment requires enablement of both huge page options.
 - Changing vRouter huge page configuration settings in a Red Hat Openstack environment typically requires a compute node reboot.
 - 1 GB pages: Reboot required.
 - 2 MB: Reboot usually required. The reboot is sometimes avoided in environments where memory isn't highly fragmented or the required number of pages can be easily allocated.
 - We recommend allotting 2GB of memory—either using the default 1024x2MB huge page size setting or the 2x1GB size setting—for huge pages in most environments. Some larger environments might require additional huge page memory settings for scale. Other huge page size settings should only be set by expert users in specialized circumstances.
 - If the **ContrailVrouterHugepages1GB:** and **ContrailVrouterHugepages2MB:** parameters are set to empty value in the **contrail-services.yaml** file, vRouter huge pages are disabled.
7. Update the overcloud by entering the **openstack overcloud update prepare** command and include the files that were updated during the previous steps with the overcloud update.

Example:

```

openstack overcloud update prepare
--templates tripleo-heat-templates/
--roles-file tripleo-heat-templates/roles_data_contrail_aio.yaml -e
environment-rhel-registration.yaml -e
tripleo-heat-templates/extraconfig/pre_deploy/rhel-registration/rhel-registrationresource-registry.yaml
-e
tripleo-heat-templates/environments/contrail/contrail-services.yaml -e
tripleo-heat-templates/environments/contrail/contrail-net-single.yaml -e
tripleo-heat-templates/environments/contrail/contrail-plugins.yaml -e
misc_opts.yaml -e
contrail-parameters.yaml -e
docker_registry.yaml

```

8. Prepare the overcloud nodes that include Contrail containers for the update.

- Pull the images in the repository onto the overcloud nodes.

There are multiple methods for performing this step. Commonly used methods for performing this operation include using the **docker pull** command for Docker containers and the **openstack overcloud container image upload** command for Openstack containers, or running the **tripleo-heat-templates/upload.containers.sh** and **tools/contrail/update_contrail_preparation.sh** scripts.

- (Not required in all setups) Provide export variables for the script if the predefined values aren't appropriate for your environment. The script location:

```
~/tripleo-heat-templates/tools/contrail/update_contrail_preparation.sh
```

The following variables within the script are particularly significant for this upgrade:

- **CONTRAIL_NEW_IMAGE_TAG**—The image tag of the target upgrade version of Contrail. The default value is **latest**.

If needed, you can obtain this parameter for a specific image from the [README Access to Contrail Registry 20XX](#).

NOTE: Some older deployments use the **CONTRAIL_IMAGE_TAG** variable in place of the **CONTRAIL_NEW_IMAGE_TAG** variable. Both variables are recognized by the **update_contrail_preparation.sh** script and perform the same function.

- **SSH_USER**—The SSH username for logging into overcloud nodes. The default value is **heat-admin**.
- **SSH_OPTIONS**—Custom SSH option values.

The default SSH options for your environment are typically pre-defined. You are typically only changing this value if you want to customize your update.

- **STOP_CONTAINERS**—The list of containers that must be stopped before the upgrade can proceed. The default value is **contrail_config_api contrail_analytics_api**.
- Run the script:



CAUTION: Contrail services stop working when the script starts running.

```
~/tripleo-heat-templates/tools/contrail/update_contrail_preparation.sh
```

9. Update the Contrail Controller nodes:

- Run the **openstack overcloud update run** command on the first Contrail controller and, if needed, on a Contrail Analytics node. The purpose of this step is to update one Contrail Controller and one Contrail Analytics node to support the environment so the other Contrail Controllers and analytics nodes can be updated without incurring additional downtime.

Example:

```
openstack overcloud update run --nodes overcloud-contrailcontroller-0
```

Ensure that the contrail status is **ok** on overcloud-contrailcontroller-0 before proceeding.

If the analytics and the analyticsdb nodes are on separate nodes, you may have to update the nodes individually:

```
openstack overcloud update run --nodes overcloud-contrailcontroller-0
openstack overcloud update run --roles
ContrailAnalytics,ContrailAnalyticsDatabase
```

- After the upgrade, check the docker container status and versions for the Contrail Controllers and the Contrail Analytics and AnalyticsDB nodes.

```
docker ps -a
```

- Update the remaining Contrail Controller nodes:

Example:

```
openstack overcloud update run --nodes overcloud-contrailcontroller-1
openstack overcloud update run --nodes overcloud-contrailcontroller-2
openstack overcloud update run --nodes overcloud-contrailcontroller-3
...
```

10. Update the Openstack Controllers using the **openstack overcloud update run** commands:

Example:

```
openstack overcloud update run --nodes overcloud-controller-0
openstack overcloud update run --nodes overcloud-controller-1
openstack overcloud update run --nodes overcloud-controller-2
...
```

11. Individually update the compute nodes.

NOTE: The compute node agent will be down during this step. The estimated downtime varies by environment, but is typically between 1 and 5 minutes.

Consider migrating workloads that can't tolerate this downtime before performing this step

```
openstack overcloud update run --nodes overcloud-novacompute-1
openstack overcloud update run --nodes overcloud-novacompute-2
openstack overcloud update run --nodes overcloud-novacompute-3
...
```

Reboot your compute node to complete the update.

NOTE: A reboot is required to complete this procedure only if a kernel update is also needed. If you would like to avoid rebooting your compute node, check the log files in the **/var/log/yum.log** file to see if kernel packages were updated during the compute node update. A reboot is required only if kernel updates occurred as part of the compute node update procedure.

```
sudo reboot
```

Use the **contrail-status** command to monitor upgrade status. Ensure all pods reach the *running* state and all services reach the *active* state.

This **contrail-status** command provides output after a successful upgrade:

NOTE: Some output fields and data have been removed from this **contrail-status** command sample for readability.

Pod	Service	Original Name	State
analytics	api	contrail-analytics-api	running
analytics	collector	contrail-analytics-collector	running
analytics	nodemgr	contrail-nodemgr	running
analytics	provisioner	contrail-provisioner	running
analytics	redis	contrail-external-redis	running
analytics-alarm	alarm-gen	contrail-analytics-alarm-gen	running
analytics-alarm	kafka	contrail-external-kafka	running
analytics-alarm	nodemgr	contrail-nodemgr	running
analytics-alarm	provisioner	contrail-provisioner	running
analytics-alarm	zookeeper	contrail-external-zookeeper	running
analytics-snmp	nodemgr	contrail-nodemgr	running
analytics-snmp	provisioner	contrail-provisioner	running
analytics-snmp	snmp-collector	contrail-analytics-snmp-collector	running
analytics-snmp	topology	contrail-analytics-snmp-topology	running
config	api	contrail-controller-config-api	running

<trimmed>

== Contrail control ==

control: active

nodemgr: active

named: active

dns: active

== Contrail analytics-alarm ==

nodemgr: active

kafka: active

alarm-gen: active

== Contrail database ==

nodemgr: active

query-engine: active

cassandra: active


```

== Contrail analytics ==
nodemgr: active
api: active
collector: active

== Contrail config-database ==
nodemgr: active
zookeeper: active
rabbitmq: active
cassandra: active

== Contrail webui ==
web: active
job: active

== Contrail analytics-snmp ==
snmp-collector: active
nodemgr: active
topology: active

== Contrail config ==
svc-monitor: active
nodemgr: active
device-manager: active
api: active
schema: active

```

12. Enter the **openstack overcloud update converge** command to finalize the update.

NOTE: The options used in the **openstack overcloud update converge** in this step will match the options used with the **openstack overcloud update prepare** command entered in 7.

```

openstack overcloud update converge
--templates tripleo-heat-templates/
--roles-file tripleo-heat-templates/roles_data_contrail_aio.yaml -e
environment-rhel-registration.yaml -e
tripleo-heat-templates/extraconfig/pre_deploy/rhel-registration/rhel-registrationresource-registry.yaml
-e
tripleo-heat-templates/environments/contrail/contrail-services.yaml -e
tripleo-heat-templates/environments/contrail/contrail-net-single.yaml -e

```

```
tripleo-heat-templates/environments/contrail/contrail-plugins.yaml -e
misc_opts.yaml -e
contrail-parameters.yaml -e
docker_registry.yaml
```

Monitor screen messages indicating *SUCCESS* to confirm that the updates made in this step are successful.

RELATED DOCUMENTATION

| [Installing Contrail with OpenStack by Using Juju Charms](#) | 412

Updating Contrail Networking using the Zero Impact Upgrade Procedure in a Canonical Openstack Deployment with Juju Charms

IN THIS SECTION

- [Prerequisites](#) | 130
- [When to Use This Procedure](#) | 131
- [Recommendations](#) | 131
- [Updating Contrail Networking in a Canonical Openstack Deployment Using Juju Charms](#) | 132

This document provides the steps needed to update a Contrail Networking deployment that is using Canonical Openstack as its orchestration platform. The procedure utilizes Juju charms and provides a zero impact upgrade (ZIU) with minimal disruption to network operations.

Prerequisites

This document makes the following assumptions about your environment:

- A Contrail Networking deployment using Canonical Openstack as the orchestration platform is already operational.
- Juju charms for Contrail services are active in your environment, and the Contrail Networking controller has access to the Juju jumphost and the Juju cluster.

When to Use This Procedure

This procedure is used to upgrade Contrail Networking when it is running in environments using Canonical Openstack.

You can use this procedure to incrementally upgrade to the next main Contrail Networking release only. This procedure is not validated for upgrades between releases that are two or more releases apart.

The procedure in this document has been validated for the following Contrail Networking upgrade scenarios:

Table 5: Contrail Networking with Canonical Openstack Validated Upgrade Scenarios

Starting Contrail Networking Release	Target Upgraded Contrail Networking Release
Any 1912.L Release	Any 1912.L Release
1912 or any 1912.L release	2003
2003	2005
2005	2008
2008	2011

Recommendations

We strongly recommend performing the following tasks before starting this procedure:

- Backup your current environment.
- Enable huge pages for the kernel-mode vRouter.

Starting in Contrail Networking Release 2005, you can enable huge pages in the kernel-mode vRouter to avoid future compute node reboots during upgrades. Huge pages in Contrail Networking are used primarily to allocate flow and bridge table memory within the vRouter. Huge pages for kernel-mode vRouters provide enough flow and bridge table memory to avoid compute node reboots to complete future Contrail Networking software upgrades.

We recommend allotting 2GB of memory—either using the default 1024x2MB huge page size setting or the 2x1GB size setting—for huge pages. Other huge page size settings should only be set by expert users in specialized circumstances.

A compute node reboot is required to initially enable huge pages. Future compute node upgrades can happen without reboots after huge pages are enabled. The 1024x2MB huge page setting is configured by default starting in Contrail Networking Release 2005, but is not active in any compute node until the compute node is rebooted to enable the setting.

2GB and 1MB huge page size settings cannot be enabled simultaneously in environments using Juju.

The huge page configurations can be changed by entering one of the following commands:

- Enable 1024 2MB huge pages (Default): **juju config contrail-agent kernel-hugepages-2m=1024**

Disable 2MB huge pages (empty value): **juju config contrail-agent kernel-hugepages-2m=""**

- Enable 2 1GB huge pages: **juju config contrail-agent kernel-hugepages-1g=2**

Disable 1GB huge pages (default. empty value): **juju config contrail-agent kernel-hugepages-1g=""**

NOTE: 1GB huge page settings can only be specified at initial deployment; you cannot modify the setting in active deployments. The 1GB huge page setting can also not be completely disabled after being activated on a compute node. Be sure that you want to use 1GB huge page settings on your compute node before enabling the setting.

Updating Contrail Networking in a Canonical Openstack Deployment Using Juju Charms

To update Contrail Networking in an environment that is using Canonical Openstack as the orchestration platform:

1. Upgrade all charms. See the [Upgrading applications](#) document from Juju.
2. From the Juju jumphost, enter the **run-action** command to place all control plane services—Contrail Controller, Contrail Analytics, & Contrail AnalyticsDB—into maintenance mode in preparation for the upgrade.

```
juju run-action --wait contrail-controller/leader upgrade-ziu
```

NOTE: The **--wait** option is not required to complete this step, but is recommended to ensure this procedure completes without interfering with the procedures in the next step.

Wait for all charms to move to the *maintenance* status. You can check the status of all charms by entering the **juju status** command.

3. Update the image tags in Juju for the Contrail Analytics, Contrail AnalyticsDB, Contrail Agent, and Contrail Openstack services.

```
juju config contrail-analytics image-tag=master-latest
juju config contrail-analyticsdb image-tag=master-latest
juju config contrail-agent image-tag=master-latest
juju config contrail-openstack image-tag=master-latest
```

If a Contrail Service node (CSN) is part of the cluster, also update the image tags in Juju for the Contrail Service node.

```
juju config contrail-agent-csn image-tag=master-latest
```

- 4. Update the image tag in Juju for the Contrail Controller service:

```
juju config contrail-controller image-tag=master-latest
```

- 5. After updating the image tags, wait for all services to complete stage 5 of the ZIU upgrade process workflow. The wait time for this step varies by environment, but often takes 30 to 90 minutes.

Enter the **juju status** command and review the **Workload** and **Message** field outputs to monitor progress. The update is complete when all services are in the maintenance state—the **Workload** field output is **maintenance**—and each individual service has completed stage 5 of the ZIU upgrade—illustrated by the **ziu is in progress - stage/done = 5/5** output in the **Message** field.

A sample output of an in-progress update that has not completed the image tag update process. The **Message** field illustrates that the ZIU processes have not completed stage 5 of the upgrade.

NOTE: Some **juju status** output fields removed for readability.

juju status			
Unit	Workload	Agent	Message
contrail-analytics/0* = 4/4	maintenance	idle	ziu is in progress - stage/done
contrail-analytics/1 = 4/4	maintenance	idle	ziu is in progress - stage/done
contrail-analytics/2 = 4/4	maintenance	idle	ziu is in progress - stage/done
contrail-analyticsdb/0* = 4/4	maintenance	idle	ziu is in progress - stage/done
contrail-analyticsdb/1 = 4/3	maintenance	idle	ziu is in progress - stage/done

```

contrail-analyticsdb/2      maintenance idle      ziu is in progress - stage/done
= 4/3
contrail-controller/0*     maintenance idle      ziu is in progress - stage/done
= 4/4
  ntp/3                    active      idle      chrony: Ready
contrail-controller/1     maintenance executing ziu is in progress - stage/done
= 4/3
  ntp/2                    active      idle      chrony: Ready
contrail-controller/2     maintenance idle      ziu is in progress - stage/done
= 4/3
  ntp/4                    active      idle      chrony: Ready
contrail-keystone-auth/0* active      idle      Unit is ready

```

A sample output of an update that has completed the image tag update process on all services. The **Workload** field is **maintenance** for all services and the **Message** field explains that stage 5 of the ZIU process is done.

NOTE: Some **juju status** output fields removed for readability.

```

juju status
Unit                                Workload    Agent Message
contrail-analytics/0*              maintenance idle      ziu is in progress - stage/done =
5/5
contrail-analytics/1              maintenance idle      ziu is in progress - stage/done =
5/5
contrail-analytics/2              maintenance idle      ziu is in progress - stage/done =
5/5
contrail-analyticsdb/0*           maintenance idle      ziu is in progress - stage/done =
5/5
contrail-analyticsdb/1           maintenance idle      ziu is in progress - stage/done =
5/5
contrail-analyticsdb/2           maintenance idle      ziu is in progress - stage/done =
5/5
contrail-controller/0*            maintenance idle      ziu is in progress - stage/done =
5/5
  ntp/3                          active      idle      chrony: Ready
contrail-controller/1            maintenance idle      ziu is in progress - stage/done =
5/5
  ntp/2                          active      idle      chrony: Ready
contrail-controller/2            maintenance idle      ziu is in progress - stage/done =

```

5/5

ntp/4	active	idle	chrony: Ready
contrail-keystone-auth/0*	active	idle	Unit is ready
glance/0*	active	idle	Unit is ready
haproxy/0*	active	idle	Unit is ready
keepalived/2	active	idle	VIP ready
haproxy/1	active	idle	Unit is ready
keepalived/0*	active	idle	VIP ready
haproxy/2	active	idle	Unit is ready
keepalived/1	active	idle	VIP ready
heat/0*	active	idle	Unit is ready
contrail-openstack/3	active	idle	Unit is ready
keystone/0*	active	idle	Unit is ready
mysql/0*	active	idle	Unit is ready
neutron-api/0*	active	idle	Unit is ready
contrail-openstack/2	active	idle	Unit is ready
nova-cloud-controller/0*	active	idle	Unit is ready
nova-compute/0*	active	idle	Unit is ready

6. Upgrade every Contrail agent on each individual compute node:

```
juju run-action contrail-agent/0 upgrade
juju run-action contrail-agent/1 upgrade
juju run-action contrail-agent/2 upgrade
...
```

If Contrail Service nodes (CSNs) are part of the cluster, also upgrade every Contrail CSN agent:

```
juju run-action contrail-agent-csn/0 upgrade
...
```

Wait for each compute node and CSN node upgrade to finish. The wait time for this step varies by environment, but typically takes around 10 minutes to complete per node.

7. If huge pages are not enabled for your vRouter, log into each individual compute node and reboot to complete the procedure.

NOTE: A compute node reboot is required to initially enable huge pages. If huge pages have been configured in Juju without a compute node reboot, you can also use this reboot to enable huge pages. You can avoid rebooting the compute node during future software upgrades after this initial reboot.

1024x2MB huge page support is configured by default starting in Contrail Networking Release 2005, which is also the first Contrail Networking release that supports huge pages. If you are upgrading to Release 2005 for the first time, a compute node reboot is always required because huge pages could not have been previously enabled.

This reboot also enables the default 1024x2MB huge page configuration unless you change the huge page configuration in Release 2005 or later.

```
sudo reboot
```

This step can be skipped if huge pages are enabled.

RELATED DOCUMENTATION

| [Installing Contrail with OpenStack by Using Juju Charms](#) | 412

Upgrading Contrail Networking using contrail-ansible Deployer

NOTE: This procedure can be used for Contrail Networking upgrades to Contrail Networking Release 2003 or earlier.

If you are upgrading to Contrail Networking Release 2005 or later using the Ansible deployer, see [“How to Perform a Zero Impact Contrail Networking Upgrade using the Ansible Deployer” on page 114](#).

Use the following procedure to upgrade Contrail Networking using contrail-ansible deployer.

The procedure supports incremental model and you can use it to upgrade from Contrail Networking Release *N-1* to *N*.

Take snapshots of your current configurations before you proceed with the upgrade process. For details, refer to [“How to Backup and Restore Contrail Databases in JSON Format” on page 142](#).

1. Navigate to the directory where the **contrail-ansible-deployer-`<xxx>.<NN>.tgz`** was untarred.

See [Sample instances.yml File on page 77](#).

Example using Contrail Networking Release 2003:

```
cd contrail-ansible-deployer-2003.33/contrail-ansible-deployer/config/
```

```
vi contrail-ansible-deployer-2003.33/contrail-ansible-deployer/config/instances.yml
```

Sample **instances.yml** files for various other deployments are available at the same directory.

2. Update `CONTRAIL_VERSION` and `CONTRAIL_CONTAINER_TAG` to the desired version tag in this **instances.yml** file.

Access `CONTRAIL_CONTAINER_TAG` located at [README Access to Contrail Networking Registry 20xx](#).

For example:

```
CONTRAIL_VERSION = 2003.33
CONTRAIL_CONTAINER_TAG = 2003.33
```

3. Run the following commands from **contrail-ansible-deployer** directory.

- For Contrail with OpenStack deployment:

```
cd contrail-ansible-deployer
```

```
ansible-playbook -e orchestrator=openstack -i inventory/ playbooks/install_openstack.yml -v
```

```
ansible-playbook -e orchestrator=openstack -i inventory/ playbooks/install_contrail.yml -v
```

- For Contrail with Kubernetes deployment:

```
cd contrail-ansible-deployer
```

```
ansible-playbook -e orchestrator=kubernetes -i inventory/ playbooks/install_k8s.yml -v
```

```
ansible-playbook -e orchestrator=kubernetes -i inventory/ playbooks/install_contrail.yml -v
```

The ansible playbook logs are available on the terminal during execution. You can also access it at `/var/log/ansible.log`.

RELATED DOCUMENTATION

[Installing a Contrail Cluster using Contrail Command and instances.yml](#) | 74

[Upgrading Contrail Networking using Contrail Command](#) | 109

Upgrading Contrail Networking using In-Place Upgrade Procedure

NOTE: This procedure can be used for Contrail Networking upgrades to Contrail Networking Release 2003 or earlier only.

If you are upgrading to Contrail Networking Release 2005 or later using an in-place upgrade procedure, see [“How to Perform a Zero Impact Contrail Networking Upgrade using the Ansible Deployer” on page 114.](#)

This document provides steps to upgrade Contrail Networking using in-place upgrade procedure.

The procedure supports incremental model and you can use it to upgrade from Contrail Networking Release N-1 to N.

BEST PRACTICE: You must take snapshots of your current system before proceeding with the upgrade process.

For a list of supported platforms for all Contrail Networking releases, see [Contrail Networking Supported Platforms List](#).

1. Update kernel version on all the compute nodes.

```
yum -y update kernel*
```

NOTE: You must not update kernel version if you are upgrading from Contrail Networking Release 1910 to Release 1911.

2. Update `CONTRAIL_VERSION` and `CONTRAIL_CONTAINER_TAG` to the desired version tag in this `instances.yml` file.

Access `CONTRAIL_CONTAINER_TAG` located at [README Access to Contrail Networking Registry 20xx](#).

3. Run the following commands from `contrail-ansible-deployer` directory.

For Contrail with OpenStack deployment:

```
cd contrail-ansible-deployer
ansible-playbook -i inventory/ -e orchestrator=openstack playbooks/configure_instances.yml
ansible-playbook -e orchestrator=openstack -i inventory/ playbooks/install_contrail.yml
```

4. Reboot the compute node.
5. Check the status of Contrail service on the compute node.

All services must be *active* .

sudo contrail-status

The ansible playbook logs are available on the terminal during execution. You can also access it at `/var/log/ansible.log`.

RELATED DOCUMENTATION

[Upgrading Contrail Networking using Contrail Command | 109](#)

[Upgrading Contrail Networking using contrail-ansible Deployer | 136](#)

[Upgrading Contrail Networking with Red Hat Openstack 13 using ISSU](#)

Upgrading Contrail Networking Release 19xx with RHOSP13 to Contrail Networking Release 2011 with RHOSP16.1

The goal of this topic is to provide a combined procedure to upgrade Red Hat OpenStack Platform (RHOSP) from RHOSP 13 to RHOSP 16.1 by leveraging Red Hat Fast Forward Upgrade (FFU) procedure while simultaneously upgrading Contrail Networking from Release 19xx to Release 2011. The procedure leverages the Zero Impact Upgrade (ZIU) procedure from Contrail to minimize the downtime.

The downtime will be reduced by not requiring extra server reboots in addition to the ones that the RHOSP FFU procedure already requires for Kernel/RHEL upgrades.

Refer to [Red Hat OpenStack Framework for Upgrades \(13 to16.1\)](#) documentation for details on RHOSP 13 to RHOSP 16.1 Fast Forward Upgrade (FFU) procedure of OpenStack Platform environment from one long life version to the next long life version.

Access `ContrailImageTag` located at [README Access to Contrail Registry 20XX](#).

1. Follow *chapter 2—Planning and preparation for an in-place upgrade* through *chapter 8.3— Copying the Leapp data to the overcloud nodes* of [Red Hat OpenStack Framework for Upgrades \(13 to16.1\)](#) procedure.
2. Use predictable NIC names for overcloud nodes
 - a. Log in to the undercloud as the stack user.
 - b. Run the **playbook-nics.yaml** playbook on all overcloud nodes.

```
$ ansible-playbook -i ~/inventory.yaml playbook-nics.yaml
```

The playbook sets the new NIC prefix to **em**. To set a different NIC prefix, set the prefix variable when running the playbook:

```
$ ansible-playbook -i ~/inventory.yaml -e prefix="mynic" playbook-nics.yaml
```

- c. Run the **playbook-nics-vhost0.yaml** playbook on all overcloud compute nodes.

```
$ ansible-playbook -i inventory.yaml -l overcloud_Compute $my_dir/playbook-nics-vhost0.yaml
```

For details, see

https://github.com/tungstenfabric/tf-deployment-test/blob/master/rhosp/ffu_ziu_13_16/tf_specific/playbook-nics-vhost0.yaml.

- d. Reboot overcloud nodes using the standard reboot procedures. For more information, see [Red Hat rebooting nodes](#) documentation.

3. Follow *chapter 8.5—Copying the Leapp data to the overcloud nodes* through *chapter 19.2—Upgrading Controller nodes with external Ceph deployments* of [Red Hat OpenStack Framework for Upgrades \(13 to 16.1\)](#) procedure.

4. Upgrade compute nodes to OpenStack Platform 16.1.

NOTE: If you are not using the default stack name (overcloud), set your stack name with the **--stack STACK NAME** option replacing STACK NAME with the name of your stack.

- a. Source the **stackrc** file.

```
$ source ~/stackrc
```

- b. Migrate your instances. For details, see [RedHat Migrating virtual machine instances between Compute nodes](#) documentation.

- c. Run the upgrade command with the **system_upgrade_prepare** and **system_upgrade_run** tags.

```
$ openstack overcloud upgrade run --stack STACK NAME --tags system_upgrade_prepare --limit overcloud-compute-0
```

```
$ openstack overcloud upgrade run --stack STACK NAME --tags system_upgrade_run --limit overcloud-compute-0
```

To upgrade multiple Compute nodes in parallel, set the **--limit** option to a comma-separated list of nodes that you want to upgrade.

- d. Run the upgrade command with no tags to perform Red Hat OpenStack Platform upgrade.

```
$ openstack overcloud upgrade run --stack STACK NAME --limit overcloud-compute-0
```

5. Follow *chapter 19.4—Synchronizing the overcloud stack* through *chapter 21.4—Synchronizing the overcloud stack* of [Red Hat OpenStack Framework for Upgrades \(13 to 16.1\)](#) procedure to complete the upgrade.

RELATED DOCUMENTATION

[Understanding Red Hat OpenStack Platform Director | 211](#)

[Setting Up the Infrastructure | 216](#)

[Setting Up the Undercloud | 225](#)

[Setting Up the Overcloud | 227](#)

Backup and Restore Contrail Software

IN THIS CHAPTER

- [How to Backup and Restore Contrail Databases in JSON Format | 142](#)

How to Backup and Restore Contrail Databases in JSON Format

IN THIS SECTION

- [Before You Begin | 142](#)
- [Simple Database Backup in JSON Format | 143](#)
- [Examples: Simple Database Backups in JSON Format | 147](#)
- [Restore Database from the Backup in JSON Format | 150](#)
- [Example: How to Restore a Database Using the JSON Backup \(Ansible Deployer Environment\) | 156](#)
- [Example: How to Restore a Database Using the JSON Backup \(Red Hat Openstack Deployer Environment\) | 160](#)

This document shows how to backup and restore the Contrail databases—Cassandra and Zookeeper—in JSON format.

Before You Begin

The backup and restore procedure must be completed for nodes running the same Contrail Networking release. The procedure is used to backup the Contrail Networking databases only; it does not include instructions for backing up orchestration system databases.



CAUTION: Database backups must be consistent across all systems because the state of the Contrail database is associated with other system databases, such as OpenStack databases. Database changes associated with northbound APIs must be stopped on all the systems before performing any backup operation. For example, you might block the external VIP for northbound APIs at the load balancer level, such as HAproxy.

Simple Database Backup in JSON Format

This procedure provides a simple database backup in JSON format. This procedure is performed using the **db_json_exim.py** script located in the **/usr/lib/python2.7/site-packages/cfgm_common** on the controller node.

To perform this database backup:

1. Log into one of the config nodes. Create the **/tmp/db-dump** directory on any of the config node hosts.

```
mkdir /tmp/db-dump
```

2. On the same config node, copy the **contrail-api.conf** file from the container to the host.

Ansible Deployer:

```
docker cp config_api_1:/etc/contrail/contrail-api.conf /tmp/db-dump/
```

Red Hat Openstack Deployer:

```
docker cp contrail_config_api:/etc/contrail/contrail-api.conf /tmp/db-dump/
```

The Cassandra database instance on any configuration node includes the complete Cassandra database for all configuration nodes in the cluster. Steps 1 and 2, therefore, only need to be performed on one configuration node.

3. Stop the following docker configuration services on all of the Contrail configuration nodes.

Ansible Deployer:

```
docker stop config_svcmonitor_1
docker stop config_devicemgr_1
```

```
docker stop config_schema_1
docker stop config_api_1
```

Red Hat Openstack Deployer:

```
docker stop contrail_config_svc_monitor
docker stop contrail_config_device_manager
docker stop contrail_config_schema
docker stop contrail_config_api
```

This step must be performed on each individual config node in the cluster.

4. Return to the config node where you performed steps 1 and 2.

List the docker image to find the name or ID of the *config api* image.

docker image ls | grep config-api

Example:

```
docker image ls | grep config-api
hub.juniper.net/contrail/contrail-controller-config-api 1909.30-ocata c9d757252a0c
4 months ago 583MB
```

5. From the same config node, start the *config api* container pointing the **entrypoint.sh** script to **/bin/bash** and mapping **/tmp/db-dump** from the host to the **/tmp** directory inside the container. You perform this step to ensure that the API services are not started on the config node.

Enter the **-v /etc/contrail/ssl:/etc/contrail/ssl:ro** command option when **cassandra_use_ssl** is used as api-server configuration parameter to ensure TLS certificates are mounted to the Contrail SSL directory. This mounting ensures that the backup procedure succeeds in environments with endpoints that require TLS authentication.

The *registry_name* and *container_tag* variables must match step 4.

```
docker run --rm -it -v /tmp/db-dump:/tmp/ -v
/etc/contrail/ssl:/etc/contrail/ssl:ro --network host --entrypoint=/bin/bash
<registry_name>/contrail-controller-config-api:<container_tag>
```

Example:


```
docker run --rm -it -v /tmp/db-dump:/tmp/ -v
/etc/contrail/ssl:/etc/contrail/ssl:ro --network host --entrypoint=/bin/bash
hub.juniper.net/contrail/contrail-controller-config-api:1909.30-ocata
```

6. From the docker container created on the config node in Step 5, use the **db_json_exim.py** script to backup data in JSON format.. The db dump file will be saved in the **/tmp/db-dump/** on this config node.

```
cd /usr/lib/python2.7/site-packages/cfgm_common
python db_json_exim.py --export-to /tmp/db-dump.json --api-conf
/tmp/contrail-api.conf
```

The Cassandra database instance on any configuration node includes the complete Cassandra database for all configuration nodes in the cluster. You, therefore, only need to perform step 4 through 6 from one of the configuration nodes.

7. (Optional. Recommended) From the same config node, enter the **cat /tmp/db-dump.json | python -m json.tool | less** command to view a more readable version of the file transfer.

```
cat /tmp/db-dump.json | python -m json.tool | less
```

8. From the same config node, exit out of the *config api* container. This will stop the container.

```
exit
```

9. Start the following configuration services on all of the Contrail configuration nodes.

Ansible Deployer:

```
docker start config_api_1
docker start config_schema_1
docker start config_svcmonitor_1
docker start config_devicemgr_1
```

Red Hat Openstack Deployer:

```
docker start contrail_config_api
docker start contrail_config_schema
```

```
docker start contrail_config_svc_monitor
docker start contrail_config_device_manager
```

This step must be performed on each individual config node.

10. On each config node, enter the **contrail-status** command to confirm that services are in the *active* or *running* states.

NOTE: Some command output and output fields are removed for readability. Output shown is from a node hosting config and analytics services.

```
contrail-status
Pod           Service      Original Name      State
analytics     api          contrail-analytics-api    running
analytics     collector    contrail-analytics-collector running
analytics     nodemgr      contrail-nodemgr         running
analytics     provisioner  contrail-provisioner      running
analytics     redis       contrail-external-redis   running
analytics-alarm alarm-gen    contrail-analytics-alarm-gen running
analytics-alarm kafka       contrail-external-kafka    running
<some output removed for readability>

== Contrail control ==
control: active
nodemgr: active
named: active
dns: active

== Contrail analytics-alarm ==
nodemgr: active
kafka: active
alarm-gen: active

== Contrail database ==
nodemgr: active
query-engine: active
cassandra: active

== Contrail analytics ==
nodemgr: active
```

```

api: active
collector: active

== Contrail config-database ==
nodemgr: active
zookeeper: active
rabbitmq: active
cassandra: active

== Contrail webui ==
web: active
job: active

== Contrail analytics-snmp ==
snmp-collector: active
nodemgr: active
topology: active

== Contrail config ==
svc-monitor: active
nodemgr: active
device-manager: active
api: active
schema: active

```

Examples: Simple Database Backups in JSON Format

These examples illustrate the process for creating a simple database backup in JSON format in both an Ansible deployer environment and a Red Hat Openstack deployer environment.

In each example, a cluster with three config nodes—**control_config1**, **control_config2**, and **control_config3**—is backed up. All tasks that need to be performed on a single config nodes are performed on **control-config1**. The tasks must be performed in the shown order.

Ansible Deployer Environment:

```

## control_config1 ##
mkdir /tmp/db-dump
docker cp config_api_1:/etc/contrail/contrail-api.conf /tmp/db-dump/
docker stop config_svcmonitor_1
docker stop config_devicemgr_1
docker stop config_schema_1

```

```

docker stop config_api_1

## control_config2 ##
docker stop config_svcmonitor_1
docker stop config_devicemgr_1
docker stop config_schema_1
docker stop config_api_1

## control_config3 ##
docker stop config_svcmonitor_1
docker stop config_devicemgr_1
docker stop config_schema_1
docker stop config_api_1

## control_config1 ##
docker run --rm -it -v /tmp/db-dump:/tmp/ -v /etc/contrail/ssl:/etc/contrail/ssl:ro
--network host --entrypoint=/bin/bash
hub.juniper.net/contrail/contrail-controller-config-api:1909.30-ocata
cd /usr/lib/python2.7/site-packages/cfgm_common
python db_json_exim.py --export-to /tmp/db-dump.json --api-conf
/tmp/contrail-api.conf
cat /tmp/db-dump.json | python -m json.tool | less
exit
docker start config_api_1
docker start config_schema_1
docker start config_svcmonitor_1
docker start config_devicemgr_1
contrail-status

## control_config2 ##
docker start config_api_1
docker start config_schema_1
docker start config_svcmonitor_1
docker start config_devicemgr_1
contrail-status

## control_config3 ##
docker start config_api_1
docker start config_schema_1
docker start config_svcmonitor_1
docker start config_devicemgr_1
contrail-status

```

Red Hat Openstack Deployer Environment:

```

## control_config1 ##
mkdir /tmp/db-dump
docker cp contrail_config_api:/etc/contrail/contrail-api.conf /tmp/db-dump/
docker stop contrail_config_svc_monitor
docker stop contrail_config_device_manager
docker stop contrail_config_schema
docker stop contrail_config_api

## control_config2 ##
docker stop contrail_config_svc_monitor
docker stop contrail_config_device_manager
docker stop contrail_config_schema
docker stop contrail_config_api

## control_config3 ##
docker stop contrail_config_svc_monitor
docker stop contrail_config_device_manager
docker stop contrail_config_schema
docker stop contrail_config_api

## control_config1 ##
docker run --rm -it -v /tmp/db-dump:/tmp/ -v /etc/contrail/ssl:/etc/contrail/ssl:ro
--network host --entrypoint=/bin/bash
hub.juniper.net/contrail/contrail-controller-config-api:1909.30-ocata
cd /usr/lib/python2.7/site-packages/cfgm_common
python db_json_exim.py --export-to /tmp/db-dump.json --api-conf
/tmp/contrail-api.conf
cat /tmp/db-dump.json | python -m json.tool | less
exit
docker start contrail_config_api
docker start contrail_config_schema
docker start contrail_config_svc_monitor
docker start contrail_config_device_manager
contrail-status

## control_config2 ##
docker start contrail_config_api
docker start contrail_config_schema
docker start contrail_config_svc_monitor
docker start contrail_config_device_manager
contrail-status

## control_config3 ##

```

```
docker start contrail_config_api
docker start contrail_config_schema
docker start contrail_config_svc_monitor
docker start contrail_config_device_manager
contrail-status
```

Restore Database from the Backup in JSON Format

This procedure provides the steps to restore a system using the simple database backup JSON file that was created in [“Simple Database Backup in JSON Format” on page 143](#).

To restore a system from a backup JSON file:

1. Copy the **contrail-api.conf** file from the container to the host on any one of the config nodes.

Ansible Deployer:

```
docker cp config_api_1:/etc/contrail/contrail-api.conf /tmp/db-dump/
```

Red Hat Openstack Deployer:

```
docker cp contrail_config_api:/etc/contrail/contrail-api.conf /tmp/db-dump/
```

2. Stop the configuration services on all of the controllers.

Ansible Deployer:

```
docker stop config_svcmonitor_1
docker stop config_devicemgr_1
docker stop config_schema_1
docker stop config_api_1
docker stop config_nodemgr_1
docker stop config_database_nodemgr_1
docker stop analytics_snmp_snmp-collector_1
docker stop analytics_snmp_topology_1
docker stop analytics_alarm_alarm-gen_1
docker stop analytics_api_1
docker stop analytics_collector_1
docker stop analytics_alarm_kafka_1
```

Red Hat Openstack Deployer—Node hosting Contrail Config containers:

```
docker stop contrail_config_svc_monitor
docker stop contrail_config_device_manager
docker stop contrail_config_schema
docker stop contrail_config_api
docker stop contrail_config_nodemgr
docker stop contrail_config_database_nodemgr
```

Red Hat Openstack Deployer—Node hosting Contrail Analytics containers:

```
docker stop contrail_analytics_snmp_collector
docker stop contrail_analytics_topology
docker stop contrail_analytics_alarmgen
docker stop contrail_analytics_api
docker stop contrail_analytics_collector
docker stop contrail_analytics_kafka
```

3. Stop the Cassandra service on all the **config-db** controllers.

Ansible Deployer:

```
docker stop config_database_cassandra_1
```

Red Hat Openstack Deployer:

```
docker stop contrail_config_database
```

4. Stop the Zookeeper service on all controllers.

Ansible Deployer:

```
docker stop config_database_zookeeper_1
```

Red Hat Openstack Deployer:

```
docker stop contrail_config_zookeeper
```

5. Backup the Zookeeper data directory on all the controllers.

Ansible Deployer:

```
cd /var/lib/docker/volumes/config_database_config_zookeeper/
cp -R _data/version-2/ version-2-save
```

Red Hat Openstack Deployer:

```
cd /var/lib/docker/volumes/config_zookeeper/
cp -R _data/version-2/ version-2-save
```

6. Delete the Zookeeper data directory contents on all the controllers.

```
rm -rf _data/version-2/*
```

7. Backup the Cassandra data directory on all the controllers.

Ansible Deployer:

```
cd /var/lib/docker/volumes/config_database_config_cassandra/
cp -R _data/ Cassandra_data-save
```

Red Hat Openstack Deployer:

```
cd /var/lib/docker/volumes/config_cassandra/
cp -R _data/ Cassandra_data-save
```

8. Delete the Cassandra data directory contents on all controllers.

```
rm -rf _data/*
```

9. Start the Zookeeper service on all the controllers.

Ansible Deployer:

```
docker start config_database_zookeeper_1
```

Red Hat Openstack Deployer:

```
docker start contrail_config_zookeeper
```


10. Start the Cassandra service on all the controllers.

Ansible Deployer:

```
docker start config_database_cassandra_1
```

Red Hat Openstack Deployer:

```
docker start contrail_config_database
```

11. List docker image to find the name or ID of the **config-api** image on the config node.

```
docker image ls | grep config-api
```

Example:

```
docker image ls | grep config-api
hub.juniper.net/contrail/contrail-controller-config-api 1909.30-ocata c9d757252a0c
4 months ago 583MB
```

12. Run a new docker container using the name or ID of the **config_api** image on the same config node.

Enter the `-v /etc/contrail/ssl:/etc/contrail/ssl:ro` command option when **cassandra_use_ssl** is used as api-server configuration parameter to ensure TLS certificates are mounted to the Contrail SSL directory. This mounting ensures that this backup procedure succeeds in environments with endpoints that require TLS authentication.

Use the *registry_name* and *container_tag* from the output of the step [11](#).

```
docker run --rm -it -v /tmp/db-dump:/tmp/ -v
/etc/contrail/ssl:/etc/contrail/ssl:ro --network host --entrypoint=/bin/bash
<registry_name>/contrail-controller-config_api:<container tag>
```

Example

```
docker run --rm -it -v /tmp/db-dump:/tmp/ -v
/etc/contrail/ssl:/etc/contrail/ssl:ro --network host --entrypoint=/bin/bash
hub.juniper.net/contrail/contrail-controller-config-api:1909.30-ocata
```

13. Restore the data in new running docker on the same config node.

```
cd /usr/lib/python2.7/site-packages/cfgm_common
python db_json_exim.py --import-from /tmp/db-dump.json --api-conf
/tmp/contrail-api.conf
```

14. Exit out of the *config api* container. This will stop the container.

```
exit
```

15. Start config services on all the controllers.

Ansible Deployer:

```
docker start config_svcmonitor_1
docker start config_devicemgr_1
docker start config_schema_1
docker start config_api_1
docker start config_nodemgr_1
docker start config_database_nodemgr_1
docker start analytics_snmp_snmp-collector_1
docker start analytics_snmp_topology_1
docker start analytics_alarm_alarm-gen_1
docker start analytics_api_1
docker start analytics_collector_1
docker start analytics_alarm_kafka_1
```

Red Hat Openstack Deployer—Node hosting Contrail Config containers:

```
docker start contrail_config_svc_monitor
docker start contrail_config_device_manager
docker start contrail_config_schema
docker start contrail_config_api
docker start contrail_config_nodemgr
docker start contrail_config_database_nodemgr
```

Red Hat Openstack Deployer—Node hosting Contrail Analytics containers:

```
docker start contrail_analytics_snmp_collector
docker start contrail_analytics_topology
docker start contrail_analytics_alarmgen
docker start contrail_analytics_api
```

```
docker start contrail_analytics_collector
docker start contrail_analytics_kafka
```

16. Enter the **contrail-status** command on each configuration node and, when applicable, on each analytics node to confirm that services are in the *active* or *running* states.

NOTE: Output shown for a config node. Some command output and output fields are removed for readability.

```
contrail-status
Pod      Service      Original Name      State
config  api          contrail-controller-config-api    running
config  device-manager  contrail-controller-config-devicemgr  running
config  dnsmasq      contrail-controller-config-dnsmasq  running
config  nodemgr      contrail-nodemgr      running
config  provisioner   contrail-provisioner   running
config  schema       contrail-controller-config-schema  running
config  stats        contrail-controller-config-stats    running
<some output removed for readability>

== Contrail control ==
control: active
nodemgr: active
named: active
dns: active

== Contrail database ==
nodemgr: active
query-engine: active
cassandra: active

== Contrail config-database ==
nodemgr: active
zookeeper: active
rabbitmq: active
cassandra: active

== Contrail webui ==
web: active
```

```

job: active

== Contrail config ==
svc-monitor: active
nodemgr: active
device-manager: active
api: active
schema: active

```

Example: How to Restore a Database Using the JSON Backup (Ansible Deployer Environment)

This example shows how to restore the databases for three controllers connected to the Contrail Configuration database (config-db). This example assumes a JSON backup file of the databases was previously created using the instructions provided in [“Simple Database Backup in JSON Format” on page 143](#). The network was deployed using Ansible and the three controllers—nodec53, nodec54, and nodec55—have separate IP addresses.

```

## Make db-dump directory. Copy contrail-api.conf to db-dump directory. ##
root@nodec54 ~]# mkdir /tmp/db-dump
root@nodec54 ~]# docker cp config_api_1:/etc/contrail/contrail-api.conf
/tmp/db-dump/

## Stop Configuration Services on All Controllers ##
[root@nodec53 ~]# docker stop config_schema_1
[root@nodec53 ~]# docker stop config_api_1
[root@nodec53 ~]# docker stop config_svcmonitor_1
[root@nodec53 ~]# docker stop config_devicemgr_1
[root@nodec53 ~]# docker stop config_nodemgr_1
[root@nodec53 ~]# docker stop config_database_nodemgr_1
[root@nodec53 ~]# docker stop analytics_snmp_snmp-collector_1
[root@nodec53 ~]# docker stop analytics_snmp_topology_1
[root@nodec53 ~]# docker stop analytics_alarm_alarm-gen_1
[root@nodec53 ~]# docker stop analytics_api_1
[root@nodec53 ~]# docker stop analytics_collector_1
[root@nodec53 ~]# docker stop analytics_alarm_kafka_1

[root@nodec54 ~]# # docker stop config_schema_1
[root@nodec54 ~]# docker stop config_api_1
[root@nodec54 ~]# docker stop config_svcmonitor_1
[root@nodec54 ~]# docker stop config_devicemgr_1
[root@nodec54 ~]# docker stop config_nodemgr_1
[root@nodec54 ~]# docker stop config_database_nodemgr_1

```

```

[root@nodec54 ~]# docker stop analytics_snmp_snmp-collector_1
[root@nodec54 ~]# docker stop analytics_snmp_topology_1
[root@nodec54 ~]# docker stop analytics_alarm_alarm-gen_1
[root@nodec54 ~]# docker stop analytics_api_1
[root@nodec54 ~]# docker stop analytics_collector_1
[root@nodec54 ~]# docker stop analytics_alarm_kafka_1

[root@nodec55 ~]# docker stop config_schema_1
[root@nodec55 ~]# docker stop config_api_1
[root@nodec55 ~]# docker stop config_svcmonitor_1
[root@nodec55 ~]# docker stop config_devicemgr_1
[root@nodec55 ~]# docker stop config_nodemgr_1
[root@nodec55 ~]# docker stop config_database_nodemgr_1
[root@nodec55 ~]# docker stop analytics_snmp_snmp-collector_1
[root@nodec55 ~]# docker stop analytics_snmp_topology_1
[root@nodec55 ~]# docker stop analytics_alarm_alarm-gen_1
[root@nodec55 ~]# docker stop analytics_api_1
[root@nodec55 ~]# docker stop analytics_collector_1
[root@nodec55 ~]# docker stop analytics_alarm_kafka_1

## Stop Cassandra ##
[root@nodec53 ~]# docker stop config_database_cassandra_1
[root@nodec54 ~]# docker stop config_database_cassandra_1
[root@nodec55 ~]# docker stop config_database_cassandra_1

## Stop Zookeeper ##
[root@nodec53 ~]# docker stop config_database_zookeeper_1
[root@nodec54 ~]# docker stop config_database_zookeeper_1
[root@nodec55 ~]# docker stop config_database_zookeeper_1

## Backup Zookeeper Directories Before Deleting Zookeeper Data Directory Contents
##
[root@nodec53 _data]# cd /var/lib/docker/volumes/config_database_config_zookeeper/
[root@nodec53 config_database_config_zookeeper]# cp -R _data/version-2/
version-2-save
[root@nodec53 config_database_config_zookeeper]# rm -rf _data/version-2/*

[root@nodec54 _data]# cd /var/lib/docker/volumes/config_database_config_zookeeper/
[root@nodec54 config_database_config_zookeeper]# cp -R _data/version-2/
version-2-save
[root@nodec54 config_database_config_zookeeper]# rm -rf _data/version-2/*

[root@nodec55 _data]# cd /var/lib/docker/volumes/config_database_config_zookeeper/
[root@nodec55 config_database_config_zookeeper]# cp -R _data/version-2/

```

```

version-2-save
[root@nodec55 config_database_config_zookeeper]# rm -rf _data/version-2/*

## Backup Cassandra Directory Before Deleting Cassandra Data Directory Contents
##
[root@nodec53 ~]# cd /var/lib/docker/volumes/config_database_config_cassandra/
[root@nodec53 config_database_config_cassandra]# cp -R _data/ Cassandra_data-save
[root@nodec53 config_database_config_cassandra]# rm -rf _data/*

[root@nodec54 ~]# cd /var/lib/docker/volumes/config_database_config_cassandra/
[root@nodec54 config_database_config_cassandra]# cp -R _data/ Cassandra_data-save
[root@nodec54 config_database_config_cassandra]# rm -rf _data/*

[root@nodec55 ~]# cd /var/lib/docker/volumes/config_database_config_cassandra/
[root@nodec55 config_database_config_cassandra]# cp -R _data/ Cassandra_data-save
[root@nodec55 config_database_config_cassandra]# rm -rf _data/*

## Start Zookeeper ##
[root@nodec53 ~]# docker start config_database_zookeeper_1
[root@nodec54 ~]# docker start config_database_zookeeper_1
[root@nodec55 ~]# docker start config_database_zookeeper_1

## Start Cassandra ##
[root@nodec53 ~]# docker start config_database_cassandra_1
[root@nodec54 ~]# docker start config_database_cassandra_1
[root@nodec55 ~]# docker start config_database_cassandra_1

## Run Docker Image & Mount Contrail TLS Certificates to Contrail SSL Directory
##
[root@nodec54 ~]# docker image ls | grep config-api
hub.juniper.net/contrail/contrail-controller-config-api 1909.30-ocata c9d757252a0c
4 months ago 583MB
[root@nodec54 ~]# docker run --rm -it -v /tmp/db-dump:/tmp/ -v
/etc/contrail/ssl:/etc/contrail/ssl:ro --network host --entrypoint=/bin/bash
hub.juniper.net/contrail/contrail-controller-config-api:1909.30-ocata

## Restore Data in New Docker Containers ##
(config_api_1)[root@nodec54 /root]$ cd /usr/lib/python2.7/site-packages/cfgm_common/
(config_api_1)[root@nodec54 /usr/lib/python2.7/site-packages/cfgm_common]$ python
db_json_exim.py --import-from /tmp/db-dump.json --api-conf /tmp/contrail-api.conf

## Start Configuration Services ##
[root@nodec53 ~]# docker start config_schema_1
[root@nodec53 ~]# docker start config_svcmonitor_1

```

```
[root@nodec53 ~]# docker start config_devicemgr_1
[root@nodec53 ~]# docker start config_nodemgr_1
[root@nodec53 ~]# docker start config_database_nodemgr_1
[root@nodec53 ~]# docker start config_api_1
[root@nodec53 ~]# docker start analytics_snmp_snmp-collector_1
[root@nodec53 ~]# docker start analytics_snmp_topology_1
[root@nodec53 ~]# docker start analytics_alarm_alarm-gen_1
[root@nodec53 ~]# docker start analytics_api_1
[root@nodec53 ~]# docker start analytics_collector_1
[root@nodec53 ~]# docker start analytics_alarm_kafka_1
```

```
[root@nodec54 ~]# docker start config_schema_1
[root@nodec54 ~]# docker start config_svcmonitor_1
[root@nodec54 ~]# docker start config_devicemgr_1
[root@nodec54 ~]# docker start config_nodemgr_1
[root@nodec54 ~]# docker start config_database_nodemgr_1
[root@nodec54 ~]# docker start config_api_1
[root@nodec54 ~]# docker start analytics_snmp_snmp-collector_1
[root@nodec54 ~]# docker start analytics_snmp_topology_1
[root@nodec54 ~]# docker start analytics_alarm_alarm-gen_1
[root@nodec54 ~]# docker start analytics_api_1
[root@nodec54 ~]# docker start analytics_collector_1
[root@nodec54 ~]# docker start analytics_alarm_kafka_1
```

```
[root@nodec55 ~]# docker start config_schema_1
[root@nodec55 ~]# docker start config_svcmonitor_1
[root@nodec55 ~]# docker start config_devicemgr_1
[root@nodec55 ~]# docker start config_nodemgr_1
[root@nodec55 ~]# docker start config_database_nodemgr_1
[root@nodec55 ~]# docker start config_api_1
[root@nodec55 ~]# docker start analytics_snmp_snmp-collector_1
[root@nodec55 ~]# docker start analytics_snmp_topology_1
[root@nodec55 ~]# docker start analytics_alarm_alarm-gen_1
[root@nodec55 ~]# docker start analytics_api_1
[root@nodec55 ~]# docker start analytics_collector_1
[root@nodec55 ~]# docker start analytics_alarm_kafka_1
```

Confirm Services are Active

```
[root@nodec53 ~]# contrail-status
[root@nodec54 ~]# contrail-status
[root@nodec55 ~]# contrail-status
```

Example: How to Restore a Database Using the JSON Backup (Red Hat Openstack Deployer Environment)

This example shows how to restore the databases from an environment that was deployed using Red Hat Openstack and includes three config nodes—*config1*, *config2*, and *config3*—connected to the Contrail Configuration database (config-db). All steps that need to be done from a single config node are performed from *config1*.

The environment also contains three analytics nodes—*analytics1*, *analytics2*, and *analytics3*—to provide analytics services.

This example assumes a JSON backup file of the databases was previously created using the instructions provided in [“Simple Database Backup in JSON Format” on page 143](#).

```
## Make db-dump directory. Copy contrail-api.conf to db-dump directory. ##
[root@config1 ~]# mkdir /tmp/db-dump
[root@config1 ~]# docker cp config_api_1:/etc/contrail/contrail-api.conf
/tmp/db-dump/

## Stop Configuration Services on All Config Nodes ##
[root@config1 ~]# docker stop contrail_config_svc_monitor
[root@config1 ~]# docker stop contrail_config_device_manager
[root@config1 ~]# docker stop contrail_config_schema
[root@config1 ~]# docker stop contrail_config_api
[root@config1 ~]# docker stop contrail_config_nodemgr
[root@config1 ~]# docker stop contrail_config_database_nodemgr

[root@config2 ~]# docker stop contrail_config_svc_monitor
[root@config2 ~]# docker stop contrail_config_device_manager
[root@config2 ~]# docker stop contrail_config_schema
[root@config2 ~]# docker stop contrail_config_api
[root@config2 ~]# docker stop contrail_config_nodemgr
[root@config2 ~]# docker stop contrail_config_database_nodemgr

[root@config3 ~]# docker stop contrail_config_svc_monitor
[root@config3 ~]# docker stop contrail_config_device_manager
[root@config3 ~]# docker stop contrail_config_schema
[root@config3 ~]# docker stop contrail_config_api
[root@config3 ~]# docker stop contrail_config_nodemgr
[root@config3 ~]# docker stop contrail_config_database_nodemgr

## Stop Analytics Services on All Analytics Nodes ##
[root@analytics1 ~]# docker stop contrail_analytics_snmp_collector
[root@analytics1 ~]# docker stop contrail_analytics_topology
```



```

[root@analytics1 ~]# docker stop contrail_analytics_alarmgen
[root@analytics1 ~]# docker stop contrail_analytics_api
[root@analytics1 ~]# docker stop contrail_analytics_collector
[root@analytics1 ~]# docker stop contrail_analytics_kafka

[root@analytics2 ~]# docker stop contrail_analytics_snmp_collector
[root@analytics2 ~]# docker stop contrail_analytics_topology
[root@analytics2 ~]# docker stop contrail_analytics_alarmgen
[root@analytics2 ~]# docker stop contrail_analytics_api
[root@analytics2 ~]# docker stop contrail_analytics_collector
[root@analytics2 ~]# docker stop contrail_analytics_kafka

[root@analytics3 ~]# docker stop contrail_analytics_snmp_collector
[root@analytics3 ~]# docker stop contrail_analytics_topology
[root@analytics3 ~]# docker stop contrail_analytics_alarmgen
[root@analytics3 ~]# docker stop contrail_analytics_api
[root@analytics3 ~]# docker stop contrail_analytics_collector
[root@analytics3 ~]# docker stop contrail_analytics_kafka

## Stop Cassandra ##
[root@config1 ~]# docker stop contrail_config_database
[root@config2 ~]# docker stop contrail_config_database
[root@config3 ~]# docker stop contrail_config_database

## Stop Zookeeper ##
[root@config1 ~]# docker stop contrail_config_zookeeper
[root@config2 ~]# docker stop contrail_config_zookeeper
[root@config3 ~]# docker stop contrail_config_zookeeper

## Backup Zookeeper Directories Before Deleting Zookeeper Data Directory Contents
##
[root@config1 _data]# cd /var/lib/docker/volumes/config_zookeeper/
[root@config1 config_zookeeper]# cp -R _data/version-2/ version-2-save
[root@config1 config_zookeeper]# rm -rf _data/version-2/*
[root@config2 _data]# cd /var/lib/docker/volumes/config_zookeeper/
[root@config2 config_zookeeper]# cp -R _data/version-2/ version-2-save
[root@config2 config_zookeeper]# rm -rf _data/version-2/*
[root@config3 _data]# cd /var/lib/docker/volumes/config_zookeeper/
[root@config3 config_zookeeper]# cp -R _data/version-2/ version-2-save
[root@config3 config_zookeeper]# rm -rf _data/version-2/*

## Backup Cassandra Directory Before Deleting Cassandra Data Directory Contents
##
[root@config1 ~]# cd /var/lib/docker/volumes/config_cassandra/

```

```

[root@config1 config_cassandra]# cp -R _data/ Cassandra_data-save
[root@config1 config_cassandra]# rm -rf _data/*

[root@config2 ~]# cd /var/lib/docker/volumes/config_cassandra/
[root@config2 config_cassandra]# cp -R _data/ Cassandra_data-save
[root@config2 config_cassandra]# rm -rf _data/*

[root@config3 ~]# cd /var/lib/docker/volumes/config_cassandra/
[root@config3 config_cassandra]# cp -R _data/ Cassandra_data-save
[root@config3 config_cassandra]# rm -rf _data/*

## Start Zookeeper ##
[root@config1 ~]# docker start contrail_config_zookeeper
[root@config2 ~]# docker start contrail_config_zookeeper
[root@config3 ~]# docker start contrail_config_zookeeper

## Start Cassandra ##
[root@config1 ~]# docker start contrail_config_database
[root@config2 ~]# docker start contrail_config_database
[root@config3 ~]# docker start contrail_config_database

## Run Docker Image & Mount Contrail TLS Certificates to Contrail SSL Directory
##
[root@config1 ~]# docker image ls | grep config-api
hub.juniper.net/contrail/contrail-controller-config-api 1909.30-ocata c9d757252a0c
4 months ago 583MB
[root@config1 ~]# docker run --rm -it -v /tmp/db-dump:/tmp/ -v
/etc/contrail/ssl:/etc/contrail/ssl:ro --network host --entrypoint=/bin/bash
hub.juniper.net/contrail/contrail-controller-config-api:1909.30-ocata

## Restore Data in New Docker Containers ##
(config_api_1)[root@config1 /root]$ cd /usr/lib/python2.7/site-packages/cfgm_common/
(config_api_1)[root@config1 /usr/lib/python2.7/site-packages/cfgm_common]$ python
db_json_exim.py --import-from /tmp/db-dump.json --api-conf /tmp/contrail-api.conf

## Start Configuration Services on All Config Nodes ##
[root@config1 ~]# docker start contrail_config_svc_monitor
[root@config1 ~]# docker start contrail_config_device_manager
[root@config1 ~]# docker start contrail_config_schema
[root@config1 ~]# docker start contrail_config_api
[root@config1 ~]# docker start contrail_config_nodemgr
[root@config1 ~]# docker start contrail_config_database_nodemgr

[root@config2 ~]# docker start contrail_config_svc_monitor

```

```

[root@config2 ~]# docker start contrail_config_device_manager
[root@config2 ~]# docker start contrail_config_schema
[root@config2 ~]# docker start contrail_config_api
[root@config2 ~]# docker start contrail_config_nodemgr
[root@config2 ~]# docker start contrail_config_database_nodemgr

[root@config3 ~]# docker start contrail_config_svc_monitor
[root@config3 ~]# docker start contrail_config_device_manager
[root@config3 ~]# docker start contrail_config_schema
[root@config3 ~]# docker start contrail_config_api
[root@config3 ~]# docker start contrail_config_nodemgr
[root@config3 ~]# docker start contrail_config_database_nodemgr

## Start Configuration Services on All Analytics Nodes ##
[root@analytics1 ~]# docker start contrail_analytics_snmp_collector
[root@analytics1 ~]# docker start contrail_analytics_topology
[root@analytics1 ~]# docker start contrail_analytics_alarmgen
[root@analytics1 ~]# docker start contrail_analytics_api
[root@analytics1 ~]# docker start contrail_analytics_collector
[root@analytics1 ~]# docker start contrail_analytics_kafka

[root@analytics2 ~]# docker start contrail_analytics_snmp_collector
[root@analytics2 ~]# docker start contrail_analytics_topology
[root@analytics2 ~]# docker start contrail_analytics_alarmgen
[root@analytics2 ~]# docker start contrail_analytics_api
[root@analytics2 ~]# docker start contrail_analytics_collector
[root@analytics2 ~]# docker start contrail_analytics_kafka

[root@analytics3 ~]# docker start contrail_analytics_snmp_collector
[root@analytics3 ~]# docker start contrail_analytics_topology
[root@analytics3 ~]# docker start contrail_analytics_alarmgen
[root@analytics3 ~]# docker start contrail_analytics_api
[root@analytics3 ~]# docker start contrail_analytics_collector
[root@analytics3 ~]# docker start contrail_analytics_kafka

## Confirm Services are Active ##
[root@config1 ~]# contrail-status
[root@config2 ~]# contrail-status
[root@config3 ~]# contrail-status

[root@analytics1 ~]# contrail-status
[root@analytics2 ~]# contrail-status
[root@analytics3 ~]# contrail-status

```

Using Contrail Networking with VMware vCenter

IN THIS CHAPTER

- [VMware vCenter with Containerized Contrail Networking | 164](#)

VMware vCenter with Containerized Contrail Networking

IN THIS SECTION

- [Integrating vCenter for Contrail | 164](#)
- [Installing and Provisioning Contrail VMware vRealize Orchestrator Plugin | 172](#)
- [Integrating Contrail Release 5.0.X with VMware vRealize Orchestrator | 182](#)
- [Managing Networks From Contrail Command and VMware vCenter User Interfaces | 187](#)
- [Configuring Underlay Network for ContrailVM | 198](#)

Integrating vCenter for Contrail

IN THIS SECTION

- [Prerequisites | 165](#)
- [ESX Agent Manager | 165](#)
- [Set Up vCenter Server | 165](#)
- [Configure Contrail Parameters | 170](#)
- [Install Contrail | 170](#)
- [Monitor and Manage ContrailVM from ESX Agent Manager | 170](#)

These topics provide instructions for integrating Contrail Release 5.1.x and microservices with VMware vCenter.

Prerequisites

Before you start the integration, ensure that the contrail controller meets the prerequisites given in “[Server Requirements and Supported Platforms](#)” on page 16.

Follow these steps to prepare Contrail controller(s):

```
yum update -y

yum install -y yum-plugin-priorities
https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm

yum install -y python-pip git gcc python-devel sshpass

yum install -y git

pip install "ansible==2.5.0" pyvmomi
```

ESX Agent Manager

VMware provides a standard vCenter solution called vSphere ESX Agent Manager (EAM), that allows you to deploy, monitor, and manage ContrailVMs on ESXi hosts.

The ContrailVM is deployed as an Agent VM that is monitored by EAM. With this integration, ContrailVMs are marked as more critical and privileged than other tenant VMs on the host.

The following are the benefits of running ContrailVM as an AgentVM from EAM:

- Auto-deploy ContrailVMs on ESXi hosts in scope (clusters).
- Manage and Monitor ContrailVMs through EAM in the vSphere web client.
- Integrate with other vCenter features like AddHos, Maintenance Mode, vSphere DRS, vSphere DPM, and VMWare HA.

These topics provide instructions for integrating Contrail Release 5.1.x and microservices with VMware vCenter.

Set Up vCenter Server

Follow these steps to set up the vCenter server.

1. Download the Contrail Ansible Deployer (**contrail-ansible-deployer-*< >*.tgz**) onto your provisioning host. You can download the deployer from <https://www.juniper.net/support/downloads/?p=contrail#sw>.
2. Untar the **tgz**.

```
- tar xvf contrail-ansible-deployer-< >.tgz
```

3. Prepare a **vcenter_vars.yml** file populated with vCenter server and ESXi hosts parameters. You can download the CentOS 7.5 and ESXi VM Host from <https://www.juniper.net/support/downloads/?p=contrail#sw>.

NOTE: You can see a sample of the **vcenter_vars.yml** file in the **contrail-ansible-deployer/playbooks/roles/vcenter/vars/vcenter_vars.yml** after you extract the image files.

NOTE: The ContrailVM's Open Virtualization Format (OVF) image must be hosted on an http or https server which runs on and is reachable from the vCenter server. The location of the OVF is provided as a URL path for **vmdk**: as shown in the example given below.

```
vcenter_servers:
- SRV1:
  hostname:
  username:
  password:
  # Optional: defaults to False
  #validate_certs: False
  datacentername:
  clusternames:
  #path to the ovf, is needed for ESX Agent Manager to deploy
  ContrailVMs
  vmdk: http://<ip-address>/centos-7.5/LATEST/ContrailVM.ovf
  # Optional: If not specified HA and DRS are turned off on the
  clusters.
  enable_ha: yes
  enable_drs: yes
```

For definition examples, refer **contrail-ansible-deployer/playbooks/roles/vcenter/vars/vcenter_vars.yml.sample**.

To enable HA and DRS in the cluster, set **enable_ha** and **enable_drs** to **yes** in the **vcenter_vars.yml** file. If these flags are not enabled, HA and DRS is turned off by default for newly created and existing clusters.

```

provider_config:
  bms:
    ssh_pwd: password
    ssh_user: root
    ntpserver: 8.8.8.8
    domainsuffix: blah.net

instances:
  bms1:
    provider: bms
    ip: <ip-address>
    roles:
      config_database:
        config:
      control:
      analytics_database:
      analytics:
      webui:
      vcenter_plugin:
  bms2:
    provider: bms
    esxi_host: <ip-address>
    ip: <ip-address>
    roles:
      vrouter:
      vcenter_manager:
        ESXI_USERNAME: root
        ESXI_PASSWORD: password
  bms3:
    provider: bms
    esxi_host: <ip-address>
    ip: <ip-address>
    roles:
      vrouter:
      vcenter_manager:
        ESXI_USERNAME: root
        ESXI_PASSWORD: password
  bms4:
    provider: bms
    esxi_host: <ip-address>
    ip: <ip-address>
    roles:
      vrouter:
      vcenter_manager:

```

```

    ESXI_USERNAME: root
    ESXI_PASSWORD: password

global_configuration:
    CONTAINER_REGISTRY: hub.juniper.net
    CONTAINER_REGISTRY_USERNAME: username
    CONTAINER_REGISTRY_PASSWORD: password
    REGISTRY_PRIVATE_INSECURE: False

contrail_configuration:
    CLOUD_ORCHESTRATOR: vcenter
    CONTROLLER_NODES: <ip-address>
    CONTRAIL_VERSION: 5.1.0-0.360
    RABBITMQ_NODE_PORT: 5673
    VCENTER_SERVER: <ip-address>
    VCENTER_USERNAME: administrator@vsphere.net
    VCENTER_PASSWORD: password
    VCENTER_DATACENTER: <DC name here>
    VCENTER_DVSWITCH: overlay
    VCENTER_WSDL_PATH: /usr/src/contrail/contrail-web-core/webroot/js/vim.wsd
    VCENTER_AUTH_PROTOCOL: https

```

NOTE: The default login credentials for Contrail OVF:

- Username: *root*
- Password: *cOntrail123*

We suggest using unique usernames and passwords in accordance with your organization's security guidelines.

```

---
vcenter_servers:
  - SRV1:
      hostname: <host-ip-address>
      username: administrator@vsphere.net
      password: password
      # Optional: defaults to False
      #validate_certs: False
      datacentername: "<your DC name here>"

```



```

    clusternames:
      - "<your cluster name here>"
    vmdk: http://<ip-address>/contrail/images/ContrailVM.ovf
    dv_switch:
      dv_switch_name: overlay
    dv_port_group:
      dv_portgroup_name: VM_pg
      number_of_ports: 1800

esxihosts:
  - name: <ip-address>
    username: root
    password: password
    datastore: <your local datastore here>
    datacenter: "<your DC name here>"
    cluster: "<your cluster name here>"
    contrail_vm:
      networks:
        - mac: 00:77:56:aa:bb:01
    vcenter_server: SRV1 #leave this
  - name: <ip-address>
    username: root
    password: password
    datastore: <your local datastore here>
    datacenter: "<your DC name here>"
    cluster: "<your cluster name here>"
    contrail_vm:
      networks:
        - mac: 00:77:56:aa:bb:02
    vcenter_server: SRV1 #leave this
  - name: <ip-address>
    username: root
    password: password
    datastore: <your local datastore here>
    datacenter: "<your DC name here>"
    cluster: "<your cluster name here>"
    contrail_vm:
      networks:
        - mac: 00:77:56:aa:bb:77
    vcenter_server: SRV1 #leave this

```

4. Run the Contrail vCenter playbook.

```
ansible-playbook playbooks/vcenter.yml
```

NOTE: Verify that the hostnames for the contrail controller(s) and the ContrailVMs (vRouters) are unique in `/etc/hostname` file.

You can verify hostname from either the DHCP options (if the management network uses DHCP) or manually (if the management network uses static IP allocation).

Configure Contrail Parameters

Populate the file `config/instances.yaml` with Contrail roles.

For an example file, see `contrail-ansible-deployer/confing/instances.yaml.vcenter_example`.

Install Contrail

Install Contrail by running the following Contrail playbooks:

```
ansible-playbook -i inventory/ -e orchestrator=vcenter
playbooks/configure_instances.yml
```

```
ansible-playbook -i inventory/ -e orchestrator=vcenter
playbooks/install_contrail.yml
```

Monitor and Manage ContrailVM from ESX Agent Manager

ContrailVMs can be monitored from EAM by using ContrailVM-Agency.

Follow these steps to monitor and manage Contrail VM from EAM:

1. Resolve issues from the ContrailVM-Agency.

The ContrailVM-Agency is in an alert state when the ContrailVM in any host is powered off or is deleted.

Click **Resolve All Issues** from the ContrailVM-Agency to correct the issue. The ContrailVM-Agency will attempt to correct the issue by bringing the ContrailVM back online or by spawning a ContrailVM from the OVF on the ESXi host.

Figure 19: vCenter Server Extensions






Figure 20: ESX Agencies



- 2. Add host.
 - a. Add ESXi host to the cluster.
 - b. Configure **Agent VM Settings** for the ESXI host.

Figure 21: Configure Agent VM Settings



For more information on configuring Agent VM, network, and datastore settings, see [Configure Agent VM Settings](#).

EAM deploys a ContrailVM (from the base OVF) on the ESXi host.

- c. Add ESXi host details to **vcenter_vars.yml** and repeat step 4 to add appropriate interfaces to the ContrailVM and to configure necessary settings in the vCenter server.
 - d. Add ContrailVM details to **instances.yml** and provision Contrail on the newly added ContrailVm (router). For more information on provisioning Contrail, see [“Install Contrail” on page 170](#).
3. Clean up the ContrailVM-Agency.
- Delete **ContrailVM-Agency** from the EAM user interface to delete ContrailVM and the agency.

RELATED DOCUMENTATION

Configuring Underlay Network for ContrailVM 198
Managing Networks From Contrail Command and VMware vCenter User Interfaces 187

Installing and Provisioning Contrail VMware vRealize Orchestrator Plugin

IN THIS SECTION

- [Accessing vRO Control Center | 173](#)
- [Installing vRO Plugin | 175](#)
- [Accessing vRO Desktop Client | 178](#)

- Connecting to vRO using the Desktop Client | 178
- Connecting to Contrail Controller | 179
- Deploying Contrail vRO Plugin | 182

A dedicated Contrail plugin is used to connect to VMware vRealize Orchestrator (vRO). Contrail Release 5.0 supported a Beta version of the plugin. Starting with Contrail Release 5.1, a fully supported version of the plugin is available.

You must install the Contrail VMware vRealize Orchestrator (vRO) plugin to connect to the vRO server.

Before you begin installation, ensure the following:

- You have administrator-level access to the Control Center of a deployed vRO appliance.
- You know the host name ({vRO}) of the deployed vRO Appliance.
- You have the login credentials of the vCenter SSO service.
- You have downloaded the vRO plugin package file to your local system.

You can download the plugin from <https://www.juniper.net/support/downloads/?p=contrail>.

You can deploy the Contrail plugin in any Java Virtual Machine (JVM) compatible environment and load it on an active vRO instance.

The following topics describe how to install and provision the Contrail vRO plugin.

Accessing vRO Control Center

Follow the steps given below to access and log in to vRO Control Center:


1. To access vRO Control Center through a Web browser, navigate to the <https://{vRO}:8283/vco-controlcenter> URL.

NOTE: Replace {vRO} given in the URL with the *host name* of the deployed vRO Appliance.

The *host name* is the IP address or the FQDN of the vRO node.

The **vCenter SSO** service page is displayed.

Figure 22: vCenter SSO service page

The image shows the vCenter SSO service page login interface. It has a blue background with the VMware logo in the top left corner. The login form is centered and includes a 'User name:' label followed by a text input field containing 'example@domain.local'. Below this is a 'Password:' label followed by a password input field. Under the password field is a checkbox labeled 'Use Windows session authentication'. At the bottom right of the form is a 'Login' button.

vmware®

User name: example@domain.local

Password:


☐ Use Windows session authentication

Login

2. On the vCenter SSO service page, enter the **User name** and **Password** in the respective fields and click **Login**. See [Figure 22 on page 174](#).

The **Orchestrator Control Center** home page is displayed.

Figure 23: Orchestrator Control Center



Installing vRO Plugin

Perform the following steps to install the vRO plugin:


1. Upload vRO plugin package.

To upload vRO plugin package:

- From the Orchestrator Control Center home page, click **Manage Plug-Ins** under the **Plug-Ins** section.

The **Manage Plug-Ins** page is displayed.

Figure 24: Manage Plug-Ins page



Manage Plug-Ins

 Install a new plug-in or manage already installed plug-ins. The preferred plug-in installation file format is .VMOAPP, but plug-ins can also be installed as .DAR files. When 'DEFAULT' logging level is selected for a specific plug-in the log level is inherited from the log level set in [Configure Logs](#) page.

Install plug-in

Plugin file (*.dar or *.vmoapp) **BROWSE** **INSTALL**


NOTE: You can install a new plugin or manage an already installed plugin from the Manage Plug-Ins page.

NOTE: *.vmoapp or *.dar file format can be used. Also, the version in this example may be different from the version you have downloaded.

- Click **Browse** in the **Install plug-in** pane and select the downloaded vRO plugin package file on your local system.
- After you select vRO plugin package file, click **Install** to upload the vRO plugin package to the vRO server.

The **EULA** page is displayed.

Figure 25: EULA page



Contrail 1.0.192
Contrail plug-in for vRealize Orchestrator

EULA:

PUT YOUR LICENSE HERE

Accept EULA ☒

CANCEL INSTALL

2. Install vRO plugin.

After you upload the vRO plugin package, select **Accept EULA** on the **EULA** page and then click **Install**.

NOTE: If you use *.vmoapp file format, you are directed to the Accept EULA page before you proceed with the installation.

If you use *.dar file format, you can directly proceed with installation.

The vRO plugin is installed.

Accessing vRO Desktop Client

After you install the VMware vRealize Orchestrator (vRO) plugin, download vRealize Orchestrator Client version 7.3.0 to access the vRO server.

To download and install the vRO desktop client application, click <https://{vRO}:8281/vco/>.

NOTE: Replace {vRO} given in the URL with the *host name* of the deployed vRO Appliance.

Figure 26: Getting Started with vRealize Orchestrator



You can download vRO desktop client applications for Windows, Mac OS X, and Linux operating systems.


Connecting to vRO using the Desktop Client

You connect to the vRO server by using the vRO desktop client.

1. Start the vRO desktop client.

The **VMware vRealize Orchestrator Login** page is displayed.

Figure 27: VMware vRealize Orchestrator Login page



2. In the VMware vRealize Orchestrator Login page, enter **Host name**, **User name**, and **Password**.

NOTE: The **Host name** also includes the port number and must be in the {vRO}:8281 format.


3. Click **Login** to connect to the vRO server. See [Figure 27 on page 179](#).

Connecting to Contrail Controller

To connect Contrail vRO to the Contrail Controller:

1. Navigate to the **Contrail > Configuration** folder in the workflow library. See [Figure 28 on page 180](#).
2. Select **Create Contrail controller connection**.

Figure 28: Workflow Library



3. Click the **Controller** tab and enter the following information:
 - **Connection name**—a unique name to identify the connection
 - **Controller host**—host name of the Contrail Connector
 - **Controller port**—port used to access the Contrail Controller

Figure 29: Controller Tab

1 Controller
2 Credentials
3 Tenant

* Connection name
Controller

* Controller host

* Controller port
8082

4. Click the **Credentials** tab and enter the following credentials to manage the Contrail Controller:

- **User name**—user name to access the Contrail Controller
- **User password**—password to access the Contrail Controller
- **Authentication server**—URL of the authentication server

Figure 30: Credentials Tab

✓ **1 Controller**
✓ **2 Credentials**
3 Tenant

User name


User password

Authentication server

5. Click the **Tenant** tab to define tenant information.

In the **Tenant** field, enter the name of the Contrail tenant.

Figure 31: Tenant Tab



6. Click **Submit** to establish connection.

Once you connect Contrail vRO to the Contrail Controller, you use Contrail workflows to make configuration changes to Contrail.

Deploying Contrail vRO Plugin

You can deploy the Contrail plugin in any Java Virtual Machine (JVM) compatible environment and load it on an active vRO instance.

SEE ALSO

| [Integrating Contrail Release 5.0.X with VMware vRealize Orchestrator | 182](#)

Integrating Contrail Release 5.0.X with VMware vRealize Orchestrator

IN THIS SECTION

- [Components of vRealize Orchestrator | 183](#)
- [Contrail Workflows | 183](#)

A dedicated Contrail plugin is used to connect to VMware vRealize Orchestrator (vRO). Contrail Release 5.0 supported a Beta version of the plugin. Starting with Contrail Release 5.0.1, a fully supported version of the plugin is available. The plugin is used to view the Contrail controller configuration in the vRO inventory, and to modify configurations by using vRO workflows. You can also create network policies, security groups, and automate both simple and complex workflows by using vRO.

Components of vRealize Orchestrator

vRO consists of the following components:

- **vRO Inventory**—The vRO inventory displays the Contrail plugin and the Contrail node or endpoint. All Contrail plugin objects that represent your system are displayed in the vRO Inventory. Objects are displayed in a hierarchical order and are based on the Contrail schema.

With Release 5.0, Contrail inventory objects such as projects, security groups, virtual networks, network IPAMs, network policies, ports, floating IP pools, and service templates are displayed in the vRO inventory. Relevant API objects are also displayed in the vRO inventory.

- **vRO Workflows**—The vRO workflow is a process that manipulates objects in a vRO Inventory. Each plugin has a set of predefined workflows. vRO combines workflows from different plugins to create complex processes and manages them. Multiple workflows are used to create blueprints in vRA.

NOTE: VMware vCenter plugin workflows are represented as simple workflows in vRO plugin.

Contrail Workflows

You must connect to the Contrail controller or an endpoint before you create Contrail workflows. You must use **Create Contrail controller connection** to connect to an endpoint. You must use **Delete Contrail controller connection** to terminate a connection with an endpoint. Once you connect to the Contrail controller, the vRO plugin accesses the Contrail inventory objects and then updates the vRO inventory with the Contrail inventory objects.

The following workflows are defined for simple networking operations in Contrail Release 5.0:

- **Network**
 - Create network
 - Delete network

NOTE: You must use the **Create network** workflow to create a network on the Contrail controller.

- **Network Policy**
 - Create network policy
 - Add rule to network policy
 - Remove network policy rule
 - Delete network policy

- Security policy
 - Create security group
 - Add rule to security group
 - Edit security group
 - Remove security group rule
 - Delete security group
- Service Instance
 - Add port tuple to service instance
 - Create service instance
 - Delete service instance
 - Remove port tuple from service instance
- Network IPAM
- Port
- Project
- Service Template
- Virtual Network
- Floating IP
 - Create floating IP
 - Delete floating IP
- Floating IP pools
 - Create floating IP pool
 - Delete floating IP pool
 - Edit floating IP pool

Starting with Contrail Release 5.0.1, the following workflows are also defined:

- Tag
 - Create global tag
 - Create tag in project
 - Delete tag
- Tag Type
 - Create tag type

- Delete tag type
- Network Policy
 - Edit network policy rule
- Security policy
 - Edit security group rule
- Service Health Check
 - Create service health check
 - Add service instance to service health check
 - Remove service instance from service health check
 - Edit service health check
 - Delete service health check
- Project
 - Add application policy set to project
 - Remove application policy set to project
 - Add tag to project
 - Remove tag from project
- Virtual Network
 - Add tag to virtual network
 - Remove tag from virtual network
- Virtual Machine Interface (VMI) - Port
 - Add tag to port
 - Remove tag from port
- Service Group
 - Create service group in policy management
 - Create service group in project
 - Add service to service group
 - Edit service of service group
 - Remove service from service group
 - Delete service group
- Address Group

- Create global address group
- Create address group in project
- Add subnet to address group
- Remove subnet from address group
- Delete address group
- Add label to address group
- Remove label from address group
- Application Policy Set
 - Create global application policy set
 - Create application policy set in project
 - Add firewall policy to application policy set
 - Remove firewall policy from application policy set
 - Add tag to application policy set
 - Remove tag from application policy set
 - Delete application policy set
- Firewall Policy
 - Create global firewall policy
 - Create firewall policy in project
 - Add firewall rule to firewall policy
 - Remove firewall rule from firewall policy
 - Delete firewall policy
- Firewall Rule
 - Create global firewall rule
 - Create firewall rule in project
 - Edit firewall rule
 - Delete firewall rule

Managing Networks From Contrail Command and VMware vCenter User Interfaces

IN THIS SECTION

- [User Interfaces for Contrail Integration with VMware vCenter | 187](#)
- [Feature Configuration for Contrail vCenter | 188](#)

You can install Contrail to work with the VMware vCenter Server version 6.5. You use the Contrail Command user interface (UI) and the vCenter UI to configure and manage the integrated Contrail system.

These topics provide instructions on how to use the Contrail Command UI and the vCenter UI to configure and manage features of a Contrail VMware integrated system.

User Interfaces for Contrail Integration with VMware vCenter

The Contrail Command and VMware vCenter UIs are available after installing the integrated Contrail system.

After Contrail is integrated with VMware vCenter, you can use the UIs to manage and configure features of the system.

Contrail Administration User Interface

Contrail Command is an administrator's user interface. It provides a view of all components managed by the Contrail controller.


1. To log in to Contrail Command, use your Contrail server main IP address URL as follows:

`https://<contrail command ip>:9091`

2. Then log in using your registered Contrail account administrator credentials.

Upon successful login, the Contrail vCenter UI is displayed. See [Figure 32 on page 188](#).

Figure 32: Contrail vCenter User Interface



Feature Configuration for Contrail vCenter

IN THIS SECTION

- Create Virtual Network | 189
- Delete Virtual Network | 191
- Create a Virtual Machine – vCenter UI | 193
- View Virtual Machine | 198

This section provides instructions on how to create, delete, and view virtual networks by using the Contrail Command UI. This section also provides instructions on how to create a virtual machine by using the vCenter client.

Create Virtual Network

Follow these steps to create a virtual network by using the Contrail Command UI.

1. Click **Overlay > Virtual Networks** to access the All Networks page.

NOTE: Ensure that you are in the vCenter project.

2. Click **Create**.

The Network tab of the Create Virtual Network page is displayed.

3. Enter the name of the virtual network in the **Name** field.
4. Select a network policy from the **Network Policies** list.
5. Select an allocation mode from the **Allocation Mode** list.
6. Click **+ Add** in the Subnets section and add the following information in the fields that are displayed.
 - a. Select a network IPAM service from the **Network IPAM** list.
 - b. Enter valid IPv4 subnet/mask in the **CIDR** field.
 - c. Enter allocation pool information in the **Allocation Pools** field.
 - d. Enter service IP address in the **Service Address** field.
 - e. Click **Create** to add subnet information.

[Figure 33 on page 190](#) shows the creation of a virtual network called Green-VN.

Figure 33: Create a Virtual Network

OVERLAYVirtual NetworksCreate Virtual Networkdefault-domainvCenterAdmin

NetworkTagsPermissions

Name*
Green-VN

Network Policies
Select Network Policies

Allocation Mode
User defined subnet only

VxLAN Network Identifier
1 - 16777215

Subnets

Network IPAM*
vCenter-ipam

CIDR*
192.0.2.0/24

Allocation Pools
xxx.xxx.xxx-xxx.xxx.xxx

Gateway*
192.0.2.1

Service Address
xxx.xxx.xxx


☒ Auto Gateway☒ DHCP☒ DNS

CreateCancel

- 7. Click **Create** to create the virtual network.

The All Networks page is displayed as given in [Figure 34 on page 191](#).

Figure 34: All Networks Page



Delete Virtual Network


You can delete a virtual network from the Contrail Command UI.

Follow these steps to delete a virtual network.

1. Click **Overlay > Virtual Networks** to access the All Networks page.
2. Select the check box next to the virtual network you want to delete and move the mouse pointer to the end of the row.

Click the **Remove** icon to delete the virtual network. See [Figure 35 on page 192](#).

Figure 35: Delete a Virtual Network




NOTE: No virtual machines or templates must be using this virtual network before it is deleted.

- A delete confirmation message is displayed.
3. Click **Delete** to delete the selected network.
- The virtual network is deleted.


Create a Virtual Machine – vCenter UI

Use the vCenter client interface to create a virtual machine for your VMware vCenter Contrail integrated system. This section describes how to create a virtual machine using a virtual machine template from the vCenter client interface.

1. From the vCenter UI, right-click on the cluster and select **Deploy OVF Template**.



2. Select the virtual machine name and click **Next**.



3. Select the compute resource on which you want to spawn the virtual machine and click **Next**.

- Specify the storage location for the virtual machine and click **Next**.

OL-cluster ACTIONS

Summary Monitor

Total Pro Total vM

Related Objects

Datacenter

vSphere HA

Tags

Assigned Tag

Assign... Remove...

Cluster Resources

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- 5 Select storage**
- 6 Select networks
- 7 Ready to complete

Select storage
Select the storage for the configuration and disk files

☐ Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format: Thin Provision

VM Storage Policy: Datastore Default


Name	Capacity	Provisioned	Free	Type
datastore-1	931.25 GB	33.53 GB	911.82 GB	VM
datastore-3	931.25 GB	1.42 GB	929.83 GB	VM
datastore-4	931.25 GB	1.42 GB	929.83 GB	VM
datastore-ssd1	216 GB	1.41 GB	214.59 GB	VM
VMotion	270 GB	67.88 GB	202.12 GB	NF

Compatibility

✓ Compatibility checks succeeded.

CANCEL BACK **NEXT**


5. Select the **Destination Network**, Green-VN, that you created by using the Contrail Command UI.



6. On the **Ready to complete** window, review all of the virtual machine definitions that you have selected for the template.

If all the selections are correct, click **Finish**.

This spawns the virtual machine.



7. You return to the main vCenter UI window. Navigate to the console of the VM (yVM) you just spawned to see that it has been assigned an IP address from the Green-VN you created in the Contrail UI.

```

yvm
(■-
//\  Core is distributed with ABSOLUTELY NO WARRANTY.
v_/_  www.tinycorelinux.com

tc@box:~$ ifconfig
eth0    Link encap:Ethernet  HWaddr XX:XX:XX:XX:XX:XX
        inet addr:192.168.1.3  Bcast:XXX.XXX.X.XXX  Mask:XXX.XXX.255.0
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:14 errors:0 dropped:0 overruns:0 frame:0
        TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:1374 (1.3 KiB)  TX bytes:1044 (1.0 KiB)

lo      Link encap:Local Loopback
        inet addr:XXX.0.0.1  Mask:XXX.0.0.0
        UP LOOPBACK RUNNING  MTU:65536  Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:0 (0.0 B)  TX bytes:0 (0.0 B)

tc@box:~$ _

```

Once the virtual machine is launched, you can view the virtual machine from the Contrail Command UI.

View Virtual Machine

You can view the virtual machines that you created by using the vCenter client, from the Contrail Command UI.

To view the virtual machines, click **Workloads > Instances**.

All the virtual machines are displayed in the Instances page.

RELATED DOCUMENTATION

[Configuring Underlay Network for ContrailVM | 198](#)

Configuring Underlay Network for ContrailVM

IN THIS SECTION


- [Standard Switch Setup | 199](#)
- [Distributed Switch Setup | 200](#)
- [PCI Pass-Through Setup | 202](#)
- [SR-IOV Setup | 205](#)

The ContrailVM can be configured in several different ways for the underlay (**ip-fabric**) connectivity:

Standard Switch Setup

In the standard switch setup, the ContrailVM is provided an interface through the standard switch port group that is used for management and control data, see [Figure 36 on page 199](#).

Figure 36: Standard Switch Setup



To set up the ContrailVM in this mode, the standard switch and port group must be configured in **vcenter_vars.yml**.

If switch name is not configured, the default values of **vSwitch0** are used for the standard switch.

The ContrailVM supports multiple NICs for management and **control_data** interfaces. The management interface must have the DHCP flag as **true** and the **control_data** interface can have DHCP set as **false**. When DHCP is set to false, the IP address of the **control_data** interface must be configured by the user and ensure connectivity. Additional configuration such as static routes and bond interface must be configured by the user.

The following is an example of configuration with standard switch.

```

-   name: <esxi_host>
      username: <username>
      password: <password>
      datastore: <datastore>
      vcenter_server: <server>
      datacenter: <datacenter>
      cluster: <cluster>
      std_switch_list:
        - pg_name: mgmt-pg
          switch_name: vSwitch0
      contrail_vm:
        networks:
          - mac: 00:77:56:aa:bb:03
            sw_type: standard
            switch_name: vSwitch0
            pg: mgmt-pg

```


Distributed Switch Setup

A distributed switch functions as a single virtual switch across associated hosts.

In the distributed switch setup, the ContrailVM is provided an interface through the distributed switch port group that is used for management and control data, see [Figure 37 on page 201](#).

The ContrailVM can be configured to use the management and control_data NICs from DVS. When the DVS configuration is specified, the standard switch configuration is ignored.

Figure 37: Distributed Switch Setup



To set up the ContrailVM in this mode, configure the distributed switch, port group, number of ports in the port group, and the uplink in the `vcenter_servers` section in `vcenter_servers.yml`.

NOTE: The uplink can be a link aggregation group (LAG). If you use LAG, then DVS and LAG should be preconfigured.

The following is an example distributed switch configuration in `vcenter_vars.yml`.

```
vcenter_servers:
- SRV1:
  hostname: <server>
  username: <username>
  password: <password>
  datacentername: <datacenter>
  clusternames:
    - <cluster>
```

```

dv_switch:
  dv_switch_name: <dvs_name>
dv_port_group:
  dv_portgroup_name: <pg_name>
  number_of_ports: <num_of_ports>
dv_switch_control_data:
  dv_switch_name: <ctrl_dvs_name>
dv_port_group_control_data:
  dv_portgroup_name: <ctrl_pg_name>
  number_of_ports: <num_of_ports>
uplink:
  - 'vmnic3'


```

PCI Pass-Through Setup

PCI pass-through is a virtualization technique in which a physical Peripheral Component Interconnect (PCI) device is directly connected to a virtual machine, bypassing the hypervisor. Drivers in the VM can directly access the PCI device, resulting in a high rate of data transfer.

In the pass-through setup, the ContrailVM is provided management and control data interfaces. Pass-through interfaces are used for control data. [Figure 38 on page 203](#) shows a PCI pass-through setup with a single **control_data** interface.

Figure 38: PCI Pass-Through with Single Control Data Interface




When setting up the ContrailVM with pass-through interfaces, upon provisioning ESXi hosts in the installation process, the PCI pass-through interfaces are exposed as Ethernet interfaces in the ContrailVM, and are identified in the **control_data** device field.

The following is an example PCI pass-through configuration with a single **control_data** interface:

```
esxihosts:
  - name: <esxi_host>
    username: <username>
    password: <password>
    datastore: <datastore>
    vcenter_server: <server>
    datacenter: <datacenter>
    cluster: <cluster>
    contrail_vm:
      networks:
        - mac: <mac_addr>
      pci_devices:
        - '0000:04:00.0'
```

Figure 39 on page 204 shows a PCI pass-through setup with a bond_control data interface, which has multiple pass-through NICs.

Figure 39: PCI Pass-Through Setup with Bond Control Interface



Update the ContrailVM section in `vcenter_vars.yml` with `pci_devices` as shown in the following example:

```
esxihosts:
  - name: <esxi_host>
    username: <username>
    password: <password>
    datastore: <datastore>
    vcenter_server: <server>
    datacenter: <datacenter>
    cluster: <cluster>
    contrail_vm:

      networks:
        - mac: <mac_addr>
      pci_devices:
```


- '0000:04:00.0'
- '0000:04:00.1'

SR-IOV Setup

A single root I/O virtualization (SR-IOV) interface allows a network adapter device to separate access to its resources among various hardware functions.

In the SR-IOV setup, the ContrailVM is provided management and control data interfaces. SR-IOV interfaces are used for control data. See [Figure 40 on page 205](#).

Figure 40: SR-IOV Setup



In VMware, the **port-group** is mandatory for SR-IOV interfaces because the ability to configure the networks is based on the active policies for the port holding the virtual machines.


To set up the ContrailVM with SR-IOV interfaces, all configurations used for the standard switch setup are also used for the pass-through setup, providing management connectivity to the ContrailVM.

To provide the **control_data** interfaces, configure the SR-IOV-enabled physical interfaces in the **contrail_vm** section, and configure the **control_data** in the global section of **vcenter_vars.yml**.

Upon provisioning ESXi hosts in the installation process, the SR-IOV interfaces are exposed as Ethernet interfaces in the Contrail VM.

Figure 41 on page 206 shows a SR-IOV setup with a single **control_data** interface.

Figure 41: SR-IOV With Single Control Data Interface



The following is an example SR-IOV configuration for the cluster and server configuration.

The cluster configuration:

```
vcenter_servers:
- SRV1:
  hostname: <server>
  username: <username>
  password: <password>
  datacentername: <datacenter>
  clusternames:
  - <cluster>
```

```

dv_switch:
  dv_switch_name: <dvs_name>
dv_port_group:
  dv_portgroup_name: <pg_name>
  number_of_ports: <num_of_ports>
dv_switch_sr_iov:
  dv_switch_name: <sriov_dvs_name>
dv_port_group_sriov:
  dv_portgroup_name: <sriov_pg_name>
  number_of_ports:

```

The server configuration:

```

esxihosts:
  - name: <esxi_host>
    username: <username>
    password: <password>
    datastore: <datastore>
    vcenter_server: <server>
    datacenter: <datacenter>
    cluster: <cluster>
    contrail_vm:

    networks:
      - mac: <mac_addr>
    sr_iov_nics:
      - 'vmnic0'

```

[Figure 42 on page 208](#) shows an SR-IOV configuration with a bond **control_data** interface, which has multiple SR-IOV NICs.


```

    number_of_ports: <num_of_ports>
  dv_switch_sr_iov:
    dv_switch_name: <sriov_dvs_name>
  dv_port_group_sriov:
    dv_portgroup_name: <sriov_pg_name>
  number_of_ports:

```

The server configuration:

```

esxihosts:
  - name: <esxi_host>
    username: <username>
    password: <password>
    datastore: <datastore>
    vcenter_server: <server>
    datacenter: <datacenter>
    cluster: <cluster>
    contrail_vm:

    networks:
      - mac: <mac_addr>
    sr_iov_nics:
      - 'vmnic0'
      - 'vmnic1'

```

SEE ALSO

| [Managing Networks From Contrail Command and VMware vCenter User Interfaces](#) | 187

Using Contrail Networking with OpenStack

IN THIS CHAPTER

- [Setting Up Contrail with Red Hat OpenStack 16.1 | 210](#)
- [Setting Up Contrail with Red Hat OpenStack 13 | 270](#)
- [Configuring Virtual Networks | 339](#)
- [Using Contrail Resources in Heat Templates | 353](#)
- [QoS Support in Contrail Networking | 358](#)
- [Load Balancers | 367](#)
- [Optimizing Contrail Networking | 399](#)
- [Contrail Networking OpenStack Analytics | 400](#)
- [Contrail OpenStack APIs | 407](#)

Setting Up Contrail with Red Hat OpenStack 16.1

IN THIS SECTION

- [Understanding Red Hat OpenStack Platform Director | 211](#)
- [Setting Up the Infrastructure | 216](#)
- [Setting Up the Undercloud | 225](#)
- [Setting Up the Overcloud | 227](#)

Understanding Red Hat OpenStack Platform Director

IN THIS SECTION

- [Red Hat OpenStack Platform Director | 211](#)
- [Contrail Networking Roles | 212](#)
- [Undercloud Requirements | 213](#)
- [Overcloud Requirements | 213](#)
- [Networking Requirements | 214](#)
- [Compatibility Matrix | 215](#)
- [Installation Summary | 215](#)

Red Hat OpenStack Platform Director

Starting with Contrail Networking Release 2008, Contrail Networking supports using Contrail with Red Hat OpenStack Platform Director 16.1.

This chapter explains how to integrate a Contrail Networking Release 2008 (or higher) installation with Red Hat OpenStack Platform Director 16.1.

Red Hat OpenStack Platform provides an installer called the Red Hat OpenStack Platform director (RHOSPd or OSPd), which is a toolset based on the OpenStack project TripleO (OOO, OpenStack on OpenStack). TripleO is an open source project that uses features of OpenStack to deploy a fully functional, tenant-facing OpenStack environment.

TripleO can be used to deploy an RDO-based OpenStack environment integrated with Tungsten Fabric. Red Hat OpenStack Platform director can be used to deploy an RHOSP-based OpenStack environment integrated with Contrail Networking.

OSPd uses the concepts of undercloud and overcloud. OSPd sets up an undercloud, a single server running an operator-facing deployment that contains the OpenStack components needed to deploy and manage an overcloud, a tenant-facing deployment that hosts user workloads.

The overcloud is the deployed solution that can represent a cloud for any purpose, such as production, staging, test, and so on. The operator can select to deploy to their environment any of the available overcloud roles, such as controller, compute, and the like.

OSPd leverages existing core components of OpenStack including Nova, Ironi, Neutron, Heat, Glance, and Ceilometer to deploy OpenStack on bare metal hardware.

- Nova and IroniC are used in the undercloud to manage the bare metal instances that comprise the infrastructure for the overcloud.
- Neutron is used to provide a networking environment in which to deploy the overcloud.
- Glance stores machine images.
- Ceilometer collects metrics about the overcloud.

For more information about OSPd architecture, see [OSPd documentation](#).

Contrail Networking Roles


OSPd supports composable roles, which are groups of services that you define through Heat templates. Composable roles allow you to integrate Contrail Networking into the overcloud environment.

The following are the Contrail Networking roles used for integrating into the overcloud:

- Contrail Controller
- Contrail Analytics
- Contrail Analytics Database
- Contrail-TSN
- Contrail-DPDK

[Figure 43 on page 213](#) shows the relationship and components of an undercloud and overcloud architecture for Contrail Networking.

Figure 43: Undercloud and Overcloud with Roles



Undercloud Requirements

The undercloud is a single server or VM that hosts the OpenStack Platform director, which is an OpenStack installation used to provision OpenStack on the overcloud.

See [Undercloud Requirements](#) for the compute requirements of the undercloud.

Overcloud Requirements

The overcloud roles can be deployed to bare metal servers or to virtual machines (VMs), but the compute nodes must be deployed to bare metal systems. Every overcloud node must support IPMI for booting up from the undercloud using PXE.

Ensure the following requirements are met for the Contrail Networking nodes per role.

- Non-high availability: A minimum of 4 overcloud nodes are needed for control plane roles for a non-high availability deployment:
 - 1x contrail-config (includes Contrail control)
 - 1x contrail-analytics
 - 1x contrail-analytics-database
 - 1x OpenStack controller

- High availability: A minimum of 12 overcloud nodes are needed for control plane roles for a high availability deployment:
 - 3x contrail-config (includes Contrail control)
 - 3x contrail-analytics
 - 3x contrail-analytics-database
 - 3x OpenStack controller

If the control plane roles are deployed to VMs, use 3 separate physical servers and deploy one role of each kind to each physical server.

See [Overcloud Requirements](#) for the compute requirements of the overcloud.

Networking Requirements

As a minimum, the installation requires two networks:

- provisioning network - This is the private network that the undercloud uses to provision the overcloud.
- external network - This is the externally-routable network you use to access the undercloud and overcloud nodes.

Ensure the following requirements are met for the provisioning network:

- One NIC from every machine must be in the same broadcast domain of the provisioning network, and it should be the same NIC on each of the overcloud machines. For example, if you use the second NIC on the first overcloud machine, you should use the second NIC on each additional overcloud machine. During installation, these NICs will be referenced by a single name across all overcloud machines.
- The provisioning network NIC should not be the same NIC that you are using for remote connectivity to the undercloud machine. During the undercloud installation, an Open vSwitch bridge will be created for Neutron, and the provisioning NIC will be bridged to the Open vSwitch bridge. Consequently, connectivity would be lost if the provisioning NIC was also used for remote connectivity to the undercloud machine.
- The provisioning NIC on the overcloud nodes must be untagged.
- You must have the MAC address of the NIC that will PXE boot the IPMI information for the machine on the provisioning network. The IPMI information will include such things as the IP address of the IPMI NIC and the IPMI username and password.
- All of the networks must be available to all of the Contrail Networking roles and computes.

While the provisioning and external networks are sufficient for basic applications, you should create additional networks in most overcloud environments to provide isolation for the different traffic types by assigning network traffic to specific network interfaces or bonds.

When isolated networks are configured, the OpenStack services are configured to use the isolated networks. If no isolated networks are configured, all services run on the provisioning network. If only some isolated networks are configured, traffic belonging to a network not configured runs on the provisioning network.

The following networks are typically deployed when using network isolation topology:

- Provisioning - used by the undercloud to provision the overcloud
- Internal API - used by OpenStack services to communicate with each other
- Tenant - used for tenant overlay data plane traffic (one network per tenant)
- Storage - used for storage data traffic
- Storage Management - used for storage control and management traffic
- External - provides external access to the undercloud and overcloud, including external access to the web UIs and public APIs
- Floating IP - provides floating IP access to the tenant network (can either be merged with external or can be a separate network)
- Management - provides access for system administration

Compatibility Matrix

The following combinations of Operating System/OpenStack/Deployer/Contrail Networking are supported:

Table 6: Compatibility Matrix

Operating System	OpenStack	Deployer	Contrail Networking
RHEL 8.2	OSP16	OSPd16	Contrail Networking 2008 or higher

Installation Summary

The general installation procedure is as follows:

- Set up the infrastructure, which is the set of servers or VMs that host the undercloud and overcloud, including the provisioning network that connects them together.
- Set up the undercloud, which is the OSPd application.
- Set up the overcloud, which is the set of services in the tenant-facing network. Contrail Networking is part of the overcloud.

For more information on installing and using the RHOSPd, see [Red Hat documentation](#).

Release History Table

Release	Description
2008	Starting with Contrail Networking Release 2008, Contrail Networking supports using Contrail with Red Hat OpenStack Platform Director 16.1.

Setting Up the Infrastructure

SUMMARY

Follow this topic to set up the infrastructure for Contrail Networking deployment with RHOSP 16.1.

IN THIS SECTION

- [Target Configuration \(Example\) | 216](#)
- [Configure the External Physical Switch | 218](#)
- [Configure KVM Hosts | 219](#)
- [Create the Overcloud VM Definitions on the Overcloud KVM Hosts | 221](#)
- [Create the Undercloud VM Definition on the Undercloud KVM Host | 223](#)

Target Configuration (Example)

Undercloud and overcloud KVM hosts require virtual switches and virtual machine definitions to be configured. You can deploy any KVM host operating system version that supports KVM and OVS. The following example shows a RHEL/CentOS based system. If you are using RHEL, you must subscribe the system.

The following example illustrates all control plane functions as Virtual Machines hosted on KVM hosts.

There are different ways to create the infrastructure providing the control plane elements. To illustrate the installation procedure, we will use four host machines for the infrastructure, each running KVM. KVM1 contains a VM running the undercloud while KVM2 through KVM4 each contains a VM running an OpenStack controller and a Contrail controller ([Table 7 on page 216](#)).

Table 7: Control Plane Infrastructure

KVM Host	Virtual Machines
KVM1	undercloud
KVM2	OpenStack Controller 1, Contrail Contoller 1

Table 7: Control Plane Infrastructure (continued)

KVM Host	Virtual Machines
KVM3	OpenStack Controller 2, Contrail Controller 2
KVM4	OpenStack Controller 3, Contrail Controller 3

Figure 44 on page 217 shows the physical connectivity where each KVM host and each compute node has two interfaces that connect to an external switch. These interfaces attach to separate virtual bridges within the VM, allowing for two physically separate networks (external and provisioning networks).

Figure 44: Physical View





Figure 45 on page 218 shows the logical view of the connectivity where VLANs are used to provide further network separation for the different OpenStack network types.

Figure 45: Logical View



The following sections describe how to configure the infrastructure, the undercloud, and finally the overcloud.

Configure the External Physical Switch

Configure the ports and VLANs on the external physical switch according to the following table:

Table 8: External Physical Switch Port and VLAN Configuration

Port	Trunked VLAN	Native VLAN
ge0	-	-
ge1	700, 720	-

Table 8: External Physical Switch Port and VLAN Configuration (*continued*)

Port	Trunked VLAN	Native VLAN
ge2	700, 710, 720, 730, 740, 750	-
ge3	-	-
ge4	710, 730	700
ge5	-	-

Configure KVM Hosts

Use this example procedure to install the required packages and start KVM and Open vSwitch on each undercloud and overcloud KVM host.

1. Log in to a KVM host.
2. Install the required packages.

```
yum install -y libguestfs \
  libguestfs-tools \
  openvswitch \
  virt-install \
  kvm libvirt \
  libvirt-python \
  python-virtualbmc \
  python-virtinst
```

3. Start KVM and Open vSwitch.

```
systemctl start libvirtd
systemctl start openvswitch
```

4. Additionally, on the overcloud nodes only, create and start the virtual switches br0 and br1.

Table 9: vSwitch Configuration

Bridge	Trunked VLAN	Native VLAN
br0	710, 720, 730 740, 750	700
br1	-	-

```

# Create the virtual switches and bind them to the respective interfaces.
ovs-vsctl add-br br0
ovs-vsctl add-br br1
ovs-vsctl add-port br0 NIC1
ovs-vsctl add-port br1 NIC2

# Create the configuration file for br0.
cat << EOF > br0.xml
<network>
  <name>br0</name>
  <forward mode='bridge'/>
  <bridge name='br0'/>
  <virtualport type='openvswitch'/>
  <portgroup name='overcloud'/>
    <vlan trunk='yes'>
      <tag id='700' nativeMode='untagged'/>
      <tag id='710'/>
      <tag id='720'/>
      <tag id='730'/>
      <tag id='740'/>
      <tag id='750'/>
    </vlan>
  </portgroup>
</network>
EOF

# Create the configuration file for br1.
cat << EOF > br1.xml
<network>
  <name>br1</name>
  <forward mode='bridge'/>
  <bridge name='br1'/>
  <virtualport type='openvswitch'/>
</network>
EOF

# Create the br0 network based on the configuration file.
virsh net-define br0.xml
virsh net-start br0
virsh net-autostart br0

# Create the br1 network based on the configuration file.
virsh net-define br1.xml
virsh net-start br1

```

```
virsh net-autostart br1
```

5. Repeat step 1 through step 4 for each KVM host.

Create the Overcloud VM Definitions on the Overcloud KVM Hosts

Use this example procedure on each overcloud KVM host (KVM2 to KVM4) to do the following:

- create the VM definitions for that overcloud KVM host
- create and start a virtual baseboard management controller for that overcloud KVM host so that the VM can be managed using IPMI
- create an **ironic_list** file to be used by the undercloud

This example procedure creates a VM definition consisting of 2 compute nodes, 1 Contrail controller node, and 1 OpenStack controller node on each overcloud KVM host.

1. Log in to an overcloud KVM host.
2. Specify the roles you want to create.

```
ROLES=compute:2,contrail-controller:1,control:1
```

3. Create the VM definitions.

```
# Initialize and specify the IPMI user and password you want to use.
num=0
ipmi_user=<user>
ipmi_password=<password>
libvirt_path=/var/lib/libvirt/images
port_group=overcloud
prov_switch=br0
/bin/rm ironic_list

# For each role and instance specified in the ROLES variable:
#   - create the VM definition
#   - create and start a virtual baseboard management controller (vbmc)
#   - store the VM information into an ironic_list file (for later use in the
undercloud)
IFS=' ' read -ra role_list <<< "${ROLES}"
for role in ${role_list[@]}; do
    role_name=`echo $role|cut -d ":" -f 1`
    role_count=`echo $role|cut -d ":" -f 2`
```

```

for count in `seq 1 ${role_count}`; do
    echo $role_name $count
    qemu-img create -f qcow2 ${libvirt_path}/${role_name}_${count}.qcow2 99G
    virsh define /dev/stdin <<EOF
$(virt-install --name ${role_name}_${count} \
    --disk ${libvirt_path}/${role_name}_${count}.qcow2 \
    --vcpus=4 \
    --ram=16348 \
    --network network=br0,model=virtio,portgroup=${port_group} \
    --network network=br1,model=virtio \
    --virt-type kvm \
    --cpu host \
    --import \
    --os-variant rhel8.2 \
    --serial pty \
    --console pty,target_type=virtio \
    --graphics vnc \
    --print-xml)
EOF
    vbmc add ${role_name}_${count} --port 1623${num} --username ${ipmi_user}
--password ${ipmi_password}
    vbmc start ${role_name}_${count}
    prov_mac=`virsh domiflist ${role_name}_${count}|grep ${prov_switch}|awk
'{print $5}'`
    vm_name=${role_name}-${count}-`hostname -s`
    kvm_ip=`ip route get 1 |grep src |awk '{print $7}'`
    echo ${prov_mac} ${vm_name} ${kvm_ip} ${role_name} 1623${num}>> ironic_list
    num=$(expr $num + 1)
done
done

```

4. Repeat step 1 through step 3 on each overcloud KVM host.



CAUTION: This procedure creates one **ironic_list** file per overcloud KVM host. Combine the contents of each file into a single **ironic_list** file on the undercloud.

The following shows the resulting **ironic_list** file after you combine the contents from each separate file:

```
52:54:00:e7:ca:9a compute-1-5b3s31 10.87.64.32 compute 16230
52:54:00:30:6c:3f compute-2-5b3s31 10.87.64.32 compute 16231
52:54:00:9a:0c:d5 contrail-controller-1-5b3s31 10.87.64.32 contrail-controller 16232
52:54:00:cc:93:d4 control-1-5b3s31 10.87.64.32 control 16233
52:54:00:28:10:d4 compute-1-5b3s30 10.87.64.31 compute 16230
52:54:00:7f:36:e7 compute-2-5b3s30 10.87.64.31 compute 16231
52:54:00:32:e5:3e contrail-controller-1-5b3s30 10.87.64.31 contrail-controller 16232
52:54:00:d4:31:aa control-1-5b3s30 10.87.64.31 control 16233
52:54:00:d1:d2:ab compute-1-5b3s32 10.87.64.33 compute 16230
52:54:00:ad:a7:cc compute-2-5b3s32 10.87.64.33 compute 16231
52:54:00:55:56:50 contrail-controller-1-5b3s32 10.87.64.33 contrail-controller 16232
52:54:00:91:51:35 control-1-5b3s32 10.87.64.33 control 16233
```

Create the Undercloud VM Definition on the Undercloud KVM Host

Use this example procedure on the undercloud KVM host (KVM1) to create the undercloud VM definition and to start the undercloud VM.

1. Create the images directory.

```
mkdir ~/images
cd images
```

2. Retrieve the image.

RHEL

Download rhel-server-8.2-update-1-x86_64-kvm.qcow2 from RedHat portal to ~/images.

```
ccloud_image=~/images/rhel-server-8.2-update-1-x86_64-kvm.qcow2
```

3. Customize the undercloud image.

```
undercloud_name=queensa
undercloud_suffix=local
root_password=<password>
```

```

stack_password=<password>
export LIBGUESTFS_BACKEND=direct
qemu-img create -f qcow2 /var/lib/libvirt/images/${undercloud_name}.qcow2 100G
virt-resize --expand /dev/sda1 ${cloud_image}
/var/lib/libvirt/images/${undercloud_name}.qcow2
virt-customize -a /var/lib/libvirt/images/${undercloud_name}.qcow2 \
--run-command 'xfs_growfs /' \
--root-password password:${root_password} \
--hostname ${undercloud_name}.${undercloud_suffix} \
--run-command 'useradd stack' \
--password stack:password:${stack_password} \
--run-command 'echo "stack ALL=(root) NOPASSWD:ALL" | tee -a /etc/sudoers.d/stack' \
--chmod 0440:/etc/sudoers.d/stack \
--run-command 'sed -i "s/PasswordAuthentication no/PasswordAuthentication yes/g"
/etc/ssh/sshd_config' \
--run-command 'systemctl enable sshd' \
--run-command 'yum remove -y cloud-init' \
--selinux-relabel

```

NOTE: As part of the undercloud definition, a user called **stack** is created. This user will be used later to install the undercloud.

4. Define the undercloud virsh template.

```

vcpus=8
vram=32000
virt-install --name ${undercloud_name} \
--disk /var/lib/libvirt/images/${undercloud_name}.qcow2 \
--vcpus=${vcpus} \
--ram=${vram} \
--network network=default,model=virtio \
--network network=br0,model=virtio,portgroup=overcloud \
--virt-type kvm \
--import \
--os-variant rhel8.2 \
--graphics vnc \
--serial pty \
--noautoconsole \
--console pty,target_type=virtio

```


5. Start the undercloud VM.

```
virsh start ${undercloud_name}
```

6. Retrieve the undercloud IP address. It might take several seconds before the IP address is available.

```
undercloud_ip=`virsh domifaddr ${undercloud_name} |grep ipv4 |awk '{print $4}'  
|awk -F"/" '{print $1}'` ssh-copy-id ${undercloud_ip}
```

Setting Up the Undercloud

SUMMARY

Follow this topic to setting up the undercloud for Contrail Networking deployment with RHOSP 16.1.

IN THIS SECTION

- [Install the Undercloud | 225](#)
- [Perform Post-Install Configuration | 226](#)

Install the Undercloud

Use this example procedure to install the undercloud.

1. Log in to the undercloud VM from the undercloud KVM host.

```
ssh ${undercloud_ip}
```

2. Configure the hostname.

```
undercloud_name=`hostname -s`  
undercloud_suffix=`hostname -d`  
hostnamectl set-hostname ${undercloud_name}.${undercloud_suffix}  
hostnamectl set-hostname --transient ${undercloud_name}.${undercloud_suffix}
```

3. Add the hostname to the **/etc/hosts** file. The following example assumes the management interface is eth0.

```
undercloud_ip=`ip addr sh dev eth0 | grep "inet " | awk '{print $2}' | awk -F"/"
'{'print $1}'`
echo ${undercloud_ip} ${undercloud_name}.${undercloud_suffix} ${undercloud_name}
>> /etc/hosts
```

4. Set up the repositories.

RHEL

```
#Register with Satellite (can be done with CDN as well)
satellite_fqdn=device.example.net
act_key=xxx
org=example
yum localinstall -y
http://${satellite_fqdn}/pub/katello-ca-consumer-latest.noarch.rpm
subscription-manager register --activationkey=${act_key} --org=${org}
```

5. Install the Tripleo client.

```
yum install -y python-tripleoclient tmux
```

6. Copy the undercloud configuration file sample and modify the configuration as required. See [Red Hat documentation](#) for information on how to modify that file.

```
su - stack
cp /usr/share/python-tripleoclient/undercloud.conf.sample ~/undercloud.conf
vi ~/undercloud.conf
```

7. Install the undercloud.

```
openstack undercloud install
source stackrc
```

Perform Post-Install Configuration

1. Configure a forwarding path between the provisioning network and the external network:

```
sudo iptables -A FORWARD -i br-ctlplane -o eth0 -j ACCEPT
sudo iptables -A FORWARD -i eth0 -o br-ctlplane -m state --state
RELATED,ESTABLISHED -j ACCEPT
```

```
sudo iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

2. Add the external API interface:

```
sudo ip link add name vlan720 link br-ctlplane type vlan id 720
sudo ip addr add 10.2.0.254/24 dev vlan720
sudo ip link set dev vlan720 up
```

3. Add the **stack** user to the docker group:

```
newgrp docker
exit
su - stack
source stackrc
```

Setting Up the Overcloud

SUMMARY

Follow this topic to setting up the overcloud for Contrail Networking deployment with RHOSP 16.1.

IN THIS SECTION

- [Configuring the Overcloud | 227](#)
- [Customizing the Contrail Service with Templates \(contrail-services.yaml\) | 232](#)
- [Customizing the Contrail Network with Templates | 233](#)
- [Installing Overcloud | 269](#)

Configuring the Overcloud

Use this example procedure on the undercloud to set up the configuration for the overcloud.

1. Specify the name server to be used:

```
undercloud_nameserver=8.8.8.8
openstack subnet set `openstack subnet show ctlplane-subnet -c id -f value`
--dns-nameserver ${undercloud_nameserver}
```

2. Retrieve and upload the overcloud images.

- a. Create the image directory:

```
mkdir images
cd images
```

- b. Retrieve the overcloud images from either the RDO project or from Red Hat.

OSP16

```
sudo yum install -y rhosp-director-images rhosp-director-images-ipa
for i in /usr/share/rhosp-director-images/overcloud-full-latest-16.0.tar
/usr/share/rhosp-director-images/ironic-python-agent-latest-16.0.tar ; do tar
-xvf $i; done
```

- c. Upload the overcloud images:

```
cd
openstack overcloud image upload --image-path /home/stack/images/
```

- ### 3. Prepare OpenStack's bare metal provisioning (Ironic).

Ironic is an integrated OpenStack program that provisions bare metal machines instead of virtual machines. It is best thought of as a bare metal hypervisor API and a set of plugins that interact with the bare metal hypervisors.

NOTE: Make sure to combine the **ironic_list** files from the three overcloud KVM hosts.

- a. Add the overcloud VMs to ironic:

[illegible]

```

--property cpu_arch=x86_64 \
--driver-info
ipmi_username=${ipmi_user} \
--driver-info ipmi_address=${kvm_ip}
\
--driver-info
ipmi_password=${ipmi_password} \
--driver-info ipmi_port=${ipmi_port}
\
--name=${name} \
--property
capabilities=profile:${profile},boot_option:local \
-c uuid -f value`
    openstack baremetal port create --node ${uuid} ${mac}
done < <(cat ironic_list)

DEPLOY_KERNEL=$(openstack image show bm-deploy-kernel -f value -c id)
DEPLOY_RAMDISK=$(openstack image show bm-deploy-ramdisk -f value -c id)

for i in `openstack baremetal node list -c UUID -f value`; do
    openstack baremetal node set $i --driver-info deploy_kernel=$DEPLOY_KERNEL
    --driver-info deploy_ramdisk=$DEPLOY_RAMDISK
done

for i in `openstack baremetal node list -c UUID -f value`; do
    openstack baremetal node show $i -c properties -f value
done

```

b. Introspect the overcloud node:

```

for node in $(openstack baremetal node list -c UUID -f value) ; do
    openstack baremetal node manage $node
done
openstack overcloud node introspect --all-manageable --provide

```

4. Create Flavor:

```

for i in compute-dpdk \
compute-sriov \
contrail-controller \
contrail-analytics \
contrail-database \

```

```

contrail-analytics-database; do
    openstack flavor create $i --ram 4096 --vcpus 1 --disk 40
    openstack flavor set --property "capabilities:boot_option"="local" \
        --property "capabilities:profile"="${i}" ${i}
    openstack flavor set --property resources:CUSTOM_BAREMETAL=1 --property
resources:DISK_GB='0'
        --property resources:MEMORY_MB='0'
        --property resources:VCPU='0' ${i}
done

```

5. Copy the TripleO heat templates.

```
cp -r /usr/share/openstack-tripleo-heat-templates/ tripleo-heat-templates
```

6. Download and copy the Contrail heat templates from <https://support.juniper.net/support/downloads>.

```

tar -xzvf contrail-tripleo-heat-templates-<version>.tgz
cp -r contrail-tripleo-heat-templates/* tripleo-heat-templates/

```

7. Create **rhsm.yaml** file with your RedHat credentials

```

parameter_defaults:
  RhsmVars:
    rhsm_repos:
      - fast-datapath-for-rhel-8-x86_64-rpms
      - openstack-16.1-for-rhel-8-x86_64-rpms
      - satellite-tools-6.5-for-rhel-8-x86_64-rpms
      - ansible-2-for-rhel-8-x86_64-rpms
      - rhel-8-for-x86_64-highavailability-rpms
      - rhel-8-for-x86_64-appstream-rpms
      - rhel-8-for-x86_64-baseos-rpms
    rhsm_username: "YOUR_REDHAT_LOGIN"
    rhsm_password: "YOUR_REDHAT_PASSWORD"
    rhsm_org_id: "YOUR_REDHAT_ID"
    rhsm_pool_ids: "YOUR_REDHAT_POOL_ID"

```

8. Create and upload the OpenStack containers.
 - a. Create the OpenStack container file.

NOTE: The container must be created based on the OpenStack program.

OSP16

```
sudo openstack tripleo container image prepare \
  -e ~/containers-prepare-parameter.yaml
  -e ~/rhsm.yaml > ~/overcloud_containers.yaml

sudo openstack overcloud container image upload --config-file
~/overcloud_containers.yaml
```

- b. Upload the OpenStack containers:

```
openstack overcloud container image upload --config-file
~/local_registry_images.yaml
```

9. Create and upload the Contrail containers.

- a. Create the Contrail container file.

NOTE: This step is optional. The Contrail containers can be downloaded from external registries later.

```
cd ~/tf-heat-templates/tools/contrail
./import_contrail_container.sh -f container_outputfile -r registry -t tag [-i
insecure] [-u username] [-p password] [-c certificate path]
```

Here are few examples of importing Contrail containers from different sources:

- Import from password protected public registry:

```
./import_contrail_container.sh -f /tmp/contrail_container -r
hub.juniper.net/contrail -u USERNAME -p PASSWORD -t 1234
```

- Import from Dockerhub:

```
./import_contrail_container.sh -f /tmp/contrail_container -r
docker.io/opencontrailnightly -t 1234
```

- Import from private secure registry:

```
./import_contrail_container.sh -f /tmp/contrail_container -r
device.example.net:5443 -c
http://device.example.net/pub/device.example.net.crt -t 1234
```

- Import from private insecure registry:

```
./import_contrail_container.sh -f /tmp/contrail_container -r 10.0.0.1:5443
-i 1 -t 1234
```

- Upload Contrail containers to the undercloud registry:

```
openstack overcloud container image upload --config-file
/tmp/contrail_container
```

Customizing the Contrail Service with Templates (contrail-services.yaml)

This section contains information to customize Contrail services for your network by modifying the **contrail-services.yaml** file.

- Contrail Services customization

```
vi ~/tripleo-heat-templates/environments/contrail-services.yaml
```

```
parameter_defaults:
  ContrailSettings:
    VROUTER_GATEWAY: 10.0.0.1
    # KEY1: value1
    # KEY2: value2

    VXLAN_VN_ID_MODE: "configured"
    ENCAP_PRIORITY: "VXLAN,MPLSoUDP,MPLSoGRE"

  ContrailControllerParameters:
    AAAMode: rbac
```


- Contrail registry settings

```
vi ~/tripleo-heat-templates/environments/contrail-services.yaml
```

Here are few examples of default values for various registries:

- Public Juniper registry

```
parameter_defaults:
  ContrailRegistry: hub.juniper.net/contrail
  ContrailRegistryUser: <USER>
  ContrailRegistryPassword: <PASSWORD>
```

- Insecure registry

```
parameter_defaults:
  ContrailRegistryInsecure: true
  DockerInsecureRegistryAddress: 10.87.64.32:5000,192.168.24.1:8787
  ContrailRegistry: 10.87.64.32:5000
```

- Private secure registry

```
parameter_defaults:
  ContrailRegistryCertUrl: http://device.example.net/pub/device.example.net.crt

  ContrailRegistry: device.example.net:5443
```

- Contrail Container image settings

```
parameter_defaults:
  ContrailImageTag: queens-5.0-104-rhel-queens
```

Customizing the Contrail Network with Templates

IN THIS SECTION

- [Overview | 234](#)
- [Roles Configuration \(roles_data_contrail_aio.yaml\) | 234](#)
- [Network Parameter Configuration \(contrail-net.yaml\) | 237](#)

- [Network Interface Configuration \(*-NIC-*.yaml\) | 239](#)
- [Advanced vRouter Kernel Mode Configuration | 250](#)
- [Advanced vRouter DPDK Mode Configuration | 252](#)
- [Advanced vRouter SRIOV + Kernel Mode Configuration | 255](#)
- [Advanced vRouter SRIOV + DPDK Mode Configuration | 258](#)
- [Advanced Scenarios | 261](#)

Overview

In order to customize the network, define different networks and configure the overcloud nodes NIC layout. TripleO supports a flexible way of customizing the network.

The following networking customization example uses network as:

Table 10: Network Customization

Network	VLAN	overcloud Nodes
provisioning	-	All
internal_api	710	All
external_api	720	OpenStack CTRL
storage	740	OpenStack CTRL, Computes
storage_mgmt	750	OpenStack CTRL
tenant	-	Contrail CTRL, Computes

Roles Configuration (roles_data_contrail_aio.yaml)

IN THIS SECTION

- [OpenStack Controller | 235](#)
- [Compute Node | 235](#)
- [Contrail Controller | 236](#)
- [Compute DPDK | 236](#)

- [Compute SRIOV | 236](#)
- [Compute CSN | 237](#)

The networks must be activated per role in the roles_data file:

```
vi ~/tripleo-heat-templates/roles_data_contrail_aio.yaml
```

OpenStack Controller

```
#####
# Role: Controller                                                                    #
#####
- name: Controller
  description: |
    Controller role that has all the controller services loaded and handles
    Database, Messaging and Network functions.
  CountDefault: 1
  tags:
    - primary
    - controller
  networks:
    - External
    - InternalApi
    - Storage
    - StorageMgmt
```

Compute Node

```
#####
# Role: Compute                                                                    #
#####
- name: Compute
  description: |
    Basic Compute Node role
  CountDefault: 1
  networks:
    - InternalApi
```

- Tenant
- Storage

Contrail Controller

```
#####
# Role: ContrailController                                     #
#####
- name: ContrailController
  description: |
    ContrailController role that has all the Contrail controller services loaded
    and handles config, control and webui functions
  CountDefault: 1
  tags:
    - primary
    - contrailcontroller
  networks:
    - InternalApi
    - Tenant
```

Compute DPDK

```
#####
# Role: ContrailDpdk                                         #
#####
- name: ContrailDpdk
  description: |
    Contrail Dpdk Node role
  CountDefault: 0
  tags:
    - contraildpdk
  networks:
    - InternalApi
    - Tenant
    - Storage
```

Compute SRIOV

```
#####
# Role: ContrailSriov
```

```
#####
- name: ContrailSriov
  description: |
    Contrail Sriov Node role
  CountDefault: 0
  tags:
    - contrailsriov
  networks:
    - InternalApi
    - Tenant
    - Storage
```

Compute CSN

```
#####
# Role: ContrailTsn
#####
- name: ContrailTsn
  description: |
    Contrail Tsn Node role
  CountDefault: 0
  tags:
    - contrailtsn
  networks:
    - InternalApi
    - Tenant
    - Storage
```

Network Parameter Configuration (contrail-net.yaml)

```
cat ~/tripleo-heat-templates/environments/contrail/contrail-net.yaml
```

```
resource_registry:
  OS::TripleO::Controller::Net::SoftwareConfig:
    ../../network/config/contrail/controller-nic-config.yaml
  OS::TripleO::ContrailController::Net::SoftwareConfig:
    ../../network/config/contrail/contrail-controller-nic-config.yaml
  OS::TripleO::ContrailControlOnly::Net::SoftwareConfig:
    ../../network/config/contrail/contrail-controller-nic-config.yaml
  OS::TripleO::Compute::Net::SoftwareConfig:
    ../../network/config/contrail/compute-nic-config.yaml
```

```

OS::TripleO::ContrailDpdk::Net::SoftwareConfig:
../../network/config/contrail/contrail-dpdk-nic-config.yaml
OS::TripleO::ContrailSriov::Net::SoftwareConfig:
../../network/config/contrail/contrail-sriov-nic-config.yaml
OS::TripleO::ContrailTsn::Net::SoftwareConfig:
../../network/config/contrail/contrail-tsn-nic-config.yaml

```

```

parameter_defaults:
  # Customize all these values to match the local environment
  TenantNetCidr: 10.0.0.0/24
  InternalApiNetCidr: 10.1.0.0/24
  ExternalNetCidr: 10.2.0.0/24
  StorageNetCidr: 10.3.0.0/24
  StorageMgmtNetCidr: 10.4.0.0/24
  # CIDR subnet mask length for provisioning network
  ControlPlaneSubnetCidr: '24'
  # Allocation pools
  TenantAllocationPools: [{'start': '10.0.0.10', 'end': '10.0.0.200'}]
  InternalApiAllocationPools: [{'start': '10.1.0.10', 'end': '10.1.0.200'}]
  ExternalAllocationPools: [{'start': '10.2.0.10', 'end': '10.2.0.200'}]
  StorageAllocationPools: [{'start': '10.3.0.10', 'end': '10.3.0.200'}]
  StorageMgmtAllocationPools: [{'start': '10.4.0.10', 'end': '10.4.0.200'}]
  # Routes
  ControlPlaneDefaultRoute: 192.168.24.1
  InternalApiDefaultRoute: 10.1.0.1
  ExternalInterfaceDefaultRoute: 10.2.0.1
  # Vlans
  InternalApiNetworkVlanID: 710
  ExternalNetworkVlanID: 720
  StorageNetworkVlanID: 730
  StorageMgmtNetworkVlanID: 740
  TenantNetworkVlanID: 3211
  # Services
  EC2MetadataIp: 192.168.24.1 # Generally the IP of the undercloud
  DnsServers: ["172.x.x.x"]
  NtpServer: 10.0.0.1

```

Network Interface Configuration (*-NIC-*.yaml)

IN THIS SECTION

- [OpenStack Controller | 239](#)
- [Contrail Controller | 243](#)
- [Compute Node | 246](#)

NIC configuration files exist per role in the following directory:

```
cd ~/tripleo-heat-templates/network/config/contrail
```

OpenStack Controller

```
heat_template_version: rocky
```

```
description: >
```

```
Software Config to drive os-net-config to configure multiple interfaces
for the compute role. This is an example for a Nova compute node using
Contrail vrouter and the vhost0 interface.
```

```
parameters:
```

```
ControlPlaneIp:
```

```
  default: ''
```

```
  description: IP address/subnet on the ctlplane network
```

```
  type: string
```

```
ExternalIpSubnet:
```

```
  default: ''
```

```
  description: IP address/subnet on the external network
```

```
  type: string
```

```
InternalApiIpSubnet:
```

```
  default: ''
```

```
  description: IP address/subnet on the internal_api network
```

```
  type: string
```

```
InternalApiDefaultRoute: # Not used by default in this template
```

```
  default: '10.0.0.1'
```

```
  description: The default route of the internal api network.
```

```
  type: string
```

```
StorageIpSubnet:
```

```

    default: ''
    description: IP address/subnet on the storage network
    type: string
StorageMgmtIpSubnet:
    default: ''
    description: IP address/subnet on the storage_mgmt network
    type: string
TenantIpSubnet:
    default: ''
    description: IP address/subnet on the tenant network
    type: string
ManagementIpSubnet: # Only populated when including
environments/network-management.yaml
    default: ''
    description: IP address/subnet on the management network
    type: string
ExternalNetworkVlanID:
    default: 10
    description: Vlan ID for the external network traffic.
    type: number
InternalApiNetworkVlanID:
    default: 20
    description: Vlan ID for the internal_api network traffic.
    type: number
StorageNetworkVlanID:
    default: 30
    description: Vlan ID for the storage network traffic.
    type: number
StorageMgmtNetworkVlanID:
    default: 40
    description: Vlan ID for the storage mgmt network traffic.
    type: number
TenantNetworkVlanID:
    default: 50
    description: Vlan ID for the tenant network traffic.
    type: number
ManagementNetworkVlanID:
    default: 60
    description: Vlan ID for the management network traffic.
    type: number
ControlPlaneSubnetCidr: # Override this via parameter_defaults
    default: '24'
    description: The subnet CIDR of the control plane network.
    type: string

```



```

    routes:
      - ip_netmask: 169.x.x.x/32
        next_hop:
          get_param: EC2MetadataIp
      - default: true
        next_hop:
          get_param: ControlPlaneDefaultRoute
    - type: vlan
      vlan_id:
        get_param: InternalApiNetworkVlanID
      device: nic1
      addresses:
        - ip_netmask:
            get_param: InternalApiIpSubnet
    - type: vlan
      vlan_id:
        get_param: ExternalNetworkVlanID
      device: nic1
      addresses:
        - ip_netmask:
            get_param: ExternalIpSubnet
    - type: vlan
      vlan_id:
        get_param: StorageNetworkVlanID
      device: nic1
      addresses:
        - ip_netmask:
            get_param: StorageIpSubnet
    - type: vlan
      vlan_id:
        get_param: StorageMgmtNetworkVlanID
      device: nic1
      addresses:
        - ip_netmask:
            get_param: StorageMgmtIpSubnet

```

```

outputs:

```

```

  OS::stack_id:
    description: The OsNetConfigImpl resource.
    value:
      get_resource: OsNetConfigImpl

```

Contrail Controller

```
heat_template_version: rocky
```

```
description: >
```

```
Software Config to drive os-net-config to configure multiple interfaces
for the compute role. This is an example for a Nova compute node using
Contrail vrouter and the vhost0 interface.
```

```
parameters:
```

```
ControlPlaneIp:
```

```
  default: ''
```

```
  description: IP address/subnet on the ctlplane network
```

```
  type: string
```

```
ExternalIpSubnet:
```

```
  default: ''
```

```
  description: IP address/subnet on the external network
```

```
  type: string
```

```
InternalApiIpSubnet:
```

```
  default: ''
```

```
  description: IP address/subnet on the internal_api network
```

```
  type: string
```

```
InternalApiDefaultRoute: # Not used by default in this template
```

```
  default: '10.0.0.1'
```

```
  description: The default route of the internal api network.
```

```
  type: string
```

```
StorageIpSubnet:
```

```
  default: ''
```

```
  description: IP address/subnet on the storage network
```

```
  type: string
```

```
StorageMgmtIpSubnet:
```

```
  default: ''
```

```
  description: IP address/subnet on the storage_mgmt network
```

```
  type: string
```

```
TenantIpSubnet:
```

```
  default: ''
```

```
  description: IP address/subnet on the tenant network
```

```
  type: string
```

```
ManagementIpSubnet: # Only populated when including
environments/network-management.yaml
```

```
  default: ''
```

```
  description: IP address/subnet on the management network
```

```
  type: string
```

```

ExternalNetworkVlanID:
  default: 10
  description: Vlan ID for the external network traffic.
  type: number
InternalApiNetworkVlanID:
  default: 20
  description: Vlan ID for the internal_api network traffic.
  type: number
StorageNetworkVlanID:
  default: 30
  description: Vlan ID for the storage network traffic.
  type: number
StorageMgmtNetworkVlanID:
  default: 40
  description: Vlan ID for the storage mgmt network traffic.
  type: number
TenantNetworkVlanID:
  default: 50
  description: Vlan ID for the tenant network traffic.
  type: number
ManagementNetworkVlanID:
  default: 60
  description: Vlan ID for the management network traffic.
  type: number
ControlPlaneSubnetCidr: # Override this via parameter_defaults
  default: '24'
  description: The subnet CIDR of the control plane network.
  type: string
ControlPlaneDefaultRoute: # Override this via parameter_defaults
  description: The default route of the control plane network.
  type: string
ExternalInterfaceDefaultRoute: # Not used by default in this template
  default: '10.0.0.1'
  description: The default route of the external network.
  type: string
ManagementInterfaceDefaultRoute: # Commented out by default in this template
  default: unset
  description: The default route of the management network.
  type: string
DnsServers: # Override this via parameter_defaults
  default: []
  description: A list of DNS servers (2 max for some implementations) that will
be added to resolv.conf.
  type: comma_delimited_list

```

```

EC2MetadataIp: # Override this via parameter_defaults
  description: The IP address of the EC2 metadata server.
  type: string

```

```

resources:
  OsNetConfigImpl:
    type: OS::Heat::SoftwareConfig
    properties:
      group: script
      config:
        str_replace:
          template:
            get_file: ../../scripts/run-os-net-config.sh
          params:
            $network_config:
              network_config:
                - type: interface
                  name: nic1
                  use_dhcp: false
                  dns_servers:
                    get_param: DnsServers
                  addresses:
                - ip_netmask:
                    list_join:
                      - '/'
                    - - get_param: ControlPlaneIp
                      - get_param: ControlPlaneSubnetCidr
                  routes:
                - ip_netmask: 169.x.x.x/32
                  next_hop:
                    get_param: EC2MetadataIp
                - default: true
                  next_hop:
                    get_param: ControlPlaneDefaultRoute
            - type: vlan
              vlan_id:
                get_param: InternalApiNetworkVlanID
              device: nic1
              addresses:
                - ip_netmask:
                    get_param: InternalApiIpSubnet
            - type: interface
              name: nic2

```

```

        use_dhcp: false
        addresses:
        - ip_netmask:
            get_param: TenantIpSubnet

```

```

outputs:
  OS::stack_id:
    description: The OsNetConfigImpl resource.
    value:
      get_resource: OsNetConfigImpl

```

Compute Node

```
heat_template_version: rocky
```

```

description: >
  Software Config to drive os-net-config to configure multiple interfaces
  for the compute role. This is an example for a Nova compute node using
  Contrail vrouter and the vhost0 interface.

```

```

parameters:
  ControlPlaneIp:
    default: ''
    description: IP address/subnet on the ctlplane network
    type: string
  ExternalIpSubnet:
    default: ''
    description: IP address/subnet on the external network
    type: string
  InternalApiIpSubnet:
    default: ''
    description: IP address/subnet on the internal_api network
    type: string
  InternalApiDefaultRoute: # Not used by default in this template
    default: '10.0.0.1'
    description: The default route of the internal api network.
    type: string
  StorageIpSubnet:
    default: ''
    description: IP address/subnet on the storage network

```

```

    type: string
StorageMgmtIpSubnet:
  default: ''
  description: IP address/subnet on the storage_mgmt network
  type: string
TenantIpSubnet:
  default: ''
  description: IP address/subnet on the tenant network
  type: string
ManagementIpSubnet: # Only populated when including
environments/network-management.yaml
  default: ''
  description: IP address/subnet on the management network
  type: string
ExternalNetworkVlanID:
  default: 10
  description: Vlan ID for the external network traffic.
  type: number
InternalApiNetworkVlanID:
  default: 20
  description: Vlan ID for the internal_api network traffic.
  type: number
StorageNetworkVlanID:
  default: 30
  description: Vlan ID for the storage network traffic.
  type: number
StorageMgmtNetworkVlanID:
  default: 40
  description: Vlan ID for the storage mgmt network traffic.
  type: number
TenantNetworkVlanID:
  default: 50
  description: Vlan ID for the tenant network traffic.
  type: number
ManagementNetworkVlanID:
  default: 60
  description: Vlan ID for the management network traffic.
  type: number
ControlPlaneSubnetCidr: # Override this via parameter_defaults
  default: '24'
  description: The subnet CIDR of the control plane network.
  type: string
ControlPlaneDefaultRoute: # Override this via parameter_defaults
  description: The default route of the control plane network.

```

```

    type: string
ExternalInterfaceDefaultRoute: # Not used by default in this template
    default: '10.0.0.1'
    description: The default route of the external network.
    type: string
ManagementInterfaceDefaultRoute: # Commented out by default in this template
    default: unset
    description: The default route of the management network.
    type: string
DnsServers: # Override this via parameter_defaults
    default: []
    description: A list of DNS servers (2 max for some implementations) that will
be added to resolv.conf.
    type: comma_delimited_list
EC2MetadataIp: # Override this via parameter_defaults
    description: The IP address of the EC2 metadata server.
    type: string

```

```

resources:
  OsNetConfigImpl:
    type: OS::Heat::SoftwareConfig
    properties:
      group: script
      config:
        str_replace:
          template:
            get_file: ../../scripts/run-os-net-config.sh
          params:
            $network_config:
              network_config:
                - type: interface
                  name: nic1
                  use_dhcp: false
                  dns_servers:
                    get_param: DnsServers
                  addresses:
                    - ip_netmask:
                        list_join:
                          - '/'
                        - - get_param: ControlPlaneIp
                          - get_param: ControlPlaneSubnetCidr
                  routes:
                    - ip_netmask: 169.x.x.x/32

```



```

      next_hop:
        get_param: EC2MetadataIp
    - default: true
      next_hop:
        get_param: ControlPlaneDefaultRoute
  - type: vlan
    vlan_id:
      get_param: InternalApiNetworkVlanID
    device: nic1
    addresses:
      - ip_netmask:
          get_param: InternalApiIpSubnet
  - type: vlan
    vlan_id:
      get_param: StorageNetworkVlanID
    device: nic1
    addresses:
      - ip_netmask:
          get_param: StorageIpSubnet
  - type: contrail_vrouter
    name: vhost0
    use_dhcp: false
    members:
      -
        type: interface
        name: nic2
        use_dhcp: false
    addresses:
      - ip_netmask:
          get_param: TenantIpSubnet

```

```

outputs:
  OS::stack_id:
    description: The OsNetConfigImpl resource.
    value:
      get_resource: OsNetConfigImpl

```

Advanced vRouter Kernel Mode Configuration

IN THIS SECTION

- [VLAN | 250](#)
- [Bond | 250](#)
- [Bond + VLAN | 251](#)

In addition to the standard NIC configuration, the vRouter kernel mode supports VLAN, Bond, and Bond + VLAN modes. The configuration snippets below only show the relevant section of the NIC template configuration for each mode.

VLAN

```
- type: vlan
  vlan_id:
    get_param: TenantNetworkVlanID
  device: nic2
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
    -
      type: interface
      name:
        str_replace:
          template: vlanVLANID
          params:
            VLANID: {get_param: TenantNetworkVlanID}
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet
```

Bond

```
- type: linux_bond
  name: bond0
  bonding_options: "mode=4 xmit_hash_policy=layer2+3"
  use_dhcp: false
```

```

members:
  -
    type: interface
    name: nic2
  -
    type: interface
    name: nic3
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
    -
      type: interface
      name: bond0
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet

```

Bond + VLAN

```

- type: linux_bond
  name: bond0
  bonding_options: "mode=4 xmit_hash_policy=layer2+3"
  use_dhcp: false
  members:
    -
      type: interface
      name: nic2
    -
      type: interface
      name: nic3
- type: vlan
  vlan_id:
    get_param: TenantNetworkVlanID
  device: bond0
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
    -
      type: interface
      name:

```

```

    str_replace:
      template: vlanVLANID
      params:
        VLANID: {get_param: TenantNetworkVlanID}
    use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet

```

Advanced vRouter DPDK Mode Configuration

IN THIS SECTION

- [Standard | 253](#)
- [VLAN | 253](#)
- [Bond | 253](#)
- [Bond + VLAN | 254](#)

In addition to the standard NIC configuration, the vRouter DPDK mode supports Standard, VLAN, Bond, and Bond + VLAN modes.

Network Environment Configuration:

```
vi ~/tripleo-heat-templates/environments/contrail/contrail-services.yaml
```

Enable the number of hugepages:

```

# For Intel CPU
ContrailDpdkParameters:
  KernelArgs: "intel_iommu=on iommu=pt default_hugepagesz=1GB hugepagesz=1G
hugepages=4 hugepagesz=2M hugepages=1024"
  ExtraSysctlSettings:
    # must be equal to value from kernel args: hugepages=4
    vm.nr_hugepages:
      value: 4
    vm.max_map_count:
      value: 128960

```

See the following NIC template configurations for vRouter DPDK mode. The configuration snippets below only show the relevant section of the NIC configuration for each mode.

Standard

```
- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
  addresses:
    - ip_netmask:
      get_param: TenantIpSubnet
```

VLAN

```
- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  vlan_id:
    get_param: TenantNetworkVlanID
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
  addresses:
    - ip_netmask:
      get_param: TenantIpSubnet
```

Bond

```
- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
```

```

driver: uio_pci_generic
cpu_list: 0x01
bond_mode: 4
bond_policy: layer2+3
members:
  -
    type: interface
    name: nic2
    use_dhcp: false
  -
    type: interface
    name: nic3
    use_dhcp: false
addresses:
- ip_netmask:
    get_param: TenantIpSubnet

```

Bond + VLAN

```

- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  vlan_id:
    get_param: TenantNetworkVlanID
  bond_mode: 4
  bond_policy: layer2+3
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
    -
      type: interface
      name: nic3
      use_dhcp: false
  addresses:
- ip_netmask:
    get_param: TenantIpSubnet

```

Advanced vRouter SRIOV + Kernel Mode Configuration

IN THIS SECTION

- VLAN | 256
- Bond | 256
- Bond + VLAN | 257

vRouter SRIOV + Kernel mode can be used in the following combinations:

- Standard
- VLAN
- Bond
- Bond + VLAN

Network environment configuration:

```
vi ~/tripleo-heat-templates/environments/contrail/contrail-services.yaml
```

Enable the number of hugepages:

```
ContrailSriovParameters:
  KernelArgs: "intel_iommu=on iommu=pt default_hugepagesz=1GB hugepagesz=1G
hugepages=4 hugepagesz=2M hugepages=1024"
  ExtraSysctlSettings:
    # must be equal to value from 1G kernel args: hugepages=4
    vm.nr_hugepages:
      value: 4
```

SRIOV PF/VF settings:

```
NovaPCIPassthrough:
- devname: "ens2f1"
  physical_network: "sriov1"
ContrailSriovNumVFs: ["ens2f1:7"]
```

The SRIOV NICs are not configured in the NIC templates. However, vRouter NICs must still be configured. See the following NIC template configurations for vRouter kernel mode. The configuration snippets below only show the relevant section of the NIC configuration for each mode.

VLAN

```
- type: vlan
  vlan_id:
    get_param: TenantNetworkVlanID
  device: nic2
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
    -
      type: interface
      name:
        str_replace:
          template: vlanVLANID
          params:
            VLANID: {get_param: TenantNetworkVlanID}
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet
```

Bond

```
- type: linux_bond
  name: bond0
  bonding_options: "mode=4 xmit_hash_policy=layer2+3"
  use_dhcp: false
  members:
    -
      type: interface
      name: nic2
    -
      type: interface
      name: nic3
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
    -
```



```

    type: interface
    name: bond0
    use_dhcp: false
addresses:
- ip_netmask:
    get_param: TenantIpSubnet

```

Bond + VLAN

```

- type: linux_bond
  name: bond0
  bonding_options: "mode=4 xmit_hash_policy=layer2+3"
  use_dhcp: false
  members:
    -
      type: interface
      name: nic2
    -
      type: interface
      name: nic3
- type: vlan
  vlan_id:
    get_param: TenantNetworkVlanID
  device: bond0
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
    -
      type: interface
      name:
        str_replace:
          template: vlanVLANID
          params:
            VLANID: {get_param: TenantNetworkVlanID}
      use_dhcp: false
addresses:
- ip_netmask:
    get_param: TenantIpSubnet

```

Advanced vRouter SRIOV + DPDK Mode Configuration

IN THIS SECTION

- [Standard | 259](#)
- [VLAN | 259](#)
- [Bond | 259](#)
- [Bond + VLAN | 260](#)

vRouter SRIOV + DPDK can be used in the following combinations:

- Standard
- VLAN
- Bond
- Bond + VLAN

Network environment configuration:

```
vi ~/tripleo-heat-templates/environments/contrail/contrail-services.yaml
```

Enable the number of hugepages

```
ContrailSriovParameters:
  KernelArgs: "intel_iommu=on iommu=pt default_hugepagesz=1GB hugepagesz=1G
hugepages=4 hugepagesz=2M hugepages=1024"
  ExtraSysctlSettings:
    # must be equal to value from 1G kernel args: hugepages=4
    vm.nr_hugepages:
      value: 4
```

SRIOV PF/VF settings

```
NovaPCIPassthrough:
- devname: "ens2f1"
  physical_network: "sriov1"
ContrailSriovNumVFs: ["ens2f1:7"]
```

The SRIOV NICs are not configured in the NIC templates. However, vRouter NICs must still be configured. See the following NIC template configurations for vRouter DPDK mode. The configuration snippets below only show the relevant section of the NIC configuration for each mode.

Standard

```
- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet
```

VLAN

```
- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  vlan_id:
    get_param: TenantNetworkVlanID
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet
```

Bond

```
- type: contrail_vrouter_dpdk
  name: vhost0
```

```

use_dhcp: false
driver: uio_pci_generic
cpu_list: 0x01
bond_mode: 4
bond_policy: layer2+3
members:
  -
    type: interface
    name: nic2
    use_dhcp: false
  -
    type: interface
    name: nic3
    use_dhcp: false
addresses:
  - ip_netmask:
      get_param: TenantIpSubnet

```

Bond + VLAN

```

- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  vlan_id:
    get_param: TenantNetworkVlanID
  bond_mode: 4
  bond_policy: layer2+3
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
    -
      type: interface
      name: nic3
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet

```

Advanced Scenarios


Remote Compute

Remote Compute extends the data plane to remote locations (POP) whilst keeping the control plane central. Each POP will have its own set of Contrail control services, which are running in the central location. The difficulty is to ensure that the compute nodes of a given POP connect to the Control nodes assigned to that POP. The Control nodes must have predictable IP addresses and the compute nodes have to know these IP addresses. In order to achieve that the following methods are used:

- Custom Roles
- Static IP assignment
- Precise Node placement
- Per Node hieradata

Each overcloud node has a unique DMI UUID. This UUID is known on the undercloud node as well as on the overcloud node. Hence, this UUID can be used for mapping node specific information. For each POP, a Control role and a Compute role has to be created.

Overview



Mapping Table

Table 11: Mapping Table

Nova Name	Ironic Name	UUID	KVM	IP Address	POP
overcloud -contrailcontrolonly -0	control-only-1- 5b3s30	Ironic UUID: 7d758dce-2784- 45fd-be09-5a41eb53e764 DMI UUID: 73F8D030-E896- 4A95-A9F5-E1A4FEBE322D	5b3s30	10.0.0.11	POP1

Table 11: Mapping Table (continued)

Nova Name	Ironic Name	UUID	KVM	IP Address	POP
overcloud -contrailcontrolonly -1	control-only-2- 5b3s30	Ironic UUID: d26abdeb-d514-4a37-a7fb-2cd2511c351f DMI UUID: 14639A66-D62C-4408-82EE-FDDC4E509687	5b3s30	10.0.0.14	POP2
overcloud -contrailcontrolonly -2	control-only-1- 5b3s31	Ironic UUID: 91dd9fa9-e8eb-4b51-8b5e-bbaffb6640e4 DMI UUID: 28AB0B57-D612-431E-B177-1C578AE0FEA4	5b3s31	10.0.0.12	POP1
overcloud -contrailcontrolonly -3	control-only-2- 5b3s31	Ironic UUID: 09fa57b8-580f-42ec-bf10-a19573521ed4 DMI UUID: 09BEC8CB-77E9-42A6-AFF4-6D4880FD87D0	5b3s31	10.0.0.15	POP2
overcloud -contrailcontrolonly -4	control-only-1- 5b3s32	Ironic UUID: 4766799-24c8-4e3b-af54-353f2b796ca4 DMI UUID: 3993957A-ECBF-4520-9F49-0AF6EE1667A7	5b3s32	10.0.0.13	POP1
overcloud -contrailcontrolonly -5	control-only-2- 5b3s32	Ironic UUID: 58a803ae-a785-470e-9789-139abbfa74fb DMI UUID: AF92F485-C30C-4D0A-BDC4-C6AE97D06A66	5b3s32	10.0.0.16	POP2

ControlOnly preparation

Add ControlOnly overcloud VMs to overcloud KVM host

NOTE: This has to be done on the overcloud KVM hosts

Two ControlOnly overcloud VM definitions will be created on each of the overcloud KVM hosts.

```

ROLES=control-only:2
num=4
ipmi_user=<user>
ipmi_password=<password>
libvirt_path=/var/lib/libvirt/images
port_group=overcloud
prov_switch=br0

/bin/rm ironic_list
IFS=' ' read -ra role_list <<< "${ROLES}"
for role in ${role_list[@]}; do
    role_name=`echo $role|cut -d ":" -f 1`
    role_count=`echo $role|cut -d ":" -f 2`
    for count in `seq 1 ${role_count}`; do
        echo $role_name $count
        qemu-img create -f qcow2 ${libvirt_path}/${role_name}_${count}.qcow2 99G
        virsh define /dev/stdin <<EOF
$(virt-install --name ${role_name}_${count} \
--disk ${libvirt_path}/${role_name}_${count}.qcow2 \
--vcpus=4 \
--ram=16348 \
--network network=br0,model=virtio,portgroup=${port_group} \
--network network=br1,model=virtio \
--virt-type kvm \
--cpu host \
--import \
--os-variant rhel7 \
--serial pty \
--console pty,target_type=virtio \
--graphics vnc \
--print-xml)
EOF
        vbmc add ${role_name}_${count} --port 1623${num} --username ${ipmi_user}
        --password ${ipmi_password}
        vbmc start ${role_name}_${count}
        prov_mac=`virsh domiflist ${role_name}_${count}|grep ${prov_switch}|awk '{print $5}'`
        vm_name=${role_name}-${count}-${hostname} -s`
        kvm_ip=`ip route get 1 |grep src |awk '{print $7}'`
        echo ${prov_mac} ${vm_name} ${kvm_ip} ${role_name} 1623${num}>> ironic_list
        num=$(expr $num + 1)
    done
done

```


NOTE: The generated *ironic_list* will be needed on the undercloud to import the nodes to Ironic.

Get the *ironic_lists* from the overcloud KVM hosts and combine them.

```
cat ironic_list_control_only
52:54:00:3a:2f:ca control-only-1-5b3s30 10.87.64.31 control-only 16234
52:54:00:31:4f:63 control-only-2-5b3s30 10.87.64.31 control-only 16235
52:54:00:0c:11:74 control-only-1-5b3s31 10.87.64.32 control-only 16234
52:54:00:56:ab:55 control-only-2-5b3s31 10.87.64.32 control-only 16235
52:54:00:c1:f0:9a control-only-1-5b3s32 10.87.64.33 control-only 16234
52:54:00:f3:ce:13 control-only-2-5b3s32 10.87.64.33 control-only 16235
```

Import:

```
ipmi_password=<password>
ipmi_user=<user>

DEPLOY_KERNEL=$(openstack image show bm-deploy-kernel -f value -c id)
DEPLOY_RAMDISK=$(openstack image show bm-deploy-ramdisk -f value -c id)

num=0
while IFS= read -r line; do
    mac=`echo $line|awk '{print $1}'`
    name=`echo $line|awk '{print $2}'`
    kvm_ip=`echo $line|awk '{print $3}'`
    profile=`echo $line|awk '{print $4}'`
    ipmi_port=`echo $line|awk '{print $5}'`
    uuid=`openstack baremetal node create --driver ipmi \
        --property cpus=4 \
        --property memory_mb=16348 \
        --property local_gb=100 \
        --property cpu_arch=x86_64 \
        --driver-info ipmi_username=${ipmi_user} \
        \
        --driver-info ipmi_address=${kvm_ip} \
        --driver-info ipmi_password=${ipmi_password} \
        \
        --driver-info ipmi_port=${ipmi_port} \
        --name=${name} \
        --property capabilities=boot_option:local
```

```
\
                                -c uuid -f value`
openstack baremetal node set ${uuid} --driver-info deploy_kernel=$DEPLOY_KERNEL
--driver-info deploy_ramdisk=$DEPLOY_RAMDISK
openstack baremetal port create --node ${uuid} ${mac}
openstack baremetal node manage ${uuid}
num=$(expr $num + 1)
done < <(cat ironic_list_control_only)
```

ControlOnly node introspection

```
openstack overcloud node introspect --all-manageable --provide
```

Get the ironic UUID of the ControlOnly nodes

```
openstack baremetal node list |grep control-only
| 7d758dce-2784-45fd-be09-5a41eb53e764 | control-only-1-5b3s30 | None | power off
| available | False |
| d26abdeb-d514-4a37-a7fb-2cd2511c351f | control-only-2-5b3s30 | None | power off
| available | False |
| 91dd9fa9-e8eb-4b51-8b5e-bbaffb6640e4 | control-only-1-5b3s31 | None | power off
| available | False |
| 09fa57b8-580f-42ec-bf10-a19573521ed4 | control-only-2-5b3s31 | None | power off
| available | False |
| f4766799-24c8-4e3b-af54-353f2b796ca4 | control-only-1-5b3s32 | None | power off
| available | False |
| 58a803ae-a785-470e-9789-139abbfa74fb | control-only-2-5b3s32 | None | power off
| available | False |
```

The first ControlOnly node on each of the overcloud KVM hosts will be used for POP1, the second for POP2, and so and so forth.

Get the ironic UUID of the POP compute nodes:

```
openstack baremetal node list |grep compute
| 91d6026c-b9db-49cb-a685-99a63da5d81e | compute-3-5b3s30 | None | power off |
available | False |
| 8028eb8c-e1e6-4357-8fcf-0796778bd2f7 | compute-4-5b3s30 | None | power off |
available | False |
| b795b3b9-c4e3-4a76-90af-258d9336d9fb | compute-3-5b3s31 | None | power off |
```

```
available | False |
| 2d4be83e-6fcc-4761-86f2-c2615dd15074 | compute-4-5b3s31 | None | power off |
available | False |
```

The first two compute nodes belong to POP1 the second two compute nodes belong to POP2.

Create an input YAML using the ironic UUIDs:

```
~/subcluster_input.yaml
---
- subcluster: subcluster1
  asn: "65413"
  control_nodes:
    - uuid: 7d758dce-2784-45fd-be09-5a41eb53e764
      ipaddress: 10.0.0.11
    - uuid: 91dd9fa9-e8eb-4b51-8b5e-bbaffb6640e4
      ipaddress: 10.0.0.12
    - uuid: f4766799-24c8-4e3b-af54-353f2b796ca4
      ipaddress: 10.0.0.13
  compute_nodes:
    - uuid: 91d6026c-b9db-49cb-a685-99a63da5d81e
      vrouter_gateway: 10.0.0.1
    - uuid: 8028eb8c-e1e6-4357-8fcf-0796778bd2f7
      vrouter_gateway: 10.0.0.1
- subcluster: subcluster2
  asn: "65414"
  control_nodes:
    - uuid: d26abdeb-d514-4a37-a7fb-2cd2511c351f
      ipaddress: 10.0.0.14
    - uuid: 09fa57b8-580f-42ec-bf10-a19573521ed4
      ipaddress: 10.0.0.15
    - uuid: 58a803ae-a785-470e-9789-139abbfa74fb
      ipaddress: 10.0.0.16
  compute_nodes:
    - uuid: b795b3b9-c4e3-4a76-90af-258d9336d9fb
      vrouter_gateway: 10.0.0.1
    - uuid: 2d4be83e-6fcc-4761-86f2-c2615dd15074
      vrouter_gateway: 10.0.0.1
```

NOTE: Only control_nodes, compute_nodes, dpdk_nodes and sriov_nodes are supported.

Generate subcluster environment:

```
~/tripleo-heat-templates/tools/contrail/create_subcluster_environment.py -i
~/subcluster_input.yaml \
    -o
~/tripleo-heat-templates/environments/contrail/contrail-subcluster.yaml
```

Check subcluster environment file:

```
cat ~/tripleo-heat-templates/environments/contrail/contrail-subcluster.yaml
parameter_defaults:
  NodeDataLookup:
    041D7B75-6581-41B3-886E-C06847B9C87E:
      contrail_settings:
        CONTROL_NODES: 10.0.0.14,10.0.0.15,10.0.0.16
        SUBCLUSTER: subcluster2
        VROUTER_GATEWAY: 10.0.0.1
    09BEC8CB-77E9-42A6-AFF4-6D4880FD87D0:
      contrail_settings:
        BGP_ASN: '65414'
        SUBCLUSTER: subcluster2
    14639A66-D62C-4408-82EE-FDDC4E509687:
      contrail_settings:
        BGP_ASN: '65414'
        SUBCLUSTER: subcluster2
    28AB0B57-D612-431E-B177-1C578AE0FEA4:
      contrail_settings:
        BGP_ASN: '65413'
        SUBCLUSTER: subcluster1
    3993957A-ECBF-4520-9F49-0AF6EE1667A7:
      contrail_settings:
        BGP_ASN: '65413'
        SUBCLUSTER: subcluster1
    73F8D030-E896-4A95-A9F5-E1A4FEBE322D:
      contrail_settings:
        BGP_ASN: '65413'
        SUBCLUSTER: subcluster1
    7933C2D8-E61E-4752-854E-B7B18A424971:
      contrail_settings:
        CONTROL_NODES: 10.0.0.14,10.0.0.15,10.0.0.16
        SUBCLUSTER: subcluster2
        VROUTER_GATEWAY: 10.0.0.1
    AF92F485-C30C-4D0A-BDC4-C6AE97D06A66:
```

```

    contrail_settings:
      BGP_ASN: '65414'
      SUBCLUSTER: subcluster2
BB9E9D00-57D1-410B-8B19-17A0DA581044:
    contrail_settings:
      CONTROL_NODES: 10.0.0.11,10.0.0.12,10.0.0.13
      SUBCLUSTER: subcluster1
      VROUTER_GATEWAY: 10.0.0.1
E1A809DE-FDB2-4EB2-A91F-1B3F75B99510:
    contrail_settings:
      CONTROL_NODES: 10.0.0.11,10.0.0.12,10.0.0.13
      SUBCLUSTER: subcluster1
      VROUTER_GATEWAY: 10.0.0.1

```

Deployment

Add `contrail-subcluster.yaml`, `contrail-ips-from-pool-all.yaml` and `contrail-scheduler-hints.yaml` to the OpenStack deploy command:

```

openstack overcloud deploy --templates ~/tripleo-heat-templates \
-e ~/overcloud_images.yaml \
-e ~/tripleo-heat-templates/environments/network-isolation.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-plugins.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-services.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-net.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-subcluster.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-ips-from-pool-all.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-scheduler-hints.yaml \
--roles-file ~/tripleo-heat-templates/roles_data_contrail_aio.yaml

```

Installing Overcloud

1. Deployment:

```

openstack overcloud deploy --templates tripleo-heat-templates/ \
--stack overcloud --libvirt-type kvm \
--roles-file $role_file \
-e tripleo-heat-templates/environments/rhsm.yaml \
-e tripleo-heat-templates/environments/network-isolation.yaml \
-e tripleo-heat-templates/environments/contrail/contrail-services.yaml \

```

```
-e tripleo-heat-templates/environments/contrail/contrail-net.yaml \  
-e tripleo-heat-templates/environments/contrail/contrail-plugins.yaml \  
-e containers-prepare-parameter.yaml \  
-e rhsm.yaml
```

2. Validation Test:

```
source overcloudrc  
curl -O http://download.cirros-cloud.net/0.3.5/cirros-0.3.5-x86_64-disk.img  
openstack image create --container-format bare --disk-format qcow2 --file  
cirros-0.3.5-x86_64-disk.img cirros  
openstack flavor create --public cirros --id auto --ram 64 --disk 0 --vcpus 1  
openstack network create net1  
openstack subnet create --subnet-range 1.0.0.0/24 --network net1 sn1  
nova boot --image cirros --flavor cirros --nic net-id=openstack network show  
net1 -c id -f value --availability-zone nova:overcloud-novacompute-0.localdomain  
c1  
nova list
```

Release History Table

Release	Description
2008	Starting with Contrail Networking Release 2008, Contrail Networking supports using Contrail with Red Hat OpenStack Platform Director 16.1.

Setting Up Contrail with Red Hat OpenStack 13

IN THIS SECTION

- [Understanding Red Hat OpenStack Platform Director | 271](#)
- [Setting Up the Infrastructure | 275](#)
- [Setting Up the Undercloud | 284](#)
- [Setting Up the Overcloud | 286](#)
- [Using Netronome SmartNIC vRouter with Contrail Networking | 331](#)
- [Installing OpenStack Octavia LBaaS with RHOSP in Contrail Networking | 334](#)

Understanding Red Hat OpenStack Platform Director

IN THIS SECTION

- [Red Hat OpenStack Platform Director | 271](#)
- [Contrail Roles | 272](#)
- [Undercloud Requirements | 273](#)
- [Overcloud Requirements | 273](#)
- [Networking Requirements | 273](#)
- [Compatibility Matrix | 275](#)
- [Installation Summary | 275](#)

Red Hat OpenStack Platform Director

This chapter explains how to integrate a Contrail 5.1.x installation (or higher) with Red Hat OpenStack Platform Director 13.

Red Hat OpenStack Platform provides an installer called the Red Hat OpenStack Platform director (RHOSPd or OSPd), which is a toolset based on the OpenStack project TripleO (OOO, OpenStack on OpenStack). TripleO is an open source project that uses features of OpenStack to deploy a fully functional, tenant-facing OpenStack environment.

TripleO can be used to deploy an RDO-based OpenStack environment integrated with Tungsten Fabric. Red Hat OpenStack Platform director can be used to deploy an RHOSP-based OpenStack environment integrated with Contrail.

OSPd uses the concepts of undercloud and overcloud. OSPd sets up an undercloud, a single server running an operator-facing deployment that contains the OpenStack components needed to deploy and manage an overcloud, a tenant-facing deployment that hosts user workloads.

The overcloud is the deployed solution that can represent a cloud for any purpose, such as production, staging, test, and so on. The operator can select to deploy to their environment any of the available overcloud roles, such as controller, compute, and the like.

OSPd leverages existing core components of OpenStack including Nova, Ironi, Neutron, Heat, Glance, and Ceilometer to deploy OpenStack on bare metal hardware.

- Nova and Ironi are used in the undercloud to manage the bare metal instances that comprise the infrastructure for the overcloud.
- Neutron is used to provide a networking environment in which to deploy the overcloud.

- Glance stores machine images.
- Ceilometer collects metrics about the overcloud.

For more information about OSPd architecture, see [OSPd documentation](#).

Contrail Roles


OSPd supports composable roles, which are groups of services that you define through Heat templates. Composable roles allow you to integrate Contrail into the overcloud environment.

The following are the Contrail roles used for integrating into the overcloud:

- Contrail Controller
- Contrail Analytics
- Contrail Analytics Database
- Contrail-TSN
- Contrail-DPDK

Figure 43 on page 213 shows the relationship and components of an undercloud and overcloud architecture for Contrail.

Figure 46: Undercloud and Overcloud with Roles



Undercloud Requirements

The undercloud is a single server or VM that hosts the OpenStack Platform director, which is an OpenStack installation used to provision OpenStack on the overcloud.

See [Undercloud Requirements](#) for the compute requirements of the undercloud.

Overcloud Requirements

The overcloud roles can be deployed to bare metal servers or to virtual machines (VMs), but the compute nodes must be deployed to bare metal systems. Every overcloud node must support IPMI for booting up from the undercloud using PXE.

Ensure the following requirements are met for the Contrail nodes per role.

- Non-high availability: A minimum of 4 overcloud nodes are needed for control plane roles for a non-high availability deployment:
 - 1x contrail-config (includes Contrail control)
 - 1x contrail-analytics
 - 1x contrail-analytics-database
 - 1x OpenStack controller
- High availability: A minimum of 12 overcloud nodes are needed for control plane roles for a high availability deployment:
 - 3x contrail-config (includes Contrail control)
 - 3x contrail-analytics
 - 3x contrail-analytics-database
 - 3x OpenStack controller

If the control plane roles are deployed to VMs, use 3 separate physical servers and deploy one role of each kind to each physical server.

See [Overcloud Requirements](#) for the compute requirements of the overcloud.

Networking Requirements

As a minimum, the installation requires two networks:

- provisioning network - This is the private network that the undercloud uses to provision the overcloud.
- external network - This is the externally-routable network you use to access the undercloud and overcloud nodes.

Ensure the following requirements are met for the provisioning network:

- One NIC from every machine must be in the same broadcast domain of the provisioning network, and it should be the same NIC on each of the overcloud machines. For example, if you use the second NIC on the first overcloud machine, you should use the second NIC on each additional overcloud machine.

During installation, these NICs will be referenced by a single name across all overcloud machines.

- The provisioning network NIC should not be the same NIC that you are using for remote connectivity to the undercloud machine. During the undercloud installation, an Open vSwitch bridge will be created for Neutron, and the provisioning NIC will be bridged to the Open vSwitch bridge. Consequently, connectivity would be lost if the provisioning NIC was also used for remote connectivity to the undercloud machine.
- The provisioning NIC on the overcloud nodes must be untagged.
- You must have the MAC address of the NIC that will PXE boot the IPMI information for the machine on the provisioning network. The IPMI information will include such things as the IP address of the IPMI NIC and the IPMI username and password.
- All of the networks must be available to all of the Contrail roles and computes.

While the provisioning and external networks are sufficient for basic applications, you should create additional networks in most overcloud environments to provide isolation for the different traffic types by assigning network traffic to specific network interfaces or bonds.

When isolated networks are configured, the OpenStack services are configured to use the isolated networks. If no isolated networks are configured, all services run on the provisioning network. If only some isolated networks are configured, traffic belonging to a network not configured runs on the provisioning network.

The following networks are typically deployed when using network isolation topology:

- Provisioning - used by the undercloud to provision the overcloud
- Internal API - used by OpenStack services to communicate with each other
- Tenant - used for tenant overlay data plane traffic (one network per tenant)
- Storage - used for storage data traffic
- Storage Management - used for storage control and management traffic
- External - provides external access to the undercloud and overcloud, including external access to the web UIs and public APIs
- Floating IP - provides floating IP access to the tenant network (can either be merged with external or can be a separate network)
- Management - provides access for system administration

For more information on the different network types, see [Planning Networks](#).

For more information on networking requirements, see [Networking Requirements](#).

Compatibility Matrix

The following combinations of Operating System/OpenStack/Deployer/Contrail are supported:

Table 12: Compatibility Matrix

Operating System	OpenStack	Deployer	Contrail
RHEL 7.5	OSP13	OSPd13	Contrail 5.1.x or higher
CentOS 7.5	RDO queens/stable	tripleo queens/stable	Tungsten Fabric (latest)

Installation Summary

The general installation procedure is as follows:

- Set up the infrastructure, which is the set of servers or VMs that host the undercloud and overcloud, including the provisioning network that connects them together.
- Set up the undercloud, which is the OSPd application.
- Set up the overcloud, which is the set of services in the tenant-facing network. Contrail is part of the overcloud.

For more information on installing and using the RHOSPd, see [Red Hat documentation](#).

Setting Up the Infrastructure

IN THIS SECTION

- [Target Configuration \(Example\) | 275](#)
- [Configure the External Physical Switch | 277](#)
- [Configure KVM Hosts | 278](#)
- [Create the Overcloud VM Definitions on the Overcloud KVM Hosts | 280](#)
- [Create the Undercloud VM Definition on the Undercloud KVM Host | 282](#)

Target Configuration (Example)

Undercloud and overcloud KVM hosts require virtual switches and virtual machine definitions to be configured. You can deploy any KVM host operating system version that supports KVM and OVS. The following example shows a RHEL/CentOS based system. If you are using RHEL, you must subscribe the system.

The following example illustrates all control plane functions as Virtual Machines hosted on KVM hosts.


There are different ways to create the infrastructure providing the control plane elements. To illustrate the installation procedure, we will use four host machines for the infrastructure, each running KVM. KVM1 contains a VM running the undercloud while KVM2 through KVM4 each contains a VM running an OpenStack controller and a Contrail controller ([Table 7 on page 216](#)).

Table 13: Control Plane Infrastructure

KVM Host	Virtual Machines
KVM1	undercloud
KVM2	OpenStack Controller 1, Contrail Controller 1
KVM3	OpenStack Controller 2, Contrail Controller 2
KVM4	OpenStack Controller 3, Contrail Controller 3


[Figure 44 on page 217](#) shows the physical connectivity where each KVM host and each compute node has two interfaces that connect to an external switch. These interfaces attach to separate virtual bridges within the VM, allowing for two physically separate networks (external and provisioning networks).

Figure 47: Physical View



[Figure 45 on page 218](#) shows the logical view of the connectivity where VLANs are used to provide further network separation for the different OpenStack network types.

Figure 48: Logical View



The following sections describe how to configure the infrastructure, the undercloud, and finally the overcloud.

Configure the External Physical Switch

Configure the ports and VLANs on the external physical switch according to the following table:

Table 14: External Physical Switch Port and VLAN Configuration

Port	Trunked VLAN	Native VLAN
ge0	-	-
ge1	700, 720	-

Table 14: External Physical Switch Port and VLAN Configuration (continued)

Port	Trunked VLAN	Native VLAN
ge2	700, 710, 720, 730, 740, 750	-
ge3	-	-
ge4	710, 730	700
ge5	-	-

Configure KVM Hosts

Use this example procedure to install the required packages and start KVM and Open vSwitch on each undercloud and overcloud KVM host.

1. Log in to a KVM host.
2. Install the required packages.

```
yum install -y libguestfs \
  libguestfs-tools \
  openvswitch \
  virt-install \
  kvm libvirt \
  libvirt-python \
  python-virtualbmc \
  python-virtinst
```

3. Start KVM and Open vSwitch.

```
systemctl start libvirtd
systemctl start openvswitch
```

4. Additionally, on the overcloud nodes only, create and start the virtual switches br0 and br1.

Table 15: vSwitch Configuration

Bridge	Trunked VLAN	Native VLAN
br0	710, 720, 730 740, 750	700
br1	-	-

```

# Create the virtual switches and bind them to the respective interfaces.
ovs-vsctl add-br br0
ovs-vsctl add-br br1
ovs-vsctl add-port br0 NIC1
ovs-vsctl add-port br1 NIC2

# Create the configuration file for br0.
cat << EOF > br0.xml
<network>
  <name>br0</name>
  <forward mode='bridge'/>
  <bridge name='br0'/>
  <virtualport type='openvswitch'/>
  <portgroup name='overcloud'/>
    <vlan trunk='yes'>
      <tag id='700' nativeMode='untagged'/>
      <tag id='710'/>
      <tag id='720'/>
      <tag id='730'/>
      <tag id='740'/>
      <tag id='750'/>
    </vlan>
  </portgroup>
</network>
EOF

# Create the configuration file for br1.
cat << EOF > br1.xml
<network>
  <name>br1</name>
  <forward mode='bridge'/>
  <bridge name='br1'/>
  <virtualport type='openvswitch'/>
</network>
EOF

# Create the br0 network based on the configuration file.
virsh net-define br0.xml
virsh net-start br0
virsh net-autostart br0

# Create the br1 network based on the configuration file.
virsh net-define br1.xml
virsh net-start br1

```

```
virsh net-autostart br1
```

5. Repeat step 1 through step 4 for each KVM host.

Create the Overcloud VM Definitions on the Overcloud KVM Hosts

Use this example procedure on each overcloud KVM host (KVM2 to KVM4) to do the following:

- create the VM definitions for that overcloud KVM host
- create and start a virtual baseboard management controller for that overcloud KVM host so that the VM can be managed using IPMI
- create an **ironic_list** file to be used by the undercloud

This example procedure creates a VM definition consisting of 2 compute nodes, 1 Contrail controller node, and 1 OpenStack controller node on each overcloud KVM host.

1. Log in to an overcloud KVM host.
2. Specify the roles you want to create.

```
ROLES=compute:2,contrail-controller:1,control:1
```

3. Create the VM definitions.

```
# Initialize and specify the IPMI user and password you want to use.
num=0
ipmi_user=<user>
ipmi_password=<password>
libvirt_path=/var/lib/libvirt/images
port_group=overcloud
prov_switch=br0
/bin/rm ironic_list

# For each role and instance specified in the ROLES variable:
#   - create the VM definition
#   - create and start a virtual baseboard management controller (vbmc)
#   - store the VM information into an ironic_list file (for later use in the
undercloud)
IFS=' ' read -ra role_list <<< "${ROLES}"
for role in ${role_list[@]}; do
    role_name=`echo $role|cut -d ":" -f 1`
    role_count=`echo $role|cut -d ":" -f 2`
```



```

for count in `seq 1 ${role_count}`; do
    echo $role_name $count
    qemu-img create -f qcow2 ${libvirt_path}/${role_name}_${count}.qcow2 99G
    virsh define /dev/stdin <<EOF
$(virt-install --name ${role_name}_${count} \
    --disk ${libvirt_path}/${role_name}_${count}.qcow2 \
    --vcpus=4 \
    --ram=16348 \
    --network network=br0,model=virtio,portgroup=${port_group} \
    --network network=br1,model=virtio \
    --virt-type kvm \
    --cpu host \
    --import \
    --os-variant rhel7 \
    --serial pty \
    --console pty,target_type=virtio \
    --graphics vnc \
    --print-xml)
EOF
    vbmc add ${role_name}_${count} --port 1623${num} --username ${ipmi_user}
--password ${ipmi_password}
    vbmc start ${role_name}_${count}
    prov_mac=`virsh domiflist ${role_name}_${count}|grep ${prov_switch}|awk
'{print $5}'`
    vm_name=${role_name}-${count}-`hostname -s`
    kvm_ip=`ip route get 1 |grep src |awk '{print $7}'`
    echo ${prov_mac} ${vm_name} ${kvm_ip} ${role_name} 1623${num}>> ironic_list
    num=$(expr $num + 1)
done
done

```

4. Repeat step 1 through step 3 on each overcloud KVM host.



CAUTION: This procedure creates one **ironic_list** file per overcloud KVM host. Combine the contents of each file into a single **ironic_list** file on the undercloud.

The following shows the resulting **ironic_list** file after you combine the contents from each separate file:

```
52:54:00:e7:ca:9a compute-1-5b3s31 10.87.64.32 compute 16230
52:54:00:30:6c:3f compute-2-5b3s31 10.87.64.32 compute 16231
52:54:00:9a:0c:d5 contrail-controller-1-5b3s31 10.87.64.32 contrail-controller 16232
52:54:00:cc:93:d4 control-1-5b3s31 10.87.64.32 control 16233
52:54:00:28:10:d4 compute-1-5b3s30 10.87.64.31 compute 16230
52:54:00:7f:36:e7 compute-2-5b3s30 10.87.64.31 compute 16231
52:54:00:32:e5:3e contrail-controller-1-5b3s30 10.87.64.31 contrail-controller 16232
52:54:00:d4:31:aa control-1-5b3s30 10.87.64.31 control 16233
52:54:00:d1:d2:ab compute-1-5b3s32 10.87.64.33 compute 16230
52:54:00:ad:a7:cc compute-2-5b3s32 10.87.64.33 compute 16231
52:54:00:55:56:50 contrail-controller-1-5b3s32 10.87.64.33 contrail-controller 16232
52:54:00:91:51:35 control-1-5b3s32 10.87.64.33 control 16233
```

Create the Undercloud VM Definition on the Undercloud KVM Host

Use this example procedure on the undercloud KVM host (KVM1) to create the undercloud VM definition and to start the undercloud VM.

1. Create the images directory.

```
mkdir ~/images
cd images
```

2. Retrieve the image.

- CentOS

```
curl
https://cloud.centos.org/centos/7/images/CentOS-7-x86_64-GenericCloud-1802.qcow2.xz
-o CentOS-7-x86_64-GenericCloud-1802.qcow2.xz
unxz -d images/CentOS-7-x86_64-GenericCloud-1802.qcow2.xz
cloud_image=~/images/CentOS-7-x86_64-GenericCloud-1802.qcow2
```

- RHEL

```
Download rhel-server-7.5-update-1-x86_64-kvm.qcow2 from the Red Hat portal to
~/images.
cloud_image=~/.images/rhel-server-7.5-update-1-x86_64-kvm.qcow2
```

3. Customize the undercloud image.

```
undercloud_name=queensa
undercloud_suffix=local
root_password=<password>
stack_password=<password>
export LIBGUESTFS_BACKEND=direct
qemu-img create -f qcow2 /var/lib/libvirt/images/${undercloud_name}.qcow2 100G
virt-resize --expand /dev/sda1 ${cloud_image}
/var/lib/libvirt/images/${undercloud_name}.qcow2
virt-customize -a /var/lib/libvirt/images/${undercloud_name}.qcow2 \
--run-command 'xfs_growfs /' \
--root-password password:${root_password} \
--hostname ${undercloud_name}.${undercloud_suffix} \
--run-command 'useradd stack' \
--password stack:password:${stack_password} \
--run-command 'echo "stack ALL=(root) NOPASSWD:ALL" | tee -a /etc/sudoers.d/stack' \
--chmod 0440:/etc/sudoers.d/stack \
--run-command 'sed -i "s/PasswordAuthentication no/PasswordAuthentication yes/g"
/etc/ssh/sshd_config' \
--run-command 'systemctl enable sshd' \
--run-command 'yum remove -y cloud-init' \
--selinux-relabel
```

NOTE: As part of the undercloud definition, a user called **stack** is created. This user will be used later to install the undercloud.

4. Define the undercloud virsh template.

```
vcpus=8
vram=32000
virt-install --name ${undercloud_name} \
--disk /var/lib/libvirt/images/${undercloud_name}.qcow2 \
--vcpus=${vcpus} \
--ram=${vram} \
--network network=default,model=virtio \
```

```
--network network=br0,model=virtio,portgroup=overcloud \
--virt-type kvm \
--import \
--os-variant rhel7 \
--graphics vnc \
--serial pty \
--noautoconsole \
--console pty,target_type=virtio
```

5. Start the undercloud VM.

```
virsh start ${undercloud_name}
```

6. Retrieve the undercloud IP address. It might take several seconds before the IP address is available.

```
undercloud_ip=`virsh domifaddr ${undercloud_name} |grep ipv4 |awk '{print $4}'
|awk -F"/" '{print $1}'` ssh-copy-id ${undercloud_ip}
```

Setting Up the Undercloud

IN THIS SECTION

- [Install the Undercloud | 284](#)
- [Perform Post-Install Configuration | 286](#)

Install the Undercloud

Use this example procedure to install the undercloud.

1. Log in to the undercloud VM from the undercloud KVM host.

```
ssh ${undercloud_ip}
```

2. Configure the hostname.

```
undercloud_name=`hostname -s`
```

```
undercloud_suffix=`hostname -d`
hostnamectl set-hostname ${undercloud_name}.${undercloud_suffix}
hostnamectl set-hostname --transient ${undercloud_name}.${undercloud_suffix}
```

3. Add the hostname to the `/etc/hosts` file. The following example assumes the management interface is `eth0`.

```
undercloud_ip=`ip addr sh dev eth0 | grep "inet " | awk '{print $2}' | awk -F"/" '{print $1}'`
echo ${undercloud_ip} ${undercloud_name}.${undercloud_suffix} ${undercloud_name}
>> /etc/hosts
```

4. Set up the repositories.

- CentOS

```
tripleo_repos=`python -c 'import requests;r =
requests.get("https://trunk.rdoproject.org/centos7-queens/current"); print r.text
' | grep python2-tripleo-repos|awk -F"href=\"\" " '{print $2}' | awk -F"\"\" " '{print
$1}'`
yum install -y
https://trunk.rdoproject.org/centos7-queens/current/${tripleo_repos}
tripleo-repos -b queens current
```

- RHEL

```
#Register with Satellite (can be done with CDN as well)
satellite_fqdn=device.example.net
act_key=xxx
org=example
yum localinstall -y
http://${satellite_fqdn}/pub/katello-ca-consumer-latest.noarch.rpm
subscription-manager register --activationkey=${act_key} --org=${org}
```

5. Install the Tripleo client.

```
yum install -y python-tripleoclient tmux
```

6. Copy the undercloud configuration file sample and modify the configuration as required. See [Red Hat documentation](#) for information on how to modify that file.

```
su - stack
cp /usr/share/instack-undercloud/undercloud.conf.sample ~/undercloud.conf
vi ~/undercloud.conf
```

7. Install the undercloud.

```
openstack undercloud install
source stackrc
```

Perform Post-Install Configuration

1. Configure a forwarding path between the provisioning network and the external network:

```
sudo iptables -A FORWARD -i br-ctlplane -o eth0 -j ACCEPT
sudo iptables -A FORWARD -i eth0 -o br-ctlplane -m state --state
RELATED,ESTABLISHED -j ACCEPT
sudo iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

2. Add the external API interface:

```
sudo ip link add name vlan720 link br-ctlplane type vlan id 720
sudo ip addr add 10.2.0.254/24 dev vlan720
sudo ip link set dev vlan720 up
```

3. Add the **stack** user to the docker group:

```
newgrp docker
exit
su - stack
source stackrc
```

Setting Up the Overcloud

IN THIS SECTION

- [Configuring the Overcloud | 287](#)
- [Customizing the Contrail Service with Templates \(contrail-services.yaml\) | 293](#)

- Customizing the Contrail Network with Templates | 294
- Installing Overcloud | 329

Configuring the Overcloud

Use this example procedure on the undercloud to set up the configuration for the overcloud.

1. Specify the name server to be used:

```
undercloud_nameserver=8.8.8.8
openstack subnet set `openstack subnet show ctlplane-subnet -c id -f value`
--dns-nameserver ${undercloud_nameserver}
```

2. Retrieve and upload the overcloud images.

- a. Create the image directory:

```
mkdir images
cd images
```

- b. Retrieve the overcloud images from either the RDO project or from Red Hat.

- TripleO

```
curl -O
https://images.rdoproject.org/queens/rdo_trunk/current-tripleo-rdo/ironic-python-agent.tar

curl -O
https://images.rdoproject.org/queens/rdo_trunk/current-tripleo-rdo/overcloud-full.tar

tar xvf ironic-python-agent.tar
tar xvf overcloud-full.tar
```

- OSP13

```
sudo yum install -y rhosp-director-images rhosp-director-images-ipa
for i in /usr/share/rhosp-director-images/overcloud-full-latest-13.0.tar
/usr/share/rhosp-director-images/ironic-python-agent-latest-13.0.tar ; do tar
-xvf $i; done
```

- c. Upload the overcloud images:

```
cd
openstack overcloud image upload --image-path /home/stack/images/
```

3. Prepare OpenStack's bare metal provisioning (Ironic).

Ironic is an integrated OpenStack program that provisions bare metal machines instead of virtual machines. It is best thought of as a bare metal hypervisor API and a set of plugins that interact with the bare metal hypervisors.

NOTE: Make sure to combine the **ironic_list** files from the three overcloud KVM hosts.

a. Add the overcloud VMs to Ironic:

```
ipmi_password=<password>
ipmi_user=<user>
while IFS= read -r line; do
    mac=`echo $line|awk '{print $1}'`
    name=`echo $line|awk '{print $2}'`
    kvm_ip=`echo $line|awk '{print $3}'`
    profile=`echo $line|awk '{print $4}'`
    ipmi_port=`echo $line|awk '{print $5}'`
    uuid=`openstack baremetal node create --driver ipmi \
                                           --property cpus=4 \
                                           --property memory_mb=16348 \
                                           --property local_gb=100 \
                                           --property cpu_arch=x86_64 \
                                           --driver-info

ipmi_username=${ipmi_user} \
                                           --driver-info ipmi_address=${kvm_ip}
\
                                           --driver-info
ipmi_password=${ipmi_password} \
                                           --driver-info ipmi_port=${ipmi_port}
\
                                           --name=${name} \
                                           --property
capabilities=profile:${profile},boot_option:local \
                                           -c uuid -f value`
    openstack baremetal port create --node ${uuid} ${mac}
done < <(cat ironic_list)

DEPLOY_KERNEL=$(openstack image show bm-deploy-kernel -f value -c id)
```



```

DEPLOY_RAMDISK=$(openstack image show bm-deploy-ramdisk -f value -c id)

for i in `openstack baremetal node list -c UUID -f value`; do
    openstack baremetal node set $i --driver-info deploy_kernel=$DEPLOY_KERNEL
    --driver-info deploy_ramdisk=$DEPLOY_RAMDISK
done

for i in `openstack baremetal node list -c UUID -f value`; do
    openstack baremetal node show $i -c properties -f value
done

```

b. Introspect the overcloud node:

```

for node in $(openstack baremetal node list -c UUID -f value) ; do
    openstack baremetal node manage $node
done
openstack overcloud node introspect --all-manageable --provide

```

c. Add Baremetal Server (BMS) to Ironic.

- Create rules for automated profiling.

Evaluate the attributes of the physical server. The server will automatically be profiled based on the rules.

The following example shows how to create a rule for system manufacturer as “Supermicro” and memory greater or equal to 128 GB.

```

cat << EOF > ~/rule_compute.json
[
  {
    "description": "set physical compute",
    "conditions": [
      {"op": "eq", "field": "data://auto_discovered", "value": true},
      {"op": "eq", "field": "data://inventory.system_vendor.manufacturer",
        "value": "Supermicro"},
      {"op": "ge", "field": "memory_mb", "value": 128000}
    ],
    "actions": [
      {"action": "set-attribute", "path": "driver_info/ipmi_username",
        "value": "<user>"},
      {"action": "set-attribute", "path": "driver_info/ipmi_password",
        "value": "<password>"}
    ]
  }
]

```

```

        {"action": "set-capability", "name": "profile", "value": "compute"},

        {"action": "set-attribute", "path":
"driver_info/ipmi_address", "value": "{data[inventory][bmc_address]}" }
    ]
}
]
EOF

```

You can import the rule by:

```
openstack baremetal introspection rule import ~/rule_compute.json
```

- Scan the BMC IP range and automatically add new servers matching the above rule by:

```

ipmi_range=10.87.122.25/32
ipmi_password=<password>
ipmi_user=<user>
openstack overcloud node discover --range ${ipmi_range} \
--credentials ${ipmi_user}:${ipmi_password} \
--introspect --provide

```

4. Create Flavor:

```

for i in compute-dpdk \
compute-sriov \
contrail-controller \
contrail-analytics \
contrail-database \
contrail-analytics-database; do
    openstack flavor create $i --ram 4096 --vcpus 1 --disk 40
    openstack flavor set --property "capabilities:boot_option"="local" \
        --property "capabilities:profile"="${i}" $i
done

```

5. Copy the TripleO heat templates.

```
cp -r /usr/share/openstack-tripleo-heat-templates/ tripleo-heat-templates
```

6. Download and copy the Contrail heat templates from <https://support.juniper.net/support/downloads>.

```
tar -xzvf contrail-tripleo-heat-templates-<version>.tgz
cp -r contrail-tripleo-heat-templates/* tripleo-heat-templates/
```

7. Create and upload the OpenStack containers.

a. Create the OpenStack container file.

NOTE: The container must be created based on the OpenStack program.

• TripleO

```
openstack overcloud container image prepare \
  --namespace docker.io/tripleoqueens \
  --tag current-tripleo \
  --tag-from-label rdo_version \
  --output-env-file=~/.overcloud_images.yaml

tag=`grep "docker.io/tripleoqueens" docker_registry.yaml |tail -1 |awk -F":" '{print $3}'`

openstack overcloud container image prepare \
  --namespace docker.io/tripleoqueens \
  --tag ${tag} \
  --push-destination 192.168.24.1:8787 \
  --output-env-file=~/.overcloud_images.yaml \
  --output-images-file=~/.local_registry_images.yaml
```

• OSP13

```
openstack overcloud container image prepare \
  --push-destination=192.168.24.1:8787 \
  --tag-from-label {version}-{release} \
  --output-images-file ~/.local_registry_images.yaml \
  --namespace=registry.access.Red Hat.com/rhosp13 \
  --prefix=openstack- \
  --tag-from-label {version}-{release} \
  --output-env-file ~/.overcloud_images.yaml
```

b. Upload the OpenStack containers:

```
openstack overcloud container image upload --config-file
~/local_registry_images.yaml
```

8. Create and upload the Contrail containers.

a. Create the Contrail container file.

NOTE: This step is optional. The Contrail containers can be downloaded from external registries later.

```
cd ~/tripleo-heat-templates/tools/contrail
./import_contrail_container.sh -f container_outputfile -r registry -t tag [-i
insecure] [-u username] [-p password] [-c certificate pat
```

Here are few examples of importing Contrail containers from different sources:

- Import from password protected public registry:

```
./import_contrail_container.sh -f /tmp/contrail_container -r
hub.juniper.net/contrail -u USERNAME -p PASSWORD -t 1234
```

- Import from Dockerhub:

```
./import_contrail_container.sh -f /tmp/contrail_container -r
docker.io/opencontrailnightly -t 1234
```

- Import from private secure registry:

```
./import_contrail_container.sh -f /tmp/contrail_container -r
device.example.net:5443 -c
http://device.example.net/pub/device.example.net.crt -t 1234
```

- Import from private insecure registry:

```
./import_contrail_container.sh -f /tmp/contrail_container -r 10.0.0.1:5443
-i 1 -t 1234
```

- b. Upload Contrail containers to the undercloud registry:

```
openstack overcloud container image upload --config-file
/tmp/contrail_container
```

Customizing the Contrail Service with Templates (contrail-services.yaml)

This section contains information to customize Contrail services for your network by modifying the **contrail-services.yaml** file.

- Contrail Services customization

```
vi ~/tripleo-heat-templates/environments/contrail-services.yaml
```

```
parameter_defaults:
  ContrailSettings:
    VRROUTER_GATEWAY: 10.0.0.1
    # KEY1: value1
    # KEY2: value2

    VXLAN_VN_ID_MODE: "configured"
    ENCAP_PRIORITY: "VXLAN,MPLSoUDP,MPLSoGRE"

  ContrailControllerParameters:
    AAAMode: rbac
```

- Contrail registry settings

```
vi ~/tripleo-heat-templates/environments/contrail-services.yaml
```

Here are few examples of default values for various registries:

- Public Juniper registry

```
parameter_defaults:
  ContrailRegistry: hub.juniper.net/contrail
  ContrailRegistryUser: <USER>
  ContrailRegistryPassword: <PASSWORD>
```

- Insecure registry

```
parameter_defaults:
  ContrailRegistryInsecure: true
  DockerInsecureRegistryAddress: 10.87.64.32:5000,192.168.24.1:8787
  ContrailRegistry: 10.87.64.32:5000
```

- Private secure registry

```
parameter_defaults:
  ContrailRegistryCertUrl: http://device.example.net/pub/device.example.net.crt

  ContrailRegistry: device.example.net:5443
```

- Contrail Container image settings

```
parameter_defaults:
  ContrailImageTag: queens-5.0-104-rhel-queens
```

Customizing the Contrail Network with Templates

IN THIS SECTION

- [Overview | 294](#)
- [Roles Configuration \(roles_data_contrail_aio.yaml\) | 295](#)
- [Network Parameter Configuration \(contrail-net.yaml\) | 298](#)
- [Network Interface Configuration \(*-NIC-*.yaml\) | 299](#)
- [Advanced vRouter Kernel Mode Configuration | 310](#)
- [Advanced vRouter DPDK Mode Configuration | 312](#)
- [Advanced vRouter SRIOV + Kernel Mode Configuration | 315](#)
- [Advanced vRouter SRIOV + DPDK Mode Configuration | 318](#)
- [Advanced Scenarios | 321](#)

Overview

In order to customize the network, define different networks and configure the overcloud nodes NIC layout. TripleO supports a flexible way of customizing the network.

The following networking customization example uses network as:

Table 16: Network Customization

Network	VLAN	overcloud Nodes
provisioning	-	All
internal_api	710	All
external_api	720	OpenStack CTRL
storage	740	OpenStack CTRL, Computes
storage_mgmt	750	OpenStack CTRL
tenant	-	Contrail CTRL, Computes

Roles Configuration (roles_data_contrail_aio.yaml)**IN THIS SECTION**

- [OpenStack Controller | 295](#)
- [Compute Node | 296](#)
- [Contrail Controller | 296](#)
- [Compute DPDK | 297](#)
- [Compute SRIOV | 297](#)
- [Compute CSN | 297](#)

The networks must be activated per role in the roles_data file:

```
vi ~/tripleo-heat-templates/roles_data_contrail_aio.yaml
```

OpenStack Controller

```
#####
# Role: Controller                                     #
#####
- name: Controller
  description: |
    Controller role that has all the controller services loaded and handles
```

```

    Database, Messaging and Network functions.
CountDefault: 1
tags:
  - primary
  - controller
networks:
  - External
  - InternalApi
  - Storage
  - StorageMgmt

```

Compute Node

```

#####
# Role: Compute                                                                    #
#####
- name: Compute
  description: |
    Basic Compute Node role
CountDefault: 1
networks:
  - InternalApi
  - Tenant
  - Storage

```

Contrail Controller

```

#####
# Role: ContrailController                                                            #
#####
- name: ContrailController
  description: |
    ContrailController role that has all the Contrail controller services loaded
    and handles config, control and webui functions
CountDefault: 1
tags:
  - primary
  - contrailcontroller
networks:
  - InternalApi
  - Tenant

```


Compute DPDK

```
#####
# Role: ContrailDpdk                                     #
#####
- name: ContrailDpdk
  description: |
    Contrail Dpdk Node role
  CountDefault: 0
  tags:
    - contraildpdk
  networks:
    - InternalApi
    - Tenant
    - Storage
```

Compute SRIOV

```
#####
# Role: ContrailSriov                                     #
#####
- name: ContrailSriov
  description: |
    Contrail Sriov Node role
  CountDefault: 0
  tags:
    - contrailsriov
  networks:
    - InternalApi
    - Tenant
    - Storage
```

Compute CSN

```
#####
# Role: ContrailTsn                                       #
#####
- name: ContrailTsn
  description: |
    Contrail Tsn Node role
  CountDefault: 0
  tags:
```

```

- contrailtsn
networks:
- InternalApi
- Tenant
- Storage

```

Network Parameter Configuration (contrail-net.yaml)

```
cat ~/tripleo-heat-templates/environments/contrail/contrail-net.yaml
```

```

resource_registry:
  OS::TripleO::Controller::Net::SoftwareConfig:
  ../../network/config/contrail/controller-nic-config.yaml
  OS::TripleO::ContrailController::Net::SoftwareConfig:
  ../../network/config/contrail/contrail-controller-nic-config.yaml
  OS::TripleO::ContrailControlOnly::Net::SoftwareConfig:
  ../../network/config/contrail/contrail-controller-nic-config.yaml
  OS::TripleO::Compute::Net::SoftwareConfig:
  ../../network/config/contrail/compute-nic-config.yaml
  OS::TripleO::ContrailDpdk::Net::SoftwareConfig:
  ../../network/config/contrail/contrail-dpdk-nic-config.yaml
  OS::TripleO::ContrailSriov::Net::SoftwareConfig:
  ../../network/config/contrail/contrail-sriov-nic-config.yaml
  OS::TripleO::ContrailTsn::Net::SoftwareConfig:
  ../../network/config/contrail/contrail-tsn-nic-config.yaml

```

```

parameter_defaults:
  # Customize all these values to match the local environment
  TenantNetCidr: 10.0.0.0/24
  InternalApiNetCidr: 10.1.0.0/24
  ExternalNetCidr: 10.2.0.0/24
  StorageNetCidr: 10.3.0.0/24
  StorageMgmtNetCidr: 10.4.0.0/24
  # CIDR subnet mask length for provisioning network
  ControlPlaneSubnetCidr: '24'
  # Allocation pools
  TenantAllocationPools: [{'start': '10.0.0.10', 'end': '10.0.0.200'}]
  InternalApiAllocationPools: [{'start': '10.1.0.10', 'end': '10.1.0.200'}]
  ExternalAllocationPools: [{'start': '10.2.0.10', 'end': '10.2.0.200'}]
  StorageAllocationPools: [{'start': '10.3.0.10', 'end': '10.3.0.200'}]
  StorageMgmtAllocationPools: [{'start': '10.4.0.10', 'end': '10.4.0.200'}]

```

```
# Routes
ControlPlaneDefaultRoute: 192.168.24.1
InternalApiDefaultRoute: 10.1.0.1
ExternalInterfaceDefaultRoute: 10.2.0.1
# Vlan
InternalApiNetworkVlanID: 710
ExternalNetworkVlanID: 720
StorageNetworkVlanID: 730
StorageMgmtNetworkVlanID: 740
TenantNetworkVlanID: 3211
# Services
EC2MetadataIp: 192.168.24.1 # Generally the IP of the undercloud
DnsServers: ["172.x.x.x"]
NtpServer: 10.0.0.1
```

Network Interface Configuration (*-NIC-*.yaml)

IN THIS SECTION

- [OpenStack Controller | 299](#)
- [Contrail Controller | 303](#)
- [Compute Node | 306](#)

NIC configuration files exist per role in the following directory:

```
cd ~/tripleo-heat-templates/network/config/contrail
```

OpenStack Controller

```
heat_template_version: queens

description: >
  Software Config to drive os-net-config to configure multiple interfaces
  for the compute role. This is an example for a Nova compute node using
  Contrail vrouter and the vhost0 interface.
```

```

parameters:
  ControlPlaneIp:
    default: ''
    description: IP address/subnet on the ctlplane network
    type: string
  ExternalIpSubnet:
    default: ''
    description: IP address/subnet on the external network
    type: string
  InternalApiIpSubnet:
    default: ''
    description: IP address/subnet on the internal_api network
    type: string
  InternalApiDefaultRoute: # Not used by default in this template
    default: '10.0.0.1'
    description: The default route of the internal api network.
    type: string
  StorageIpSubnet:
    default: ''
    description: IP address/subnet on the storage network
    type: string
  StorageMgmtIpSubnet:
    default: ''
    description: IP address/subnet on the storage_mgmt network
    type: string
  TenantIpSubnet:
    default: ''
    description: IP address/subnet on the tenant network
    type: string
  ManagementIpSubnet: # Only populated when including
environments/network-management.yaml
    default: ''
    description: IP address/subnet on the management network
    type: string
  ExternalNetworkVlanID:
    default: 10
    description: Vlan ID for the external network traffic.
    type: number
  InternalApiNetworkVlanID:
    default: 20
    description: Vlan ID for the internal_api network traffic.
    type: number
  StorageNetworkVlanID:
    default: 30

```

```

    description: Vlan ID for the storage network traffic.
    type: number
StorageMgmtNetworkVlanID:
    default: 40
    description: Vlan ID for the storage mgmt network traffic.
    type: number
TenantNetworkVlanID:
    default: 50
    description: Vlan ID for the tenant network traffic.
    type: number
ManagementNetworkVlanID:
    default: 60
    description: Vlan ID for the management network traffic.
    type: number
ControlPlaneSubnetCidr: # Override this via parameter_defaults
    default: '24'
    description: The subnet CIDR of the control plane network.
    type: string
ControlPlaneDefaultRoute: # Override this via parameter_defaults
    description: The default route of the control plane network.
    type: string
ExternalInterfaceDefaultRoute: # Not used by default in this template
    default: '10.0.0.1'
    description: The default route of the external network.
    type: string
ManagementInterfaceDefaultRoute: # Commented out by default in this template
    default: unset
    description: The default route of the management network.
    type: string
DnsServers: # Override this via parameter_defaults
    default: []
    description: A list of DNS servers (2 max for some implementations) that will
be added to resolv.conf.
    type: comma_delimited_list
EC2MetadataIp: # Override this via parameter_defaults
    description: The IP address of the EC2 metadata server.
    type: string

```

```

resources:
  OsNetConfigImpl:
    type: OS::Heat::SoftwareConfig
    properties:
      group: script

```

```

config:
  str_replace:
    template:
      get_file: ../../scripts/run-os-net-config.sh
    params:
      $network_config:
        network_config:
          - type: interface
            name: nic1
            use_dhcp: false
            dns_servers:
              get_param: DnsServers
            addresses:
          - ip_netmask:
              list_join:
                - '/'
                - - get_param: ControlPlaneIp
                  - get_param: ControlPlaneSubnetCidr
            routes:
          - ip_netmask: 169.x.x.x/32
            next_hop:
              get_param: EC2MetadataIp
          - default: true
            next_hop:
              get_param: ControlPlaneDefaultRoute
          - type: vlan
            vlan_id:
              get_param: InternalApiNetworkVlanID
            device: nic1
            addresses:
          - ip_netmask:
              get_param: InternalApiIpSubnet
          - type: vlan
            vlan_id:
              get_param: ExternalNetworkVlanID
            device: nic1
            addresses:
          - ip_netmask:
              get_param: ExternalIpSubnet
          - type: vlan
            vlan_id:
              get_param: StorageNetworkVlanID
            device: nic1
            addresses:

```

```

        - ip_netmask:
            get_param: StorageIpSubnet
    - type: vlan
      vlan_id:
        get_param: StorageMgmtNetworkVlanID
      device: nic1
      addresses:
        - ip_netmask:
            get_param: StorageMgmtIpSubnet

```

```

outputs:
  OS::stack_id:
    description: The OsNetConfigImpl resource.
    value:
      get_resource: OsNetConfigImpl

```

Contrail Controller

```
heat_template_version: queens
```

```

description: >
  Software Config to drive os-net-config to configure multiple interfaces
  for the compute role. This is an example for a Nova compute node using
  Contrail vrouter and the vhost0 interface.

```

```

parameters:
  ControlPlaneIp:
    default: ''
    description: IP address/subnet on the ctlplane network
    type: string
  ExternalIpSubnet:
    default: ''
    description: IP address/subnet on the external network
    type: string
  InternalApiIpSubnet:
    default: ''
    description: IP address/subnet on the internal_api network
    type: string
  InternalApiDefaultRoute: # Not used by default in this template
    default: '10.0.0.1'

```

```

    description: The default route of the internal api network.
    type: string
StorageIpSubnet:
    default: ''
    description: IP address/subnet on the storage network
    type: string
StorageMgmtIpSubnet:
    default: ''
    description: IP address/subnet on the storage_mgmt network
    type: string
TenantIpSubnet:
    default: ''
    description: IP address/subnet on the tenant network
    type: string
ManagementIpSubnet: # Only populated when including
environments/network-management.yaml
    default: ''
    description: IP address/subnet on the management network
    type: string
ExternalNetworkVlanID:
    default: 10
    description: Vlan ID for the external network traffic.
    type: number
InternalApiNetworkVlanID:
    default: 20
    description: Vlan ID for the internal_api network traffic.
    type: number
StorageNetworkVlanID:
    default: 30
    description: Vlan ID for the storage network traffic.
    type: number
StorageMgmtNetworkVlanID:
    default: 40
    description: Vlan ID for the storage mgmt network traffic.
    type: number
TenantNetworkVlanID:
    default: 50
    description: Vlan ID for the tenant network traffic.
    type: number
ManagementNetworkVlanID:
    default: 60
    description: Vlan ID for the management network traffic.
    type: number
ControlPlaneSubnetCidr: # Override this via parameter_defaults

```



```

    default: '24'
    description: The subnet CIDR of the control plane network.
    type: string
ControlPlaneDefaultRoute: # Override this via parameter_defaults
    description: The default route of the control plane network.
    type: string
ExternalInterfaceDefaultRoute: # Not used by default in this template
    default: '10.0.0.1'
    description: The default route of the external network.
    type: string
ManagementInterfaceDefaultRoute: # Commented out by default in this template
    default: unset
    description: The default route of the management network.
    type: string
DnsServers: # Override this via parameter_defaults
    default: []
    description: A list of DNS servers (2 max for some implementations) that will
be added to resolv.conf.
    type: comma_delimited_list
EC2MetadataIp: # Override this via parameter_defaults
    description: The IP address of the EC2 metadata server.
    type: string

```

```

resources:
  OsNetConfigImpl:
    type: OS::Heat::SoftwareConfig
    properties:
      group: script
      config:
        str_replace:
          template:
            get_file: ../../scripts/run-os-net-config.sh
        params:
          $network_config:
            network_config:
              - type: interface
                name: nic1
                use_dhcp: false
                dns_servers:
                  get_param: DnsServers
                addresses:
                  - ip_netmask:
                      list_join:

```

```

        - '/'
        - - get_param: ControlPlaneIp
          - get_param: ControlPlaneSubnetCidr
    routes:
    - ip_netmask: 169.x.x.x/32
      next_hop:
        get_param: EC2MetadataIp
    - default: true
      next_hop:
        get_param: ControlPlaneDefaultRoute
    - type: vlan
      vlan_id:
        get_param: InternalApiNetworkVlanID
      device: nic1
      addresses:
      - ip_netmask:
          get_param: InternalApiIpSubnet
    - type: interface
      name: nic2
      use_dhcp: false
      addresses:
      - ip_netmask:
          get_param: TenantIpSubnet

```

```

outputs:
  OS::stack_id:
    description: The OsNetConfigImpl resource.
    value:
      get_resource: OsNetConfigImpl

```

Compute Node

```
heat_template_version: queens
```

```

description: >
  Software Config to drive os-net-config to configure multiple interfaces
  for the compute role. This is an example for a Nova compute node using
  Contrail vrouter and the vhost0 interface.

```

```

parameters:
  ControlPlaneIp:
    default: ''
    description: IP address/subnet on the ctlplane network
    type: string
  ExternalIpSubnet:
    default: ''
    description: IP address/subnet on the external network
    type: string
  InternalApiIpSubnet:
    default: ''
    description: IP address/subnet on the internal_api network
    type: string
  InternalApiDefaultRoute: # Not used by default in this template
    default: '10.0.0.1'
    description: The default route of the internal api network.
    type: string
  StorageIpSubnet:
    default: ''
    description: IP address/subnet on the storage network
    type: string
  StorageMgmtIpSubnet:
    default: ''
    description: IP address/subnet on the storage_mgmt network
    type: string
  TenantIpSubnet:
    default: ''
    description: IP address/subnet on the tenant network
    type: string
  ManagementIpSubnet: # Only populated when including
environments/network-management.yaml
    default: ''
    description: IP address/subnet on the management network
    type: string
  ExternalNetworkVlanID:
    default: 10
    description: Vlan ID for the external network traffic.
    type: number
  InternalApiNetworkVlanID:
    default: 20
    description: Vlan ID for the internal_api network traffic.
    type: number
  StorageNetworkVlanID:
    default: 30

```

```

    description: Vlan ID for the storage network traffic.
    type: number
StorageMgmtNetworkVlanID:
    default: 40
    description: Vlan ID for the storage mgmt network traffic.
    type: number
TenantNetworkVlanID:
    default: 50
    description: Vlan ID for the tenant network traffic.
    type: number
ManagementNetworkVlanID:
    default: 60
    description: Vlan ID for the management network traffic.
    type: number
ControlPlaneSubnetCidr: # Override this via parameter_defaults
    default: '24'
    description: The subnet CIDR of the control plane network.
    type: string
ControlPlaneDefaultRoute: # Override this via parameter_defaults
    description: The default route of the control plane network.
    type: string
ExternalInterfaceDefaultRoute: # Not used by default in this template
    default: '10.0.0.1'
    description: The default route of the external network.
    type: string
ManagementInterfaceDefaultRoute: # Commented out by default in this template
    default: unset
    description: The default route of the management network.
    type: string
DnsServers: # Override this via parameter_defaults
    default: []
    description: A list of DNS servers (2 max for some implementations) that will
be added to resolv.conf.
    type: comma_delimited_list
EC2MetadataIp: # Override this via parameter_defaults
    description: The IP address of the EC2 metadata server.
    type: string

```

```

resources:
  OsNetConfigImpl:
    type: OS::Heat::SoftwareConfig
    properties:
      group: script

```

```

config:
  str_replace:
    template:
      get_file: ../../scripts/run-os-net-config.sh
    params:
      $network_config:
        network_config:
          - type: interface
            name: nic1
            use_dhcp: false
            dns_servers:
              get_param: DnsServers
            addresses:
          - ip_netmask:
              list_join:
                - '/'
                - - get_param: ControlPlaneIp
                  - get_param: ControlPlaneSubnetCidr
            routes:
          - ip_netmask: 169.x.x.x/32
            next_hop:
              get_param: EC2MetadataIp
          - default: true
            next_hop:
              get_param: ControlPlaneDefaultRoute
          - type: vlan
            vlan_id:
              get_param: InternalApiNetworkVlanID
            device: nic1
            addresses:
          - ip_netmask:
              get_param: InternalApiIpSubnet
          - type: vlan
            vlan_id:
              get_param: StorageNetworkVlanID
            device: nic1
            addresses:
          - ip_netmask:
              get_param: StorageIpSubnet
          - type: contrail_vrouter
            name: vhost0
            use_dhcp: false
            members:
              -

```

```

        type: interface
        name: nic2
        use_dhcp: false
    addresses:
    - ip_netmask:
        get_param: TenantIpSubnet

```

```

outputs:
  OS::stack_id:
    description: The OsNetConfigImpl resource.
    value:
      get_resource: OsNetConfigImpl

```

Advanced vRouter Kernel Mode Configuration

IN THIS SECTION

- [VLAN | 310](#)
- [Bond | 311](#)
- [Bond + VLAN | 311](#)

In addition to the standard NIC configuration, the vRouter kernel mode supports VLAN, Bond, and Bond + VLAN modes. The configuration snippets below only show the relevant section of the NIC template configuration for each mode.

VLAN

```

- type: vlan
  vlan_id:
    get_param: TenantNetworkVlanID
  device: nic2
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
  -
    type: interface
    name:
      str_replace:

```

```

        template: vlanVLANID
        params:
            VLANID: {get_param: TenantNetworkVlanID}
        use_dhcp: false
    addresses:
    - ip_netmask:
        get_param: TenantIpSubnet

```

Bond

```

- type: linux_bond
  name: bond0
  bonding_options: "mode=4 xmit_hash_policy=layer2+3"
  use_dhcp: false
  members:
    -
      type: interface
      name: nic2
    -
      type: interface
      name: nic3
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
    -
      type: interface
      name: bond0
      use_dhcp: false
  addresses:
  - ip_netmask:
      get_param: TenantIpSubnet

```

Bond + VLAN

```

- type: linux_bond
  name: bond0
  bonding_options: "mode=4 xmit_hash_policy=layer2+3"
  use_dhcp: false
  members:
    -
      type: interface

```

```

    name: nic2
  -
    type: interface
    name: nic3
- type: vlan
  vlan_id:
    get_param: TenantNetworkVlanID
  device: bond0
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
    -
      type: interface
      name:
        str_replace:
          template: vlanVLANID
          params:
            VLANID: {get_param: TenantNetworkVlanID}
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet

```

Advanced vRouter DPDK Mode Configuration

IN THIS SECTION

- [Standard | 313](#)
- [VLAN | 313](#)
- [Bond | 314](#)
- [Bond + VLAN | 314](#)

In addition to the standard NIC configuration, the vRouter DPDK mode supports Standard, VLAN, Bond, and Bond + VLAN modes.

Network Environment Configuration:

```
vi ~/tripleo-heat-templates/environments/contrail/contrail-services.yaml
```


Enable the number of hugepages:

```
parameter_defaults:
  ContrailDpdkHugepages1GB: 10
```

See the following NIC template configurations for vRouter DPDK mode. The configuration snippets below only show the relevant section of the NIC configuration for each mode.

Standard

```
- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet
```

VLAN

```
- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  vlan_id:
    get_param: TenantNetworkVlanID
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet
```

Bond

```

- type: contrail_vrouter_dpkg
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  bond_mode: 4
  bond_policy: layer2+3
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
    -
      type: interface
      name: nic3
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet

```

Bond + VLAN

```

- type: contrail_vrouter_dpkg
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  vlan_id:
    get_param: TenantNetworkVlanID
  bond_mode: 4
  bond_policy: layer2+3
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
    -
      type: interface
      name: nic3
      use_dhcp: false
  addresses:

```

```
- ip_netmask:
  get_param: TenantIpSubnet
```

Advanced vRouter SRIOV + Kernel Mode Configuration

IN THIS SECTION

- [VLAN | 316](#)
- [Bond | 316](#)
- [Bond + VLAN | 317](#)

vRouter SRIOV + Kernel mode can be used in the following combinations:

- Standard
- VLAN
- Bond
- Bond + VLAN

Network environment configuration:

```
vi ~/tripleo-heat-templates/environments/contrail/contrail-services.yaml
```

Enable the number of hugepages:

```
parameter_defaults:
  ContrailSriovHugepages1GB: 10
```

SRIOV PF/VF settings:

```
NovaPCIPassthrough:
- devname: "ens2f1"
  physical_network: "sriov1"
ContrailSriovNumVFs: ["ens2f1:7"]
```

The SRIOV NICs are not configured in the NIC templates. However, vRouter NICs must still be configured. See the following NIC template configurations for vRouter kernel mode. The configuration snippets below only show the relevant section of the NIC configuration for each mode.

VLAN

```
- type: vlan
  vlan_id:
    get_param: TenantNetworkVlanID
  device: nic2
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
    -
      type: interface
      name:
        str_replace:
          template: vlanVLANID
          params:
            VLANID: {get_param: TenantNetworkVlanID}
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet
```

Bond

```
- type: linux_bond
  name: bond0
  bonding_options: "mode=4 xmit_hash_policy=layer2+3"
  use_dhcp: false
  members:
    -
      type: interface
      name: nic2
    -
      type: interface
      name: nic3
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
    -
```

```

    type: interface
    name: bond0
    use_dhcp: false
addresses:
- ip_netmask:
    get_param: TenantIpSubnet

```

Bond + VLAN

```

- type: linux_bond
  name: bond0
  bonding_options: "mode=4 xmit_hash_policy=layer2+3"
  use_dhcp: false
  members:
    -
      type: interface
      name: nic2
    -
      type: interface
      name: nic3
- type: vlan
  vlan_id:
    get_param: TenantNetworkVlanID
  device: bond0
- type: contrail_vrouter
  name: vhost0
  use_dhcp: false
  members:
    -
      type: interface
      name:
        str_replace:
          template: vlanVLANID
          params:
            VLANID: {get_param: TenantNetworkVlanID}
      use_dhcp: false
addresses:
- ip_netmask:
    get_param: TenantIpSubnet

```

Advanced vRouter SRIOV + DPDK Mode Configuration

IN THIS SECTION

- [Standard | 319](#)
- [VLAN | 319](#)
- [Bond | 319](#)
- [Bond + VLAN | 320](#)

vRouter SRIOV + DPDK can be used in the following combinations:

- Standard
- VLAN
- Bond
- Bond + VLAN

Network environment configuration:

```
vi ~/tripleo-heat-templates/environments/contrail/contrail-services.yaml
```

Enable the number of hugepages

```
parameter_defaults:
  ContrailSriovMode: dpdk
  ContrailDpdkHugepages1GB: 10
  ContrailSriovHugepages1GB: 10
```

SRIOV PF/VF settings

```
NovaPCIPassthrough:
- devname: "ens2f1"
  physical_network: "sriov1"
ContrailSriovNumVFs: ["ens2f1:7"]
```

The SRIOV NICs are not configured in the NIC templates. However, vRouter NICs must still be configured. See the following NIC template configurations for vRouter DPDK mode. The configuration snippets below only show the relevant section of the NIC configuration for each mode.

Standard

```
- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet
```

VLAN

```
- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  vlan_id:
    get_param: TenantNetworkVlanID
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet
```

Bond

```
- type: contrail_vrouter_dpdk
  name: vhost0
```

```

use_dhcp: false
driver: uio_pci_generic
cpu_list: 0x01
bond_mode: 4
bond_policy: layer2+3
members:
  -
    type: interface
    name: nic2
    use_dhcp: false
  -
    type: interface
    name: nic3
    use_dhcp: false
addresses:
  - ip_netmask:
      get_param: TenantIpSubnet

```

Bond + VLAN

```

- type: contrail_vrouter_dpdk
  name: vhost0
  use_dhcp: false
  driver: uio_pci_generic
  cpu_list: 0x01
  vlan_id:
    get_param: TenantNetworkVlanID
  bond_mode: 4
  bond_policy: layer2+3
  members:
    -
      type: interface
      name: nic2
      use_dhcp: false
    -
      type: interface
      name: nic3
      use_dhcp: false
  addresses:
    - ip_netmask:
        get_param: TenantIpSubnet

```


Advanced Scenarios


Remote Compute

Remote Compute extends the data plane to remote locations (POP) whilst keeping the control plane central. Each POP will have its own set of Contrail control services, which are running in the central location. The difficulty is to ensure that the compute nodes of a given POP connect to the Control nodes assigned to that POP. The Control nodes must have predictable IP addresses and the compute nodes have to know these IP addresses. In order to achieve that the following methods are used:

- Custom Roles
- Static IP assignment
- Precise Node placement
- Per Node hieradata

Each overcloud node has a unique DMI UUID. This UUID is known on the undercloud node as well as on the overcloud node. Hence, this UUID can be used for mapping node specific information. For each POP, a Control role and a Compute role has to be created.

Overview



Mapping Table

Table 17: Mapping Table

Nova Name	Ironic Name	UUID	KVM	IP Address	POP
overcloud -contrailcontrolonly -0	control-only-1- 5b3s30	Ironic UUID: 7d758dce-2784- 45fd-be09-5a41eb53e764 DMI UUID: 73F8D030-E896- 4A95-A9F5-E1A4FEBE322D	5b3s30	10.0.0.11	POP1

Table 17: Mapping Table (continued)

Nova Name	Ironic Name	UUID	KVM	IP Address	POP
overcloud -contrailcontrolonly -1	control-only-2- 5b3s30	Ironic UUID: d26abdeb-d514-4a37-a7fb-2cd2511c351f DMI UUID: 14639A66-D62C-4408-82EE-FDDC4E509687	5b3s30	10.0.0.14	POP2
overcloud -contrailcontrolonly -2	control-only-1- 5b3s31	Ironic UUID: 91dd9fa9-e8eb-4b51-8b5e-bbaffb6640e4 DMI UUID: 28AB0B57-D612-431E-B177-1C578AE0FEA4	5b3s31	10.0.0.12	POP1
overcloud -contrailcontrolonly -3	control-only-2- 5b3s31	Ironic UUID: 09fa57b8-580f-42ec-bf10-a19573521ed4 DMI UUID: 09BEC8CB-77E9-42A6-AFF4-6D4880FD87D0	5b3s31	10.0.0.15	POP2
overcloud -contrailcontrolonly -4	control-only-1- 5b3s32	Ironic UUID: 4766799-24c8-4e3b-af54-353f2b796ca4 DMI UUID: 3993957A-ECBF-4520-9F49-0AF6EE1667A7	5b3s32	10.0.0.13	POP1
overcloud -contrailcontrolonly -5	control-only-2- 5b3s32	Ironic UUID: 58a803ae-a785-470e-9789-139abbfa74fb DMI UUID: AF92F485-C30C-4D0A-BDC4-C6AE97D06A66	5b3s32	10.0.0.16	POP2

ControlOnly preparation

Add ControlOnly overcloud VMs to overcloud KVM host

NOTE: This has to be done on the overcloud KVM hosts

Two ControlOnly overcloud VM definitions will be created on each of the overcloud KVM hosts.

```

ROLES=control-only:2
num=4
ipmi_user=<user>
ipmi_password=<password>
libvirt_path=/var/lib/libvirt/images
port_group=overcloud
prov_switch=br0

/bin/rm ironic_list
IFS=' ' read -ra role_list <<< "${ROLES}"
for role in ${role_list[@]}; do
    role_name=`echo $role|cut -d ":" -f 1`
    role_count=`echo $role|cut -d ":" -f 2`
    for count in `seq 1 ${role_count}`; do
        echo $role_name $count
        qemu-img create -f qcow2 ${libvirt_path}/${role_name}_${count}.qcow2 99G
        virsh define /dev/stdin <<EOF
$(virt-install --name ${role_name}_${count} \
--disk ${libvirt_path}/${role_name}_${count}.qcow2 \
--vcpus=4 \
--ram=16348 \
--network network=br0,model=virtio,portgroup=${port_group} \
--network network=br1,model=virtio \
--virt-type kvm \
--cpu host \
--import \
--os-variant rhel7 \
--serial pty \
--console pty,target_type=virtio \
--graphics vnc \
--print-xml)
EOF
        vbmc add ${role_name}_${count} --port 1623${num} --username ${ipmi_user}
--password ${ipmi_password}
        vbmc start ${role_name}_${count}
        prov_mac=`virsh domiflist ${role_name}_${count}|grep ${prov_switch}|awk '{print $5}'`
        vm_name=${role_name}-${count}-${hostname} -s`
        kvm_ip=`ip route get 1 |grep src |awk '{print $7}'`
        echo ${prov_mac} ${vm_name} ${kvm_ip} ${role_name} 1623${num}>> ironic_list
        num=$(expr $num + 1)
    done
done

```

NOTE: The generated *ironic_list* will be needed on the undercloud to import the nodes to Ironic.

Get the *ironic_lists* from the overcloud KVM hosts and combine them.

```
cat ironic_list_control_only
52:54:00:3a:2f:ca control-only-1-5b3s30 10.87.64.31 control-only 16234
52:54:00:31:4f:63 control-only-2-5b3s30 10.87.64.31 control-only 16235
52:54:00:0c:11:74 control-only-1-5b3s31 10.87.64.32 control-only 16234
52:54:00:56:ab:55 control-only-2-5b3s31 10.87.64.32 control-only 16235
52:54:00:c1:f0:9a control-only-1-5b3s32 10.87.64.33 control-only 16234
52:54:00:f3:ce:13 control-only-2-5b3s32 10.87.64.33 control-only 16235
```

Import:

```
ipmi_password=<password>
ipmi_user=<user>

DEPLOY_KERNEL=$(openstack image show bm-deploy-kernel -f value -c id)
DEPLOY_RAMDISK=$(openstack image show bm-deploy-ramdisk -f value -c id)

num=0
while IFS= read -r line; do
    mac=`echo $line|awk '{print $1}'`
    name=`echo $line|awk '{print $2}'`
    kvm_ip=`echo $line|awk '{print $3}'`
    profile=`echo $line|awk '{print $4}'`
    ipmi_port=`echo $line|awk '{print $5}'`
    uuid=`openstack baremetal node create --driver ipmi \
        --property cpus=4 \
        --property memory_mb=16348 \
        --property local_gb=100 \
        --property cpu_arch=x86_64 \
        --driver-info ipmi_username=${ipmi_user} \
        \
        --driver-info ipmi_address=${kvm_ip} \
        --driver-info ipmi_password=${ipmi_password} \
        \
        --driver-info ipmi_port=${ipmi_port} \
        --name=${name} \
        --property capabilities=boot_option:local
```

```
\
                                -c uuid -f value`
openstack baremetal node set ${uuid} --driver-info deploy_kernel=$DEPLOY_KERNEL
--driver-info deploy_ramdisk=$DEPLOY_RAMDISK
openstack baremetal port create --node ${uuid} ${mac}
openstack baremetal node manage ${uuid}
num=$(expr $num + 1)
done < <(cat ironic_list_control_only)
```

ControlOnly node introspection

```
openstack overcloud node introspect --all-manageable --provide
```

Get the ironic UUID of the ControlOnly nodes

```
openstack baremetal node list |grep control-only
| 7d758dce-2784-45fd-be09-5a41eb53e764 | control-only-1-5b3s30 | None | power off
| available | False |
| d26abdeb-d514-4a37-a7fb-2cd2511c351f | control-only-2-5b3s30 | None | power off
| available | False |
| 91dd9fa9-e8eb-4b51-8b5e-bbaffb6640e4 | control-only-1-5b3s31 | None | power off
| available | False |
| 09fa57b8-580f-42ec-bf10-a19573521ed4 | control-only-2-5b3s31 | None | power off
| available | False |
| f4766799-24c8-4e3b-af54-353f2b796ca4 | control-only-1-5b3s32 | None | power off
| available | False |
| 58a803ae-a785-470e-9789-139abbfa74fb | control-only-2-5b3s32 | None | power off
| available | False |
```

The first ControlOnly node on each of the overcloud KVM hosts will be used for POP1, the second for POP2, and so and so forth.

Get the ironic UUID of the POP compute nodes:

```
openstack baremetal node list |grep compute
| 91d6026c-b9db-49cb-a685-99a63da5d81e | compute-3-5b3s30 | None | power off |
available | False |
| 8028eb8c-e1e6-4357-8fcf-0796778bd2f7 | compute-4-5b3s30 | None | power off |
available | False |
| b795b3b9-c4e3-4a76-90af-258d9336d9fb | compute-3-5b3s31 | None | power off |
```

```
available | False |
| 2d4be83e-6fcc-4761-86f2-c2615dd15074 | compute-4-5b3s31 | None | power off |
available | False |
```

The first two compute nodes belong to POP1 the second two compute nodes belong to POP2.

Create an input YAML using the ironic UUIDs:

```
~/subcluster_input.yaml
---
- subcluster: subcluster1
  asn: "65413"
  control_nodes:
    - uuid: 7d758dce-2784-45fd-be09-5a41eb53e764
      ipaddress: 10.0.0.11
    - uuid: 91dd9fa9-e8eb-4b51-8b5e-bbaffb6640e4
      ipaddress: 10.0.0.12
    - uuid: f4766799-24c8-4e3b-af54-353f2b796ca4
      ipaddress: 10.0.0.13
  compute_nodes:
    - uuid: 91d6026c-b9db-49cb-a685-99a63da5d81e
      vrouter_gateway: 10.0.0.1
    - uuid: 8028eb8c-e1e6-4357-8fcf-0796778bd2f7
      vrouter_gateway: 10.0.0.1
- subcluster: subcluster2
  asn: "65414"
  control_nodes:
    - uuid: d26abdeb-d514-4a37-a7fb-2cd2511c351f
      ipaddress: 10.0.0.14
    - uuid: 09fa57b8-580f-42ec-bf10-a19573521ed4
      ipaddress: 10.0.0.15
    - uuid: 58a803ae-a785-470e-9789-139abbfa74fb
      ipaddress: 10.0.0.16
  compute_nodes:
    - uuid: b795b3b9-c4e3-4a76-90af-258d9336d9fb
      vrouter_gateway: 10.0.0.1
    - uuid: 2d4be83e-6fcc-4761-86f2-c2615dd15074
      vrouter_gateway: 10.0.0.1
```

NOTE: Only control_nodes, compute_nodes, dpdk_nodes and sriov_nodes are supported.

Generate subcluster environment:

```
~/tripleo-heat-templates/tools/contrail/create_subcluster_environment.py -i
~/subcluster_input.yaml \
    -o
~/tripleo-heat-templates/environments/contrail/contrail-subcluster.yaml
```

Check subcluster environment file:

```
cat ~/tripleo-heat-templates/environments/contrail/contrail-subcluster.yaml
parameter_defaults:
  NodeDataLookup:
    041D7B75-6581-41B3-886E-C06847B9C87E:
      contrail_settings:
        CONTROL_NODES: 10.0.0.14,10.0.0.15,10.0.0.16
        SUBCLUSTER: subcluster2
        VROUTER_GATEWAY: 10.0.0.1
    09BEC8CB-77E9-42A6-AFF4-6D4880FD87D0:
      contrail_settings:
        BGP_ASN: '65414'
        SUBCLUSTER: subcluster2
    14639A66-D62C-4408-82EE-FDDC4E509687:
      contrail_settings:
        BGP_ASN: '65414'
        SUBCLUSTER: subcluster2
    28AB0B57-D612-431E-B177-1C578AE0FEA4:
      contrail_settings:
        BGP_ASN: '65413'
        SUBCLUSTER: subcluster1
    3993957A-ECBF-4520-9F49-0AF6EE1667A7:
      contrail_settings:
        BGP_ASN: '65413'
        SUBCLUSTER: subcluster1
    73F8D030-E896-4A95-A9F5-E1A4FEBE322D:
      contrail_settings:
        BGP_ASN: '65413'
        SUBCLUSTER: subcluster1
    7933C2D8-E61E-4752-854E-B7B18A424971:
      contrail_settings:
        CONTROL_NODES: 10.0.0.14,10.0.0.15,10.0.0.16
        SUBCLUSTER: subcluster2
        VROUTER_GATEWAY: 10.0.0.1
    AF92F485-C30C-4D0A-BDC4-C6AE97D06A66:
```



```

    contrail_settings:
      BGP_ASN: '65414'
      SUBCLUSTER: subcluster2
BB9E9D00-57D1-410B-8B19-17A0DA581044:
    contrail_settings:
      CONTROL_NODES: 10.0.0.11,10.0.0.12,10.0.0.13
      SUBCLUSTER: subcluster1
      VROUTER_GATEWAY: 10.0.0.1
E1A809DE-FDB2-4EB2-A91F-1B3F75B99510:
    contrail_settings:
      CONTROL_NODES: 10.0.0.11,10.0.0.12,10.0.0.13
      SUBCLUSTER: subcluster1
      VROUTER_GATEWAY: 10.0.0.1

```

Deployment

Add `contrail-subcluster.yaml`, `contrail-ips-from-pool-all.yaml` and `contrail-scheduler-hints.yaml` to the OpenStack deploy command:

```

openstack overcloud deploy --templates ~/tripleo-heat-templates \
-e ~/overcloud_images.yaml \
-e ~/tripleo-heat-templates/environments/network-isolation.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-plugins.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-services.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-net.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-subcluster.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-ips-from-pool-all.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-scheduler-hints.yaml \
--roles-file ~/tripleo-heat-templates/roles_data_contrail_aio.yaml

```

Installing Overcloud

1. Deployment:

```

openstack overcloud deploy --templates ~/tripleo-heat-templates \
-e ~/overcloud_images.yaml \
-e ~/tripleo-heat-templates/environments/network-isolation.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-plugins.yaml \
-e ~/tripleo-heat-templates/environments/contrail/contrail-services.yaml \

```

```
-e ~/tripleo-heat-templates/environments/contrail/contrail-net.yaml \
--roles-file ~/tripleo-heat-templates/roles_data_contrail_aio.yaml
```

2. Validation Test:

```
source overcloudrc
curl -O http://download.cirros-cloud.net/0.3.5/cirros-0.3.5-x86_64-disk.img
openstack image create --container-format bare --disk-format qcow2 --file
cirros-0.3.5-x86_64-disk.img cirros
openstack flavor create --public cirros --id auto --ram 64 --disk 0 --vcpus 1
openstack network create net1
openstack subnet create --subnet-range 1.0.0.0/24 --network net1 sn1
nova boot --image cirros --flavor cirros --nic net-id=`openstack network show
net1 -c id -f value` --availability-zone nova:overcloud-novacompute-0.localdomain
c1
nova list
```

RELATED DOCUMENTATION

| *Installing a Nested Red Hat OpenShift Container Platform 3.11 Cluster Using Contrail Ansible Deployer*

Using Netronome SmartNIC vRouter with Contrail Networking

NOTE: The Netronome SmartNIC vRouter technology covered in this document is available for evaluation purposes only. It is not intended for deployment in production networks.

Contrail supports Netronome Agilio CX SmartNICs for Contrail Networking deployment with Red Hat OpenStack Platform Director (RHOSPd) 13 environment.

This feature will enable service providers to improve the forwarding performance which includes packets per second (PPS) of vRouter. This will optimize server CPU usage and you can deploy more Virtual network functions (VNFs) per server.

Benefits:

- Increased PPS capacity of Contrail vRouter datapath allowing applications to reach their full processing capacity.
- Reclaimed CPU cores from Contrail vRouter off-loading allowing more VMs and VNFs to be deployed per server.

The goal of this topic is to provide a procedure for deploying accelerated vRouter compute nodes.

Before you begin:

- Equip compute nodes with Netronome Agilio CX SmartNIC.

For details, refer to [Agilio CX SmartNICs](#).

- Retrieve *Agilio heat-template plugin*.

Register on Netronome support site at <https://help.netronome.com> and provide Docker Hub credentials.

Netronome will provide the TripleO templates for SmartNIC vRouter deployment on compute nodes. Also, Netronome will authorize Docker Hub registry access.

For details, refer to [Netronome Agilio vRouter 19xx deployment guide](#).

- Note the following version tags:

AGILIO_TAG="2.38-rhel-queens

FORWARDER_TAG="2.38-rhel-queens

Procedure:

NOTE: If you have multiple undercloud nodes deployed, you must perform the following procedure on the same node.

1. Configure *Agilio* plugin.

For details, refer to [Netronome agilio-ovs-openstack-plugin GitHub Repository](#).

- a. Extract the *Agilio* plugin archive and copy the **agilio-plugin** folder into the **tripleo-heat-templates** directory.

```
[stack@queensa ~]$ tar -xzf rhosp-contrail-agilio-heat-plugin-5-34.tgz
agilio-plugin/
agilio-plugin/agilio-vrouter.yaml
agilio-plugin/agilio_upgrade.sh
agilio-plugin/deploy_rhosp.sh
agilio-plugin/nfp_udev.sh
agilio-plugin/agilio-env.yaml
agilio-plugin/version
agilio-plugin/README.md
[stack@queensa ~]$ cp -r agilio-plugin/ tripleo-heat-templates/
```

- b. Navigate to the **agilio-plugin** directory on the undercloud node.

```
[tripleo-heat-templates]$ cd agilio-plugin/
```

- c. Modify **agilio-env.yaml** file as per your environment.

NOTE: Reserve at least 1375*2 MB hugepages for *virtio-forwarder*.

Sample **agilio-env.yaml** file:

```
resource_registry:
  OS::TripleO::NodeExtraConfigPost: agilio-vrouter.yaml

parameter_defaults:
  # Hugepages
  ContrailVrouterHugepages2MB: "8192"
  # IOMMU
```

```

ComputeParameters:
  KernelArgs: "intel_iommu=on iommu=pt isolcpus=1,2 "

ComputeCount: 3

# Additional config
ControlPlaneDefaultRoute: 10.0.x.1
EC2MetadataIp: 10.0.x.1 # Generally the IP of the Undercloud
DnsServers: ["8.8.8.8", "192.168.3.3"]
NtpServer: ntp.is.co.za
ContrailRegistryInsecure: true
DockerInsecureRegistryAddress: 172.x.x.150:6666,10.0.x.1:8787
ContrailRegistry: 172.x.x.150:6666
ContrailImageTag: <container_tag>-rhel-queens

# Fix DB Diskspace too low issue
ContrailAnalyticsDBMinDiskGB: 40

```

- d. Add Docker Hub credentials to **tripleo-heat-templates/agilio-plugin/agililo_upgrade.sh** file to retrieve containers from **AGILIO_REPO="docker.io/netronomesystems/"** repository.

#GENERAL DOCKER CONFIG

DOCKER_USR=*** #ENTER_DOCKER_USERNAME_HERE**

DOCKER_PASS=*** #ENTER_DOCKER_PASSWORD_HERE**

```

[root@overcloud-novacompute-2 heat-admin]# docker ps -a | grep virtio_for
7d5af8a2591d docker.io/netronomesystems/virtio-forwarder:2.38-rhel-queens "/entrypoint.sh"
30 seconds ago Up 15 seconds virtio_forwarder

```

```

[root@overcloud-novacompute-2 heat-admin]# docker ps -a | grep agilio
c7c611b5168b docker.io/netronomesystems/agilio-vrouter:2.38-rhel-queens "/entrypoint.sh"
46 seconds ago Up 38 seconds agilio_vrouter

```

2. Prepare the Contrail Networking cluster for deployment.

Refer to the following topics for deployment:

- [Understanding Red Hat OpenStack Platform Director on page 271](#)
- [Setting Up the Infrastructure on page 275](#)
- [Setting Up the Undercloud on page 284](#)
- [Setting Up the Overcloud on page 286](#)

NOTE: Do not perform steps for [“Installing Overcloud” on page 269](#).

3. Deploy the cluster by one of the following ways:

- Add **agilio-env.yaml** to installing overcloud step as mentioned in [“Installing Overcloud” on page 269](#) topic.

```
openstack overcloud deploy --templates ~/tripleo-heat-templates
-e ~/overcloud_images.yaml
-e ~/tripleo-heat-templates/environments/network-isolation.yaml
-e ~/tripleo-heat-templates/environments/contrail/contrail-plugins.yaml
-e ~/tripleo-heat-templates/environments/contrail/contrail-services.yaml
-e ~/tripleo-heat-templates/environments/contrail/contrail-net.yaml
-e ~/tripleo-heat-templates/agilio-plugin/agilio-env.yaml
--roles-file ~/tripleo-heat-templates/roles_data_contrail_aio.yaml
```

Or

- Run the following command:

```
deploy_rhosp.sh
-e ~/tripleo-heat-templates/agilio-plugin/agilio-env.yaml
```

On completing above steps successfully, refer to [Netronome agilio-ovs-openstack-plugin GitHub Repository](#) on how to spin up the accelerated VMs.

SEE ALSO

[Understanding Red Hat OpenStack Platform Director | 271](#)

[Setting Up the Infrastructure | 275](#)

[Setting Up the Undercloud | 284](#)

[Setting Up the Overcloud | 286](#)

Installing OpenStack Octavia LBaaS with RHOSP in Contrail Networking

Contrail Networking Release 2005 supports Octavia as LBaaS. The deployment supports RHOSP and Juju platforms.

With Octavia as LBaaS, Contrail Networking is only maintaining network connectivity and is not involved in any load balancing functions.

For each OpenStack load balancer creation, Octavia launches a VM known as *amphora* VM. The VM starts the HAPROXY when listener is created for the load balancer in OpenStack. Whenever the load balancer gets updated in OpenStack, *amphora* VM updates the running HAPROXY configuration. The *amphora* VM is deleted on deleting the load balancer.

Contrail Networking provides connectivity to *amphora* VM interfaces. *Amphora* VM has two interfaces; one for management and the other for data. The management interface is used by the Octavia services for the management communication. Since, Octavia services are running in the underlay network and *amphora* VM is running in the overlay network, SDN gateway is needed to reach the overlay network. The data interface is used for load balancing.

Follow the procedure to install OpenStack Octavia LBaaS with Contrail Networking:

1. Deploy RHOSP13 with Contrail Networking without Octavia.

```
openstack overcloud deploy --templates tripleo-heat-templates/ \
--roles-file tripleo-heat-templates/roles_data_contrail_aio.yaml \
-e environment-rhel-registration.yaml \
-e \
tripleo-heat-templates/extraconfig/pre_deploy/rhel-registration/rhel-registration-resource-registry.yaml \
-e tripleo-heat-templates/environments/contrail/contrail-services.yaml \
-e tripleo-heat-templates/environments/contrail/contrail-net-single.yaml \
-e tripleo-heat-templates/environments/contrail/contrail-plugins.yaml \
-e misc_opts.yaml \
-e contrail-parameters.yaml \
-e docker_registry.yaml
```

2. Make a copy of **tripleo-heat-templates/docker/services/octavia/octavia-deployment-config.yaml** file.

```
cp
tripleo-heat-templates/docker/services/octavia/octavia-deployment-config.yaml
tripleo-heat-templates/docker/services/octavia/octavia-deployment-config.bak
```

3. Make the following changes in *generate_certs* section of the **tripleo-heat-templates/docker/services/octavia/octavia-deployment-config.yaml** file.

```
conditions:

  generate_certs:
    and:
      - get_param: OctaviaGenerateCerts
      - or:
```

```

- equals:
  - get_param: StackAction
  - CREATE
- equals:
  - get_param: StackAction
  - UPDATE

```

4. Deploy RHOSP13 with Octavia services.

```

openstack overcloud deploy --templates tripleo-heat-templates/ \ --roles-file
tripleo-heat-templates/roles_data_contrail_aio.yaml \
-e environment-rhel-registration.yaml \
-e
tripleo-heat-templates/extraconfig/pre_deploy/rhel-registration/rhel-registration-resource-registry.yaml \
-e tripleo-heat-templates/environments/contrail/contrail-services.yaml \
-e tripleo-heat-templates/environments/contrail/contrail-net-single.yaml \
-e tripleo-heat-templates/environments/contrail/contrail-plugins.yaml \
-e tripleo-heat-templates/environments/services/octavia.yaml \
-e misc_opts.yaml \
-e contrail-parameters.yaml \
-e docker_registry.yaml

```

5. Rollback changes in **tripleo-heat-templates/docker/services/octavia/octavia-deployment-config.yaml** file.

```

cp
tripleo-heat-templates/docker/services/octavia/octavia-deployment-config.bak
tripleo-heat-templates/docker/services/octavia/octavia-deployment-config.yaml

```


Here is an example for creating and testing load balancer:

Prerequisites:

- You must have connectivity between Octavia controller and amphora instances,
- You must have OpenStack services into LXD containers.
- You must have separate interfaces for control plane and data plane.

1. Create private network.

```
openstack network create private
openstack subnet create private --network private --subnet-range
10.10.10.0/24 --allocation-pool
start=10.10.10.50,end=10.10.10.70 --gateway none
```

2. Create security group.

```
openstack security group create allow_all
openstack security group rule create --ingress --protocol any --prefix
'0.0.0.0/0' allow_all
```

3. Check available flavors and images. You can create them, if needed.

```
openstack flavor list
openstack image list
```

4. Create two servers for load balancer.

```
openstack server create --flavor test_flavor --image cirros --security-group
allow_all --network private cirros1
openstack server create --flavor test_flavor --image cirros --security-group
allow_all --network private cirros2
```

5. Create additional server to test load balancer.

```
openstack server create --flavor test_flavor --image cirros --security-group
allow_all --network private cirros-test
```

6. Check status and IP addresses.

```
openstack server list --long
```

7. Create simple HTTP server on every cirros. Login on both the cirros instances and run following commands:

```
MYIP=$(ifconfig eth0|grep 'inet addr'|awk -F: '{print $2}'| awk '{print $1}') while true; do echo -e "HTTP/1.0 200 OK\r\n\r\nWelcome to $MYIP" | sudo nc -l -p 80 ; done&
```

8. Create load balancer

```
openstack loadbalancer create --name lb1 --vip-subnet-id private
```

Make sure *provisioning_status* is *Active*.

```
openstack loadbalancer show lb1
```

9. Setup load balancer

```
openstack loadbalancer listener create --protocol HTTP --protocol-port 80
--name listener1 lb1
```

```
openstack loadbalancer show lb1 # Wait for the provisioning_status to be
ACTIVE.
```

```
openstack loadbalancer pool create --lb-algorithm ROUND_ROBIN --listener
listener1 --protocol HTTP --name pool1
```

```
openstack loadbalancer healthmonitor create --delay 5 --timeout 2
--max-retries 1 --type HTTP pool1
```

```
openstack loadbalancer member create --subnet-id private --address
10.10.10.50 --protocol-port 80 pool1
```

```
openstack loadbalancer member create --subnet-id private --address
10.10.10.51 --protocol-port 80 pool1
```

IP addresses 10.10.10.50 and 10.10.10.51 belong to VMs created with test http server in step [7](#).

10. Check the status of load balancer.

```
openstack loadbalancer show lb1 # Wait for the provisioning_status to be
ACTIVE.
```

```
openstack loadbalancer pool list
```

```
openstack loadbalancer pool show pool1
```

```
openstack loadbalancer member list pool1
```

```
openstack loadbalancer listener list
```

11. Login to load balancer client and verify if round robin works.

cirros@169.x.0.9's password:

```
$ curl 10.10.10.50
```

```
Welcome to 10.10.10.52
```

```
$ curl 10.10.10.50
```

```
Welcome to 10.10.10.53
```

```
$ curl 10.10.10.50
```

```
Welcome to 10.10.10.52
$ curl 10.10.10.50
Welcome to 10.10.10.53
$ curl 10.10.10.50
Welcome to 10.10.10.52
$ curl 10.10.10.50
Welcome to 10.10.10.53
```

SEE ALSO

[Support for OpenStack LBaaS | 381](#)

[Using Load Balancers in Contrail | 368](#)

[Installing OpenStack Octavia LBaaS with Juju Charms in Contrail Networking | 455](#)

Configuring Virtual Networks

IN THIS SECTION

- [Creating Projects in OpenStack for Configuring Tenants in Contrail | 340](#)
- [Creating a Virtual Network with OpenStack Contrail | 342](#)
- [Creating an Image for a Project in OpenStack Contrail | 345](#)
- [Using Security Groups with Virtual Machines \(Instances\) | 349](#)


Creating Projects in OpenStack for Configuring Tenants in Contrail

In Contrail, a tenant configuration is called a project. A project is created for each set of virtual machines (VMs) and virtual networks (VNs) that are configured as a discrete entity for the tenant.

Projects are created, managed, and edited at the OpenStack **Projects** page.

1. Click the **Admin** tab on the OpenStack dashboard, then click the **Projects** link to access the **Projects** page; see [Figure 49 on page 340](#).

Figure 49: OpenStack Projects



2. In the upper right, click the **Create Project** button to access the **Add Project** window; see [Figure 50 on page 341](#).

Figure 50: Add Project

Add Project [X]

Project Info | Project Members | Quota

Name
customer 1

Description
Additional information here...

Enabled
☒

From here you can create a new project to organize users.

Cancel Finish

3. In the **Add Project** window, on the **Project Info** tab, enter a **Name** and a **Description** for the new project, and select the **Enabled** check box to activate this project.
4. In the **Add Project** window, select the **Project Members** tab, and assign users to this project. Designate each user as **admin** or as **Member**.

As a general rule, one person should be a super user in the **admin** role for all projects and a user with a **Member** role should be used for general configuration purposes.
5. Click **Finish** to create the project.

Refer to OpenStack documentation for more information about creating and managing projects.

SEE ALSO

Creating a Virtual Network with Juniper Networks Contrail

[Creating a Virtual Network with OpenStack Contrail | 342](#)


[OpenStack documentation](#)

Creating a Virtual Network with OpenStack Contrail

You can create virtual networks in Contrail Networking from the OpenStack. The following procedure shows how to create a virtual network when using OpenStack.

1. To create a virtual network when using OpenStack Contrail, select **Project > Network > Networks**. The **Networks** page is displayed. See [Figure 51 on page 342](#).

Figure 51: Networks Page



2. Click **Create Network**. The **Create Network** window is displayed. See [Figure 52 on page 343](#) and [Figure 53 on page 343](#).

Figure 52: Create Networks

Create Network

Network

Subnet

Subnet Details

Network Name

test-network

☒ Enable Admin State

☐ Shared

☒ Create Subnet

Create a new network. In addition, a subnet associated with the network can be created in the following steps of this wizard.

Cancel

« Back

Next »

Figure 53: Subnet and Gateway Details

Create Network

Network

Subnet

Subnet Details

Subnet Name

Network Address

IP Version

IPv4

Gateway IP

☐ Disable Gateway

Creates a subnet associated with the network. You need to enter a valid "Network Address" and "Gateway IP". If you did not enter the "Gateway IP", the first value of a network will be assigned by default. If you do not want gateway please check the "Disable Gateway" checkbox. Advanced configuration is available by clicking on the "Subnet Details" tab.

Cancel

« Back

Next »

3. Click the **Network** and **Subnet** tabs to complete the fields in the **Create Network** window. See field descriptions in [Table 18 on page 344](#).

Table 18: Create Network Fields

Field	Description
Network Name	Enter a name for the network.
Subnet Name	Enter a name for the subnetwork.
Network Address	Enter the network address in CIDR format.
IP Version*	Select IPv4 or IPv6.
Gateway IP	Optionally, enter an explicit gateway IP address for the IP address block. Check the Disable Gateway box if no gateway is to be used.

4. Click the **Subnet Details** tab to specify the Allocation Pool, DNS Name Servers, and Host Routes.

Figure 54: Additional Subnet Attributes

Create Network

Network

Subnet

Subnet Details

☒ Enable DHCP

Specify additional attributes for the subnet.

Allocation Pools ?

DNS Name Servers ?

Host Routes ?

Cancel

« Back

Create


5. To save your network, click **Create** , or click **Cancel** to discard your work and start over.

Creating an Image for a Project in OpenStack Contrail

To specify an image to upload to the Image Service for a project in your system by using the OpenStack dashboard:

- 1. In OpenStack, select **Project > Compute > Images**. The Images window is displayed. See [Figure 55 on page 346](#).

Figure 55: OpenStack Images Window



- 2. Make sure you have selected the correct project to which you are associating an image.
- 3. Click **Create Image**.

The **Create An Image** window is displayed. See [Figure 56 on page 347](#).

Figure 56: OpenStack Create An Image Window

Create An Image

Name *

Description

Image Source

Image Location ▼

Image Location ?

http://example.com/image.iso

Format *

Select format ▼

Architecture

Minimum Disk (GB) ?

Minimum RAM (MB) ?

☐ Public

☐ Protected

Cancel

Create Image

Description:

Currently only images available via an HTTP URL are supported. The image location must be accessible to the Image Service. Compressed image binaries are supported (.zip and .tar.gz.)

Please note: The Image Location field MUST be a valid and direct URL to the image binary. URLs that redirect or serve error pages will result in unusable images.

s018515

- Complete the fields to specify your image. [Table 19 on page 348](#) describes each of the fields on the window.

NOTE: Only images available through an HTTP URL are supported, and the image location must be accessible to the Image Service. Compressed image binaries are supported (*.zip and *.tar.gz).

Table 19: Create an Image Fields

Field	Description
Name	Enter a name for this image.
Description	Enter a description for the image.
Image Source	Select Image File or Image Location. If you select Image File, you are prompted to browse to the local location of the file.
Image Location	Enter an external HTTP URL from which to load the image. The URL must be a valid and direct URL to the image binary. URLs that redirect or serve error pages result in unusable images.
Format	Required field. Select the format of the image from a list: AKI- Amazon Kernel Image AMI- Amazon Machine Image ARI- Amazon Ramdisk Image ISO- Optical Disk Image QCOW2- QEMU Emulator Raw- An unstructured image format VDI- Virtual Disk Image VHD- Virtual Hard Disk VMDK- Virtual Machine Disk
Architecture	Enter the architecture.
Minimum Disk (GB)	Enter the minimum disk size required to boot the image. If you do not specify a size, the default is 0 (no minimum).
Minimum Ram (MB)	Enter the minimum RAM required to boot the image. If you do not specify a size, the default is 0 (no minimum).
Public	Select this check box if this is a public image. Leave unselected for a private image.

Table 19: Create an Image Fields (continued)

Field	Description
Protected	Select this check box for a protected image.

5. When you are finished, click **Create Image**.

Using Security Groups with Virtual Machines (Instances)

IN THIS SECTION

- [Security Groups Overview | 349](#)
- [Creating Security Groups and Adding Rules | 349](#)

Security Groups Overview

A **security group** is a container for security group rules. Security groups and security group rules allow administrators to specify the type of traffic that is allowed to pass through a port. When a virtual machine (VM) is created in a virtual network (VN), a security group can be associated with the VM when it is launched. If a security group is not specified, a port is associated with a default security group. The default security group allows both ingress and egress traffic. Security rules can be added to the default security group to change the traffic behavior.

Creating Security Groups and Adding Rules


A default security group is created for each project. You can add security rules to the default security group and you can create additional security groups and add rules to them. The security groups are then associated with a VM, when the VM is launched or at a later date.

To add rules to a security group:

1. From the OpenStack interface, click the **Project** tab, select **Access & Security**, and click the **Security Groups** tab.

Any existing security groups are listed under the **Security Groups** tab, including the default security group; see [Figure 57 on page 350](#).


Figure 57: Security Groups



2. Select the **default-security-group** and click **Edit Rules** in the **Actions** column.

The **Edit Security Group Rules** window is displayed; see [Figure 58 on page 351](#). Any rules already associated with the security group are listed.

Figure 58: Edit Security Group Rules



3. Click **Add Rule** to add a new rule; see [Figure 59 on page 351](#).

Figure 59: Add Rule




Table 20: Add Rule Fields

Column	Description
IP Protocol	Select the IP protocol to apply for this rule: TCP, UDP, ICMP.
From Port	Select the port from which traffic originates to apply this rule. For TCP and UDP, enter a single port or a range of ports. For ICMP rules, enter an ICMP type code.
To Port	The port to which traffic is destined that applies to this rule, using the same options as in the From Port field.
Source	Select the source of traffic to be allowed by this rule. Specify subnet—the CIDR IP address or address block of the inter-domain source of the traffic that applies to this rule, or you can choose security group as source. Selecting security group as source allows any other instance in that security group access to any other instance via this rule.

4. Click **Create Security Group** to create additional security groups.

The **Create Security Group** window is displayed; see [Figure 60 on page 352](#).

Each new security group has a unique 32-bit security group ID and an ACL is associated with the configured rules.

Figure 60: Create Security Group

Create Security Group

Name: SG1

Description: Security Group 1

Description: From here you can create a new security group


sg-1861

Cancel Create Security Group

5. When an instance is launched, there is an opportunity to associate a security group; see [Figure 61 on page 353](#).

In the **Security Groups** list, select the security group name to associate with the instance.

Figure 61: Associate Security Group at Launch Instance



6. You can verify that security groups are attached by viewing the **SgListReq** and **IntfReq** associated with the **agent.xml**.

Using Contrail Resources in Heat Templates

IN THIS SECTION

- [Using the Contrail Heat Template | 353](#)

Using the Contrail Heat Template

IN THIS SECTION

- [Introduction to Heat | 354](#)
- [Heat Architecture | 354](#)
- [Support for Heat Version 2 Resources | 354](#)
- [Heat Version 2 with Service Templates and Port Tuple Sample Workflow | 355](#)
- [Example: Creating a Service Template Using Heat | 355](#)

Heat is the orchestration engine of the OpenStack program. Heat enables launching multiple cloud applications based on templates that are comprised of text files.

Introduction to Heat

A Heat template describes the infrastructure for a cloud application, such as networks, servers, floating IP addresses, and the like, and can be used to manage the entire life cycle of that application.

When the application infrastructure changes, the Heat templates can be modified to automatically reflect those changes. Heat can also delete all application resources if the system is finished with an application.

Heat templates can record the relationships between resources, for example, which networks are connected by means of policy enforcements, and consequently call OpenStack REST APIs that create the necessary infrastructure, in the correct order, needed to launch the application managed by the Heat template.

Heat Architecture

Heat is implemented by means of Python applications, including the following:

- **heat-client**—The CLI tool that communicates with the **heat-api** application to run Heat APIs.
- **heat-api**—Provides an OpenStack native REST API that processes API requests by sending them to the Heat engine over remote procedure calls (RPCs).
- **heat-engine**—Responsible for orchestrating the launch of templates and providing events back to the API consumer.

Support for Heat Version 2 Resources

Starting with Contrail Release 3.0.2, Contrail Heat resources and templates are autogenerated from the Contrail schema, using Heat Version 2 resources. Contrail Release 3.0.2 is the minimum required version for using Heat with Contrail in 3.x releases. The Contrail Heat Version 2 resources are of the following hierarchy: **OS::ContrailV2::<ResourceName>**.

The generated resources and templates are part of the Contrail Python package, and are located in the following directory in the target installation:

/usr/lib/python2.7/dist-packages/vnc_api/gen/heat/

The **heat/** directory has the following subdirectories:

- **resources/**—Contains all the resources for the contrail-heat plugin, which runs in the context of the Heat engine service.
- **templates/**—Contains sample templates for each resource. Each sample template presents every possible parameter in the schema. Use the sample templates as a reference when you build up more complex templates for your network design.
- **env/**—Contains the environment for input to each template.

The following contains a list of all the generated plug-in resources that are supported by contrail-heat :

https://github.com/tungstenfabric/tf-heat-plugin/tree/master/contrail_heat/new_templates

Deprecation of Heat Version 1 Resources

Heat Version 1 resources within the hierarchy **OS::Contrail::<ResourceName>** are being deprecated, and you should not create new service templates using the Heat Version 1 templates.

Heat Version 2 with Service Templates and Port Tuple Sample Workflow

With Contrail service templates Version 2, the user can create ports and bind them to a virtual machine (VM)-based service instance, by means of a port-tuple object. All objects created with the Version 2 service template are directly visible to the Contrail Heat engine, and are directly managed by Heat.

The following shows the basic workflow steps for creating a port tuple and service instance that will be managed by Heat:

1. Create a service template. Select 2 in the Version field.
2. Create a service instance for the service template just created.
3. Create a port-tuple object.
4. Create ports, using Nova VM launch or without a VM launch.
5. Label each port as left, right, mgmt, and so on, and add the ports to the port-tuple object.

Use a unique label for each of the ports in a single port tuple. The labels named left and right are used for forwarding.

6. Link the port tuple to a service instance.
7. Launch the service instance.

Example: Creating a Service Template Using Heat

The following is an example of how to create a service template using Heat.

1. Define a template to create the service template.

```
service_template.yaml
heat_template_version: 2013-05-23
description: >
  HOT template to create a service template
parameters:
  name:
    type: string
    description: Name of service template
```

```

mode:
  type: string
  description: service mode
type:
  type: string
  description: service type
image:
  type: string
  description: Name of the image
flavor:
  type: string
  description: Flavor
service_interface_type_list:
  type: string
  description: List of interface types
shared_ip_list:
  type: string
  description: List of shared ip enabled- disabled
static_routes_list:
  type: string
  description: List of static routes enabled- disabled

resources:
  service_template:
    type: OS::ContrailV2::ServiceTemplate
    properties:
      name: { get_param: name }
      service_mode: { get_param: mode }
      service_type: { get_param: type }
      image_name: { get_param: image }
      flavor: { get_param: flavor }
      service_interface_type_list: { "Fn::Split" : [ ",", Ref:
service_interface_type_list ] }
      shared_ip_list: { "Fn::Split" : [ ",", Ref: shared_ip_list ] }
      static_routes_list: { "Fn::Split" : [ ",", Ref: static_routes_list ] }
    }
  outputs:
    service_template_fq_name:
      description: FQ name of the service template
      value: { get_attr: [ service_template, fq_name] }
}

```

2. Create an environment file to define the values to put in the variables in the template file.

```

service_template.env

parameters:

    name: contrail_svc_temp

    mode: transparent

    type: firewall

    image: cirros

    flavor: m1.tiny

    service_interface_type_list: management, left, right, other

    shared_ip_list: True, True, False, False

    static_routes_list: False, True, False, False

```

3. Create the Heat stack by launching the template and the environment file, using the following command:

```
heat stack create stack1 -f service_template.yaml -e service_template.env
```

OR use this command for recent versions of OpenStack

```
openstack stack create -e <env-file-name> -t <template-file-name> <stack-name>
```

SEE ALSO

| *Service Chain Version 2 with Port Tuple*

QoS Support in Contrail Networking

IN THIS SECTION

- [Quality of Service in Contrail | 358](#)
- [Configuring Network QoS Parameters | 365](#)

Quality of Service in Contrail

IN THIS SECTION

- [Overview: Quality of Service | 358](#)
- [Contrail QoS Model | 358](#)
- [Features of Fabric Interfaces | 359](#)
- [QoS Configuration Parameters for Provisioning | 359](#)
- [Configuring QoS in Contrail Networking Release 5.0 and Later | 359](#)
- [Queuing Implementation | 360](#)
- [Contrail QoS Configuration Objects | 360](#)
- [Example: Mapping Traffic to Forwarding Classes | 361](#)
- [QoS Configuration Object Marking on the Packet | 362](#)
- [Queuing | 363](#)

Overview: Quality of Service

Quality of service (QoS) in networking provides the ability to control reliability, bandwidth, latency, and other traffic management features. Network traffic can be marked with QoS bits (DSCP, 802.1p, and MPLS EXP) that intermediate network switches and routers can use to provide service guarantees.

Contrail QoS Model

The QoS model in Contrail Networking has the following features:

- All packet forwarding devices, such as vRouter and the gateway, combine to form a system.
- Interfaces to the system are the ports from which the system sends and receives packets, such as tap interfaces and physical ports.

- Fabric interfaces are where the overlay traffic is tunneled.
- QoS is applied at the ingress to the system, for example, upon traffic from the interfaces to the fabric.
- At egress, packets are stripped of their tunnel headers and sent to interface queues, based on the forwarding class. No marking from the outer packet to the inner packet is considered at this time.

Features of Fabric Interfaces

Fabric interfaces, unlike other interfaces, are always shared. Therefore, fabric interfaces are common property. Consequently, traffic classes and QoS marking on the fabric must be controlled by the system administrator. The administrator might choose to provision different classes of service on the fabric.

In Contrail, classes of service are determined by both of the following:

- Queueing on the fabric interface, including queues, scheduling of queues, and drop policies, and
- forwarding class, a method of marking that controls how packets are sent to the fabric, including marking and identifying which queue to use.

Tenants can define which forwarding class their traffic can use, deciding which packets use which forwarding class. The Contrail QoS configuration object has a mapping table, mapping the incoming DSCP or 802.1p value to the forwarding class mapping.

The QoS configuration can also be applied to a virtual network, an interface, or a network policy.

QoS Configuration Parameters for Provisioning

Configuring QoS in Contrail Networking Release 5.0 and Later

This section describes how to provision QoS in Contrail Networking release 5.0 and later.

1. Define the hardware queues and priority group in the **instances.yaml** file under the vrouter role as shown below.

```
nodeh5:
  ip: 10.xxx.xxx.109
  provider: bms
  roles:
    vrouter:
      VROUTER_GATEWAY: 192.168.1.45
      PRIORITY_ID: 0,1,2,3,4,5,6,7
      PRIORITY_BANDWIDTH: 0,10,0,20,0,30,0,40
      PRIORITY_SCHEDULING: strict,rr,strict,rr,strict,rr,strict,rr
      QOS_QUEUE_ID: 3,11,18,28,36,43,61,53
      QOS_LOGICAL_QUEUES: "[ 1, 6-10, 12-15];[40-46];[70-74, 75,
80-95];[115];[140-143, 145];[175];[245];[215]"
      QOS_DEF_HW_QUEUE: True
  openstack_compute:
```

2. In the already provisioned setup, define the QoS configuration in the `/etc/contrail/common_vrouter.env` file as shown in the following sample.

```
PRIORITY_ID=0,1,2,3,4,5,6,7
PRIORITY_BANDWIDTH=0,10,0,20,0,30,0,40
PRIORITY_SCHEDULING=strict,rr,strict,rr,strict,rr,strict,rr
QOS_QUEUE_ID=3,11,18,28,36,43,61,53
QOS_LOGICAL_QUEUES="[ 1, 6-10, 12-15];[40-46];[70-74, 75, 80-95];[115];[140-143, 145];[175];[245];[215]"
QOS_DEF_HW_QUEUE=True
```

3. Execute the **execute docker-compose up -d** under `/etc/contrail/vrouter/` command.

Queuing Implementation

The vRouter provides the infrastructure to use queues supplied by the network interface, a method that is also called hardware queueing. Network interface cards (NICs) that implement hardware queueing have their own set of scheduling algorithms associated with the queues. The Contrail implementation is designed to work with most NICs, however, the method is tested only on an Intel-based 10G NIC, also called Niantic.

Contrail QoS Configuration Objects

Contrail QoS configuration objects include the:

- forwarding class
- QoS configuration object (**qos-config**)

The forwarding class object specifies parameters for marking and queuing, including:

- The DSCP, 802.1p, and MPLS EXP values to be written on packets.
- The queue index to be used for the packet.

The QoS configuration object specifies a mapping from DSCP, 802.1p, and MPLS EXP values to the corresponding forwarding class.

The QoS configuration has an option to specify the default forwarding class ID to use to select the forwarding class for all unspecified DSCP, 802.1p, and MPLS EXP values.

If the default forwarding class ID is not specified by the user, it defaults to the forwarding class with ID 0.


Processing of QoS marked packets to look up the corresponding forwarding class to be applied works as follows:

- For an IP packet, the DSCP map is used .
- For a Layer 2 packet, the 802.1p map is used.

- For an MPLS-tunneled packet with MPLS EXP values specified, the EXP bit value is used with the MPLS EXP map.
- If the QoS configuration is untrusted, only the default forwarding class is specified, and all incoming values of the DSCP, 802.1p, and EXP bits in the packet are mapped to the same default forwarding class.


Figure 62 on page 361 shows the processing of QoS packets.

Figure 62: Processing of QoS Packets



A virtual machine interface, virtual network, and network policy can refer to the QoS configuration object. The QoS configuration object can be specified on the vhost so that underlay traffic can also be subjected to marking and queuing. See Figure 63 on page 361.

Figure 63: Referring to the QoS Object



Example: Mapping Traffic to Forwarding Classes

This example shows how traffic forwarding classes are defined and how the QoS configuration object is defined to map the QoS bits to forwarding classes.

Table 21 on page 362 shows two forwarding class objects defined. FC1 marks the traffic with high priority values and queues it to Queue 0. FC2 marks the traffic as best effort and queues the traffic to Queue 1.

Table 21: Forwarding Class Mapping

Name	ID	DSCP	802.1p	MPLS EXP	Queue
FC1	1	10	7	7	0
FC2	2	38	0	0	1

In [Table 22 on page 362](#), the QoS configuration object DSCP values of 10, 18, and 26 are mapped to a forwarding class with ID 1, which is forwarding class FC1. All other IP packets are mapped to the forwarding class with ID 2, which is FC2. All traffic with an 802.1p value of 6 or 7 are mapped to forwarding class FC1, and the remaining traffic is mapped to FC2.

Table 22: QoS Configuration Object Mapping

DSCP	Forwarding Class ID	802.1p	Forwarding Class ID	MPLS EXP	Forwarding Class ID
10	1	6	1	5	1
18	1	7	1	7	1
26	1	*	2	*	1
*	2				

QoS Configuration Object Marking on the Packet

IN THIS SECTION

- [Traffic Originated by a Virtual Machine Interface | 363](#)
- [Traffic Destined to a Virtual Machine Interface | 363](#)
- [Traffic from a vhost Interface | 363](#)
- [Traffic from fabric interface | 363](#)
- [QoS Configuration Priority by Level | 363](#)

The following sections describes how QoS configuration object marking is handled in various circumstances.

Traffic Originated by a Virtual Machine Interface

- If a VM interface sends an IP packet to another VM in a remote compute node, the DSCP value in the IP header is used to look into the qos-config table, and the tunnel header is marked with DSCP, 802.1p, and MPLS EXP bits as specified by the forwarding class.
- If a VM sends a Layer 2 non-IP packet with an 802.1p value, the 802.1p value is used to look into the qos-config table, and the corresponding forwarding class DSCP, 802.1p, and MPLS EXP value is written to the tunnel header.
- If a VM sends an IP packet to a VM in same compute node, the packet headers are not changed while forwarding. The original packet remains unchanged.

Traffic Destined to a Virtual Machine Interface

For traffic destined to a VMI, if a tunneled packet is received, the tunnel headers are stripped off and the packet is sent to the interface. No marking is done from the outer packet to inner packet.

Traffic from a vhost Interface

The QoS configuration can be applied on IP traffic coming from a vhost interface. The DSCP value in the packet is used to look into the qos-config object specified on the vhost, and the corresponding forwarding class DSCP and 802.1p values are overwritten on the packet.

Traffic from fabric interface

The QoS configuration can be applied while receiving the packet on an Ethernet interface of a compute node, and the corresponding forwarding class DSCP and 802.1p values are overwritten on the packet.

QoS Configuration Priority by Level

The QoS configuration can be specified at different levels.

The levels that can be configured with QoS and their order of priority:

1. in policy
2. on **virtual-network**
3. on **virtual-machine-interface**

Queuing

IN THIS SECTION

- [Queue Selection in Datapath | 364](#)
- [Hardware Queueing in Linux kernel based vRouter | 364](#)
- [Parameters for QoS Scheduling Configuration | 364](#)

Contrail Networking supports QoS. These sections provide an overview of the queuing features available in Contrail Networking.

The queue to which a packet is sent is specified by the forwarding class.

Queue Selection in Datapath

In vRouter, in the data path, the forwarding class number specifies the actual physical hardware queue to which the packet needs to be sent, not to a logical selection as in other parts of Contrail. There is a mapping table in the vRouter configuration file, to translate the physical queue number from the logical queue number.

Hardware Queueing in Linux kernel based vRouter

If Xmit-Packet-Steering (XPS) is enabled, the kernel chooses the queue, from those available in a list of queues. If the kernel selects the queue, packets will not be sent to the vRouter-specified queue.

To disable this mapping:

- have a kernel without CONFIG_XPS option
- write zeros to the mapping file in `/sys/class/net//queues/tx-X/xps_cpus`.

When this mapping is disabled, the kernel will send packets to the specific hardware queue.

To verify:

See individual queue statistics in the output of 'ethtool -S ' command.

Parameters for QoS Scheduling Configuration

The following shows sample scheduling configuration for hardware queues on the compute node.

The priority group ID and the corresponding scheduling algorithm and bandwidth to be used by the priority group can be configured.

Possible values for the scheduling algorithm include:

- strict
- rr (round-robin)

When round-robin scheduling is used, the percentage of total hardware queue bandwidth that can be used by the priority group is specified in the bandwidth parameter.

The following configuration and provisioning is applicable only for compute nodes running Niantic NICs and running kernel based vrouter.

```
qos_niantic = {
    'compute1': [
        { 'priority_id': '1', 'scheduling': 'strict', 'bandwidth':
        '0'},
```

```

        { 'priority_id': '2', 'scheduling': 'rr', 'bandwidth':
'20'},
        { 'priority_id': '3', 'scheduling': 'rr', 'bandwidth': '10'}
    ],
    'compute2' :[
        { 'priority_id': '1', 'scheduling': 'strict', 'bandwidth':
'0'},
        { 'priority_id': '1', 'scheduling': 'rr', 'bandwidth': '30'}
    ]
}

```

SEE ALSO

[Configuring Network QoS Parameters | 365](#)

<https://github.com/Juniper/contrail-controller/wiki/QoS>

Configuring Network QoS Parameters

IN THIS SECTION

- [Overview | 365](#)
- [QoS Configuration Examples | 366](#)
- [Limitations | 367](#)

Overview

You can use the OpenStack Nova command-line interface (CLI) to specify a quality of service (QoS) setting for a virtual machine's network interface, by setting the **quota** of a Nova flavor. Any virtual machine created with that Nova flavor will inherit all of the specified QoS settings. Additionally, if the virtual machine that was created with the QoS settings has multiple interfaces in different virtual networks, the same QoS settings will be applied to all of the network interfaces associated with the virtual machine. The QoS settings can be specified in unidirectional or bidirectional mode.

The **quota** driver in Neutron converts QoS parameters into **libvirt** network settings of the virtual machine.

The QoS parameters available in the quota driver only cover rate limiting the network interface. There are no specifications available for policy-based QoS at this time.

QoS Configuration Examples

Although the QoS setting can be specified in quota by using either Horizon or CLI, quota creation using CLI is more robust and stable, therefore, creating by CLI is the recommended method.

Example

CLI for Nova flavor has the following format:

```
nova flavor-key <flavor_name> set quota:vif_<direction> _<param_name> = value
```

where:

<flavor_name> is the name of an existing Nova flavor.

vif_<direction>_<param_name> is the inbound or outbound QoS data name.

QoS vif types include the following:

- **vif_inbound_average** lets you specify the average rate of inbound (receive) traffic, in kilobytes/sec.
- **vif_outbound_average** lets you specify the average rate of outbound (transmit) traffic, in kilobytes/sec.
- Optional: **vif_inbound_peak** and **vif_outbound_peak** specify the maximum rate of inbound and outbound traffic, respectively, in kilobytes/sec.
- Optional: **vif_inbound_burst** and **vif_outbound_peak** specify the amount of kilobytes that can be received or transmitted, respectively, in a single burst at the peak rate.

Details for various QoS parameters for **libvirt** can be found at <http://libvirt.org/formatnetwork.html>.

The following example shows an inbound average of 800 kilobytes/sec, a peak of 1000 kilobytes/sec, and a burst amount of 30 kilobytes.

```
nova flavor-key m1.small set quota:vif_inbound_average=800
nova flavor-key m1.small set quota:vif_inbound_peak=1000
nova flavor-key m1.small set quota:vif_inbound_burst=30
```

The following is an example of specified outbound parameters:

```
nova flavor-key m1.small set quota:vif_outbound_average=800
nova flavor-key m1.small set quota:vif_outbound_peak=1000
nova flavor-key m1.small set quota:vif_outbound_burst=30
```

After the Nova flavor is configured for QoS, a virtual machine instance can be created, using either Horizon or CLI. The instance will have network settings corresponding to the nova flavor-key, as in the following:

```
<interface type="ethernet">
  <mac address="02:a3:a0:87:7f:61"/>
  <model type="virtio"/>
  <script path=""/>
  <target dev="tapa3a0877f-61"/>
  <bandwidth>
    <inbound average="800" peak="1000" burst="30"/>
    <outbound average="800" peak="1000" burst="30"/>
  </bandwidth>
</interface>
```

Limitations

- The stock **libvirt** does not support rate limiting of **ethernet** interface types. Consequently, settings like those in the example for the guest interface will not result in any **tc qdisc** settings for the corresponding tap device in the host.
- The **nova flavor-key rxtx_factor** takes a float as an input and acts as a scaling factor for receive (inbound) and transmit (outbound) throughputs. This key is only available to Neutron extensions (private extensions). The Contrail Neutron plugin doesn't implement this private extension. Consequently, setting the **nova flavor-key rxtx_factor** will not have any effect on the QoS setting of the network interface(s) of any virtual machine created with that nova flavor.
- The outbound rate limits of a virtual machine interface are not strictly achieved. The outbound throughput of a virtual machine network interface is always less than the average outbound limit specified in the virtual machine's libvirt configuration file. The same behavior is also seen when using a Linux bridge.

Load Balancers

IN THIS SECTION

- [Using Load Balancers in Contrail | 368](#)
- [Support for OpenStack LBaaS | 381](#)
- [Configuring Load Balancing as a Service in Contrail | 385](#)

Using Load Balancers in Contrail

IN THIS SECTION

- [Invoking LBaaS Drivers | 368](#)
- [Benefits of Creating Configuration Objects | 368](#)
- [Using a Service Appliance Set as the LBaaS Provider | 369](#)
- [Understanding the Load Balancer Agent | 371](#)
- [F5 Networks Load Balancer Integration in Contrail | 372](#)
- [Example: Creating a Load Balancer | 375](#)
- [Using the Avi Networks Load Balancer for Contrail | 376](#)

As of Contrail Release 3.0, load balancer LBaaS features are available. This topic includes:

Invoking LBaaS Drivers

The provider field specified in the pool configuration determines which load balancer drivers are selected. The load balancer driver selected is responsible for configuring the external hardware or virtual machine load balancer.

Supported load balancer drivers include:

- HAProxy
- A10 Networks
- F5 Networks
- Avi Networks

Benefits of Creating Configuration Objects

Starting with Contrail 3.0, the Neutron LBaaS plugin creates required configuration objects (such as pool, VIP, members, and monitor) in the Contrail API server, instead of within the Neutron plugin context, as in previous releases.

This method of configuration has the following benefits:


- Configuration objects can be created in multiple ways: from Neutron, from virtual controller APIs, or from the Contrail UI.
- The load balancer driver can make inline calls, such as REST or SUDS, to configure the external load balancer device.

- The load balancer driver can use Contrail service monitor infrastructure, such as database, logging, and API server.

NOTE: The *Neutron* LBaaS plugin is not supported in OpenStack Train release.

Figure 64 on page 369 provides an overview of the Contrail LBaaS components.

Figure 64: Contrail LBaaS components with neutron-lbaas



Using a Service Appliance Set as the LBaaS Provider

In OpenStack Neutron, the load balancer provider is statically configured in **neutron.conf**, which requires restart of the Neutron server when configuring a new provider. The following is an example of the service provider configuration in **neutron.conf**.

```
[service_providers]
service_provider =
```

```
LOADBALANCER:Opencontrail:neutron_plugin_contrail.plugins.opencontrail.
loadbalancer.driver.OpencontrailLoadbalancerDriver:default
```

In Contrail Release 3.0 and greater, the Neutron LBaaS provider is configured by using the object **service-appliance-set**. All of the configuration parameters of the LBaaS driver are populated to the **service-appliance-set** object and passed to the driver.

During initialization, the service monitor creates a default service appliance set with a default LBaaS provider, which uses an HAProxy-based load balancer. The service appliance set consists of individual service appliances for load balancing the traffic. The service appliances can be physical devices or virtual machines.

Sample Configuration: Service Appliance Set

The following is a sample configuration of the service appliance set for the LBaaS provider:

```
{
  "service-appliance-set": {
    "fq_name": [
      "default-global-system-config",
      "f5"
    ],
    "service_appliance_driver":
      "svc_monitor.services.loadbalancer.drivers.f5.f5_driver.OpencontrailF5LoadbalancerDriver",

    "parent_type": "global-system-config",
    "service_appliance_set_properties": {
      "key_value_pair": [
        {
          "key": "sync_mode",
          "value": "replication"
        },
        {
          "key": "global_routed_mode",
          "value": "True"
        }
      ]
    },
    "name": "f5"
  }
}
```

Sample Configuration: Single Service Appliance

The following is a sample configuration of a single service appliance:

```
{
  "service-appliance": {
    "fq_name": [
      "default-global-system-config",
      "f5",
      "bigip"
    ],
    "parent_type": "service-appliance-set",
    "service_appliance_ip_address": "<ip address>",
    "service_appliance_user_credentials": {
      "username": "admin",
      "password": "<password>"
    },
    "name": "bigip"
  }
}
```

Understanding the Load Balancer Agent

The load balancer agent is a module in the service monitor. The service monitor listens on the RabbitMQ configuration messaging queue (**vnc_config.object-update**) to get configuration objects. The dependency tracker triggers changes to all related objects, based on configuration updates.

The dependency tracker is informed to notify the pool object whenever the VIP, member, or health monitor object is modified.

Whenever there is an update to the pool object, either directly due to a pool update or due to a dependency update, the load balancer agent in the service monitor is notified.

The load balancer agent module handles the following:

- Loading and unloading LBaaS driver-based service appliance set configuration.
- Providing the abstract driver class for the load balancer driver.
- Invoking the LBaaS driver.
- Load balancer-related configuration.

F5 Networks Load Balancer Integration in Contrail

IN THIS SECTION

- [F5 Load Balancer Global Routed Mode | 372](#)
- [Initial Configuration on an F5 Device | 374](#)
- [Initial Configuration on an MX Series Device Used as DC Gateway | 374](#)

This section details use of the F5 load balancer driver with Contrail.

Contrail Release 3.0 implements an LBaaS driver that supports a physical or virtual F5 Networks load balancer, using the abstract load balancer driver class, **ContrailLoadBalancerAbstractDriver**.

This driver is invoked from the load balancer agent of the **contrail-svc-monitor**. The driver makes a BIG-IP interface call to configure the F5 Networks device. All of the configuration parameters used to tune the driver are configured in the **service-appliance-set** object and passed to the driver by the load balancer agent while loading the driver.

The F5 load balancer driver uses the BIG-IP interface version V1.0.6, which is a Python package extracted from the load balancer plugin provided by F5 Networks. The driver uses either a SOAP API or a REST API.

F5 Load Balancer Global Routed Mode

The F5 load balancer driver is programmed in **global routed** mode using a property of the **service-appliance-set**.

This section describes the features and requirements of the F5 load balancer driver configured in global routed mode.

The following are features of the global routed mode.

- All virtual IP addresses (VIPs) are assumed to be routable from clients and all members are routable from the F5 device.
- All access to and from the F5 device is assumed to be globally routed, with no segregation between tenant services on the F5 device. Consequently, do NOT configure overlapping addresses across tenants and networks.
- The F5 device can be attached to the corporate network or to the IP fabric.

The following are requirements to support global routed mode of an F5 device used with LBaaS:

- The entire configuration of the F5 device for Layer 2 and Layer 3 is preprovisioned.
- All tenant networks and all IP fabrics are in the same namespace as the corporate network.
- All VIPs are in the same namespace as the tenant and corporate networks.

Traffic Flow in Global Routed Mode

This section describes and illustrates the behavior of traffic flow in global routed mode.

The information in this section is based on a model that includes the following network topology:

Corporate Network --- DC Gateway (MX device) --- IP Fabric --- Compute nodes

The Corporate Network, the IP Fabric and all tenant networks use IP addresses from a single namespace, there is no overlap of the addresses in the networks. The F5 devices can be attached to the Corporate Network or to the IP Fabric, and are configured to use the global routed mode.

The role of the MX Series device is to route post-proxy traffic, coming from the F5 device in the underlay, to the pool members in the overlay. In the reverse direction, the MX device takes traffic coming from the pool members in the overlay and routes it back to the F5 device in the underlay.

The MX device is preprovisioned with the following:

- VRF connected to pool network 2
- ability to route traffic from inet.0 to the pool network

The MX routes the traffic from inet.0 to public VRF and sends traffic to the compute node where the pool member is instantiated.

The F5 device is preprovisioned with the following:


- publish route to attract VIP traffic
- pool network subnet route that points to the MX device

The F5 device is responsible for attracting traffic destined to all the VIPs, by advertising a subnet route that covers all VIPs using IGP.

The F5 device load balances among different pool members and sends traffic to the chosen member.

[Figure 65 on page 374](#) shows the traffic flow in global routed mode.

Figure 65: Global Routed Traffic Flow



A similar result can also be achieved on the switch to which the F5 is attached, by publishing the VIP subnet in IGP and using a static route to point the VIP traffic to the F5 device.

The MX should attract the reverse traffic from the pool members going back to the F5.

Routing Traffic to Pool Members

For post load balancing traffic going from the F5 device to the pool members, the MX Series device needs to attract traffic for all the tenant networks.

Routing Reverse Traffic from Pool Members to the F5 Device

The MX should attract the reverse traffic from the pool members going back to the F5.

Initial Configuration on an F5 Device

- The operator is responsible for ensuring that the F5 device attracts traffic to all VIP subnets by injecting the route for the VIP subnet into IGP. Alternately, the switch to which F5 is connected can advertise the VIP subnet route and use the static route to send VIP traffic to the F5 device.
- In the global routed mode, the F5 uses AutoMap SNAT for all VIP traffic.

Initial Configuration on an MX Series Device Used as DC Gateway

- The operator must identify a super-net that contains all tenant network subnets (pool members across multiple pools) and advertise its route into corporate and fabric networks, using IGP (preferred) or static routes.
- The operator must add a static route for the super-net into inet.0 with a next-hop of public.inet.0.

- The operator must create a public VRF and get its default route imported into the VRF. This is to attract the return traffic from pool members to the F5 device (VIP destination).

Configuration on MX Device for Each Pool Member

- For each member virtual network, the operator adds a policy to connect the member pool virtual network to the public virtual network.
- As new member virtual networks are connected to the public virtual network by policy, corresponding targets are imported by the public VRF on MX. The Contrail Device Manager generates the configuration of import, export targets for public VRF on the MX device.
- The operator must ensure that security group rules for the member virtual network ports allow traffic coming from the F5 device.

Example: Creating a Load Balancer

Use the following steps to create a load balancer in Contrail Release 3.0 and greater.

1. To configure a service appliance set, use the script in **/opt/contrail/utils** to create a load balancer provider. With the script, you specify the driver and name of the selected provider. Additional configuration can be performed using the key-value pair property configuration.

```
/opt/contrail/utils/service_appliance_set.py --api_server_ip <ip address>--api_server_port 8082 --oper
add --admin_user admin --admin_password <password> --admin_tenant_name admin --name f5
--driver "svc_monitor.services.loadbalancer.drivers.f5.f5_driver.OpencontrailF5LoadbalancerDriver"
--properties '{"use_snat": "True", "num_snat": "1", "global_routed_mode": "True", "sync_mode":
"replication", "vip_vlan": "trial2"}'
```

2. Add the actual device information of the load balancer.

```
/opt/contrail/utils/service_appliance.py --api_server_ip <ip address>--api_server_port 8082 --oper
add --admin_user admin --admin_password <password> --admin_tenant_name admin --name bigip
--service_appliance_set f5 --device_ip 10.204.216.113 --user_credential '{"user": "admin", "password":
"<password>"}'
```

3. Refer to the load balancer provider while configuring the pool.

```
neutron lb-pool-create --lb-method ROUND_ROBIN --name web_service --protocol HTTP --provider
"f5" --subnet-id <subnet id>
```

4. Add members to the load balancer pool. Both bare metal webserver and overlay webserver are allowed as pool members. The F5 device can load balance the traffic among all pool members.

```
neutron lb-member-create --address <ip address>--protocol-port 8080 --weight 3 web_service
neutron lb-member-create --address <ip address> --protocol-port 8080 --weight 2 web_service
```

5. Create a VIP for the load balancer pool.

```
neutron lb-vip-create --name httpserver --protocol-port 80 --protocol HTTP web_service --subnet-id
<subnet id>
```

6. Create the health monitor and associate it with the load balancer pool.

```
neutron lb-healthmonitor-create --delay 3 --type HTTP --max-retries 3 --timeout 3
neutron lb-healthmonitor-associate <nnnnn-nnnnn-nnnn-> web_service
```

Using the Avi Networks Load Balancer for Contrail

If you are using the Avi LBaaS driver in an OpenStack Contrail environment, there are two possible modes that are mutually-exclusive. The Avi Vantage cloud configuration is exactly the same in both modes:

- Neutron-based Avi LBaaS driver
In this mode, the Avi LBaaS driver derives from Neutron and resides in the Neutron server process. This mode enables coexistence of multiple Neutron LBaaS providers.
- Contrail-based Avi LBaaS driver
In this mode, the Avi LBaaS driver derives from Contrail and resides in the service-monitor process. This mode enables coexistence of multiple Contrail LBaaS providers.

NOTE: In a Contrail environment, you cannot have a mix of Contrail LBaaS and Neutron LBaaS. You must select a mode that is compatible with the current environment.

Installing the Avi LBaaS Neutron Driver

Use the following procedure to install the Avi Networks LBaaS load balancer driver for the Neutron server for Contrail.

The following steps are performed on the Neutron server host.

1. Determine the installed version of the Contrail Neutron plugin.

```
$ contrail-version neutron-plugin-contrail
Package Version
-----
neutron-plugin-contrail 3.0.2.0-51
```

2. Adjust the **neutron.conf** database connection URL.

```
$ vi /etc/neutron/neutron.conf
# if using mysql
connection = mysql+pymysql://neutron:c0ntrail123@127.0.0.1/neutron
```


3. Populate and upgrade the Neutron database schema.

```
# to upgrade to head
$ neutron-db-manage upgrade head
# to upgrade to a specific version
$ neutron-db-manage --config-file /etc/neutron/neutron.conf upgrade liberty
```

4. Drop foreign key constraints.

```
# obtain current mysql token
$ cat /etc/contrail/mysql.token
fabe17d9dd5ae798f7ea

$ mysql -u root -p
Enter password: fabe17d9dd5ae798f7ea

mysql> use neutron;

mysql> show create table vips;
# CONSTRAINT `vips_ibfk_1` FOREIGN KEY (`port_id`) REFERENCES `ports` (`id`) -
ports table is not used by Contrail
mysql> alter table vips drop FOREIGN KEY vips_ibfk_1;

mysql> show create table lbaas_loadbalancers;
# CONSTRAINT `fk_lbaas_loadbalancers_ports_id` FOREIGN KEY (`vip_port_id`)
REFERENCES `ports` (`id`)
mysql> alter table lbaas_loadbalancers drop FOREIGN KEY
fk_lbaas_loadbalancers_ports_id;
```

5. To install the Avi LBaaS plugin, continue with steps from the readme file that downloads with the Avi LBaaS software. You can perform either a local installation or a manual installation. The following are sample installation steps.

- For a local installation:

```
# LBaaS v1 driver
$ ./install.sh --aname avi_adc --aip

<controller_ip|controller_vip>
--auser

--apass
```

```
# LBaaS v2 driver
$ ./install.sh --aname avi_adc_v2 --aip
    <controller_ip|controller_vip>
    --auser

    --apass

    --v2
```

- For a manual installation:

```
# LBaaS v1 driver
$ vi /etc/neutron/neutron.conf
#service_plugins =
neutron_plugin_contrail.plugins.opencontrail.loadbalancer.plugin.LoadBalancerPlugin
service_plugins = neutron_lbaas.services.loadbalancer.plugin.LoadBalancerPlugin
[service_providers]
service_provider =
LOADBALANCER:Avi_ADC:neutron_lbaas.services.loadbalancer.drivers.avi.avi_driver.AviLbaaSDriver

[avi_adc]
address=10.1.11.4
user=admin
password=avi123
cloud=jcos

# LBaaS v2 driver
$ vi /etc/neutron/neutron.conf
#service_plugins =
neutron_plugin_contrail.plugins.opencontrail.loadbalancer.plugin.LoadBalancerPlugin
service_plugins =
neutron_lbaas.services.loadbalancer.plugin.LoadBalancerPluginv2
[service_providers]
service_provider =
LOADBALANCERV2:avi_adc_v2:neutron_lbaas.drivers.avi.driver.AviDriver

[avi_adc_v2]
controller_ip=10.1.11.3
username=admin
password=avi123

$ service neutron-server restart
$ neutron service-provider-list
```

Installing the Avi LBaaS Contrail Driver

Use the following procedure to install the Avi Networks LBaaS load balancer driver for Contrail.

The following steps are performed on the Contrail **api-server** host.

1. Determine the installed version of the Contrail Neutron plugin.

```
$ contrail-version neutron-plugin-contrail
Package Version
-----
neutron-plugin-contrail 3.0.2.0-51
```

2. Install the Avi driver.

```
# LBaaS v2 driver
$ ./install.sh --aname ocavi_adc_v2 --aip

<controller_ip|controller_vip>
--auser

--apass

--v2 --no-restart --no-confmodify
```

3. Set up the service appliance set.

NOTE: If **neutron_lbaas** doesn't exist on the **api-server** node, adjust the driver path to the correct path location for **neutron_lbaas**.

```
$ /opt/contrail/utils/service_appliance_set.py --api_server_ip 10.xx.xx.100 --api_server_port 8082
--oper add --admin_user admin --admin_password <password> --admin_tenant_name admin --name
ocavi_adc_v2 --driver "neutron_lbaas.drivers.avi.avi_ocdriver.OpencontrailAviLoadbalancerDriver"
--properties '{"address": "10.1.xx.3", "user": "admin", "password": "avi123", "cloud": "Default-Cloud"}'
```

4. To delete the service appliance set.

```
$ /opt/contrail/utils/service_appliance_set.py --api_server_ip 10.xx.xx.100 --api_server_port 8082
--oper del --admin_user admin --admin_password <password> --admin_tenant_name admin --name
ocavi_adc_v2
```

Configuring the Avi Controller

1. If OpenStack endpoints are private IPs and Contrail provides a public front-end IP to those endpoints, use iptables to DNAT. On the AviController only, perform iptable NAT to reach the private IPs.

\$ iptables -t nat -I OUTPUT --dest 17x.xx.xx.50 -j DNAT --to-dest 10.xx.xx.100

2. To configure the Avi controller during cloud configuration, select the “Integration with Contrail” checkbox and provide the endpoint URL of the Contrail VNC api-server. Use the Keystone credentials from the OpenStack configuration to authenticate with the api-server service.

```
: > show cloud jcos
```

Field	Value
uuid	cloud-104bb7e6-a9d2-4b34-a4c5-d94be659bb91
name	jcos
vtype	CLOUD_OPENSTACK
openstack_configuration	
username	admin
admin_tenant	demo
keystone_host	17x.xx.xx.50
mgmt_network_name	mgmtnw
privilege	WRITE_ACCESS
use_keystone_auth	True
region	RegionOne
hypervisor	KVM
tenant_se	True
import_keystone_tenants	True
anti_affinity	True
port_security	False
security_groups	True
allowed_address_pairs	True
free_floatingips	True
img_format	OS_IMG_FMT_AUTO
use_admin_url	True
use_internal_endpoints	False
config_drive	True
insecure	True
intf_sec_ips	False
external_networks	False
neutron_rbac	True
nuage_port	8443
contrail_endpoint	http://10.10.10.100:8082
apic_mode	False

dhcp_enabled	True	
mtu	1500 bytes	
prefer_static_routes	False	
enable_vip_static_routes	False	
license_type	LIC_CORES	
tenant_ref	admin	
+-----+-----+		

SEE ALSO

Configuring Load Balancing as a Service in Contrail 385
Support for OpenStack LBaaS 381
Installing OpenStack Octavia LBaaS with RHOSP in Contrail Networking 334
Installing OpenStack Octavia LBaaS with Juju Charms in Contrail Networking 455

Support for OpenStack LBaaS

IN THIS SECTION

- [OpenStack Neutron LBaaS Version 2.0 | 381](#)
- [OpenStack Octavia LBaaS | 383](#)

Starting with Release 3.1, Contrail provides support for the OpenStack Load Balancer as a Service (LBaaS) Version 2.0 APIs in the Liberty release of OpenStack.

OpenStack Neutron LBaaS Version 2.0

IN THIS SECTION

- [Platform Support | 382](#)
- [Using OpenStack LBaaS Version 2.0 | 382](#)
- [Support for Multiple Certificates per Listener | 382](#)
- [Neutron Load-Balancer Creation | 383](#)

Platform Support

Table 23 on page 382 shows which Contrail with OpenStack release combinations support which version of OpenStack LBaaS APIs.

Table 23: Contrail OpenStack Platform Support for LBaaS Versions

Contrail OpenStack Platform	LBaaS Support
Contrail-3.1-Liberty (and subsequent OS releases)	Only LBaaS v2 is supported.
Contrail-3.0-Liberty (and subsequent OS releases)	LBaaS v1 is default. LBaaS v2 is Beta.
<Contrail-any-release>-Kilo (and previous OS releases)	Only LBaaS v1 is supported.

Using OpenStack LBaaS Version 2.0

The OpenStack LBaaS Version 2.0 extension enables tenants to manage load balancers for VMs, for example, load-balancing client traffic from a network to application services, such as VMs, on the same network. The LBaaS Version 2.0 extension is used to create and manage load balancers, listeners, pools, members of a pool, and health monitors, and to view the status of a resource.

For LBaaS v2.0, the Contrail controller aggregates the configuration by provider. For example, if **haproxy** is the provider, the controller generates the configuration for **haproxy** and eliminates the need to send all of the load-balancer resources to the **vrouter-agent**; only the generated configuration is sent, as part of the service instance.

For more information about OpenStack v2.0 APIs, refer to the section *LBaaS 2.0 (STABLE) (lbaas, loadbalancers, listeners, health_monitors, pools, members)*, at <http://developer.openstack.org/api-ref-networking-v2-ext.html>.

LBaaS v2.0 also allows users to listen to multiple ports for the same virtual IP, by decoupling the virtual IP address from the port.

The object model has the following resources:

- Load balancer—Holds the virtual IP address
- Listeners—One or many listeners with different ports, protocols, and so on
- Pools
- Members
- Health monitors

Support for Multiple Certificates per Listener

Multiple certificates per listener are supported, with OpenStack Barbican as the storage for certificates. OpenStack Barbican is a REST API designed for the secure storage, provisioning, and management of secrets such as passwords, encryption keys, and X.509 certificates.

The following is an example CLI to store certificates in Barbican:

```
- barbican --os-identity-api-version 2.0 secret store --payload-content-type='text/plain' --name='certificate'
--payload="$(cat server.crt)"
```

For more information about OpenStack Barbican, see: <https://wiki.openstack.org/wiki/Barbican>.

Neutron Load-Balancer Creation

NOTE: The following procedure is written using the *Neutron* LBaaS plugin v1.0. Starting with the OpenStack Train release, *neutron-lbaas* is replaced by *Octavia*.

The following is an example of Neutron load-balancer creation:

```
- neutron net-create private-net

- neutron subnet-create --name private-subnet private-net 10.30.30.0/24

- neutron lbaas-loadbalancer-create $(neutron subnet-list | awk '/ private-subnet /
{print $2}') --name lb1

- neutron lbaas-listener-create --loadbalancer lb1 --protocol-port 443 --protocol
TERMINATED_HTTPS --name listener1 --default-tls-container=$(barbican
--os-identity-api-version 2.0 container list | awk '/ tls_container / {print $2}')

- neutron lbaas-pool-create --name pool1 --protocol HTTP --listener listener1
--lb-algorithm ROUND_ROBIN

- neutron lbaas-member-create --subnet private-subnet --address 30.30.30.10
--protocol-port 80 mypool

- neutron lbaas-member-create --subnet private-subnet --address 30.30.30.11
--protocol-port 80 mypool
```

OpenStack Octavia LBaaS

IN THIS SECTION

- [Using Octavia Load-Balancer | 384](#)
- [Octavia Load-Balancer Creation | 384](#)

Using Octavia Load-Balancer

Contrail Networking Release 2005 supports Octavia as LBaaS. The deployment supports RHOSP and Juju platforms.

With Octavia as LBaaS, Contrail Networking is only maintaining network connectivity and is not involved in any load balancing functions.

For each OpenStack load balancer creation, Octavia launches a VM known as *amphora* VM. The VM starts the HAPROXY when listener is created for the load balancer in OpenStack. Whenever the load balancer gets updated in OpenStack, *amphora* VM updates the running HAPROXY configuration. The *amphora* VM is deleted on deleting the load balancer.

Contrail Networking provides connectivity to *amphora* VM interfaces. *Amphora* VM has two interfaces; one for management and the other for data. The management interface is used by the Octavia services for the management communication. Since, Octavia services are running in the underlay network and *amphora* VM is running in the overlay network, SDN gateway is needed to reach the overlay network. The data interface is used for load balancing the traffic.

If the load balancer service is exposed to public, you must create the load balancer VIP in the public subnet. The load balancer members can be in the public or private subnet.

You must create network policy between public network and private network if the load balancer members are in the private network.

Octavia Load-Balancer Creation

The following is an example of Octavia load-balancer creation:

```
openstack loadbalancer listener create --protocol HTTP --protocol-port 80 --name
listener1 lb1
openstack loadbalancer show lb1 # Wait for the provisioning_status to be ACTIVE.
openstack loadbalancer pool create --lb-algorithm ROUND_ROBIN --listener listener1
--protocol HTTP --name pool1
openstack loadbalancer healthmonitor create --delay 5 --timeout 2 --max-retries 1
--type HTTP pool1
openstack loadbalancer member create --subnet-id private --address 10.10.10.50
--protocol-port 80 pool1
openstack loadbalancer member create --subnet-id private --address 10.10.10.51
--protocol-port 80 pool1
```

SEE ALSO

OpenStack Barbican: <https://wiki.openstack.org/wiki/Barbican>

OpenStack LBaaS v2.0 APIs: *LBaaS 2.0 (STABLE) (lbaas, loadbalancers, listeners, health_monitors, pools, members)*, at <http://developer.openstack.org/api-ref-networking-v2-ext.html>.

OpenStack Basic Load Balancing Cookbook at <https://docs.openstack.org/octavia/queens/user/guides/basic-cookbook.html>

[Installing OpenStack Octavia LBaaS with RHOSP in Contrail Networking | 334](#)

[Installing OpenStack Octavia LBaaS with Juju Charms in Contrail Networking | 455](#)

[Using Load Balancers in Contrail | 368](#)

[Configuring Load Balancing as a Service in Contrail | 385](#)

Configuring Load Balancing as a Service in Contrail

IN THIS SECTION

- [Overview: Load Balancing as a Service | 385](#)
- [Contrail LBaaS Implementation | 386](#)
- [Configuring LBaaS Using CLI | 387](#)
- [Configuring LBaaS using the Contrail Command UI | 390](#)

Overview: Load Balancing as a Service

Load Balancing as a Service (LBaaS) is a feature available through OpenStack Neutron. Contrail Release 1.20 and greater allows the use of the Neutron API for LBaaS to apply open source load balancing technologies to provision a load balancer in the Contrail system.

The LBaaS load balancer enables the creation of a pool of virtual machines serving applications, all front-ended by a virtual-ip. The LBaaS implementation has the following features:


- Load balancing of traffic from clients to a pool of backend servers. The load balancer proxies all connections to its virtual IP.
- Provides load balancing for HTTP, TCP, and HTTPS.
- Provides health monitoring capabilities for applications, including HTTP, TCP, and ping.
- Enables floating IP association to **virtual-ip** for public access to the backend pool.

In [Figure 66 on page 386](#), the load balancer is launched with the virtual IP address 198.51.100.2. The backend pool of virtual machine applications (App Pool) is on the subnet 203.0.113.0/24. Each of the application virtual machines gets an IP address (virtual-ip) from the pool subnet. When a client connects to the **virtual-ip** for accessing the application, the load balancer proxies the TCP connection on its **virtual-ip**, then creates a new TCP connection to one of the virtual machines in the pool.

The pool member is selected using one of following methods:

- weighted round robin (WRR), based on the weight assignment
- least connection, selects the member with the fewest connections
- source IP selects based on the **source-ip** of the packet

Figure 66: Load Balancing as a Service in Contrail



Additionally, the load balancer monitors the health of each pool member using the following methods:

- Monitors TCP by creating a TCP connection at intervals.
- Monitors HTTP by creating a TCP connection and issuing an HTTP request at intervals.
- Monitors ping by checking if a member can be reached by pinging.

Contrail LBaaS Implementation

Contrail supports the OpenStack LBaaS Neutron APIs and creates relevant objects for LBaaS, including **virtual-ip**, **loadbalancer-pool**, **loadbalancer-member**, and **loadbalancer-healthmonitor**. Contrail creates a service instance when a **loadbalancer-pool** is associated with a **virtual-ip** object. The service scheduler then launches a namespace on a randomly selected virtual router and spawns HAProxy into that namespace. The configuration for HAProxy is picked up from the load balancer objects. Contrail supports high availability of namespaces and HAProxy by spawning active and standby on two different routers.

A Note on Installation

To use the LBaaS feature, HAProxy, version 1.5 or greater and **iproute2**, version 3.10.0 or greater must both be installed on the Contrail compute nodes.

If you are using fab commands for installation, the haproxy and iproute2 packages will be installed automatically with LBaaS if you set the following:

```
env.enable_lbaas=True
```

Use the following to check the version of the **iproute2** package on your system:

```
root@nodeh5:/var/log# ip -V
ip utility, iproute2-ss130716
root@nodeh5:/var/log#
```

Limitations

LBaaS currently has these limitations:

- A pool should not be deleted before deleting the VIP.
- Multiple VIPs cannot be associated with the same pool. If pool needs to be reused, create another pool with the same members and bind it to the second VIP.
- Members cannot be moved from one pool to another. If needed, first delete the members from one pool, then add to a different pool.
- In case of active-standby failover, namespaces might not get cleaned up when the agent restarts.
- The floating-ip association needs to select the VIP port and not the service ports.

Configuring LBaaS Using CLI

The LBaaS feature is enabled on Contrail through Neutron API calls. The following procedure shows how to create a pool network and a VIP network using CLI. The VIP network is created in the public network and members are added in the pool network.

Creating a Load Balancer

Use the following steps to create a load balancer in Contrail.

1. Create a VIP network.

```
neutron net-create vipnet
```

```
neutron subnet-create --name vipsubnet vipnet 198.51.100.2
```

2. Create a pool network.

```
neutron net-create poolnet
```

```
neutron subnet-create --name poolsubnet poolnet 203.0.113.0/24
```

3. Create a pool for HTTP.

```
neutron lb-pool-create --lb-method ROUND_ROBIN --name mypool --protocol HTTP --subnet-id poolsubnet
```


4. Add members to the pool.

```
neutron lb-member-create --address 203.0.113.3 --protocol-port 80 mypool
```

```
neutron lb-member-create --address 203.0.113.4 --protocol-port 80 mypool
```

5. Create a VIP for HTTP and associate it to the pool.

```
neutron lb-vip-create --name myvip --protocol-port 80 --protocol HTTP--subnet-id vipsubnet mypool
```



Deleting a Load Balancer

Use the following steps to delete a load balancer in Contrail.

1. Delete the VIP.

```
neutron lb-vip-delete <vip-uuid>
```

2. Delete members from the pool.

```
neutron lb-member-delete <member-uuid>
```

3. Delete the pool.

```
neutron lb-pool-delete <pool-uuid>
```

Managing Healthmonitor for Load Balancer

Use the following commands to create a healthmonitor, associate a healthmonitor to a pool, disassociate a healthmonitor, and delete a healthmonitor.

1. Create a healthmonitor.

```
neutron lb-healthmonitor-create --delay 20 --timeout 10 --max-retries 3 --type HTTP
```

2. Associate a healthmonitor to a pool.

```
neutron lb-healthmonitor-associate <healthmonitor-uuid> mypool
```

3. Disassociate a healthmonitor from a pool.

```
neutron lb-healthmonitor-disassociate <healthmonitor-uuid> mypool
```

Configuring an SSL VIP with an HTTP Backend Pool

Use the following steps to configure an SSL VIP with an HTTP backend pool.

1. Copy an SSL certificate to all compute nodes.

```
scp ssl_certificate.pem <compute-node-ip> <certificate-path>
```

2. Update the information in `/etc/contrail/contrail-vrouter-agent.conf`.

```
# SSL certificate path haproxy
```

```
haproxy_ssl_cert_path=<certificate-path>
```

3. Restart `contrail-vrouter-agent`.

```
service contrail-vrouter-agent restart
```

4. Create a VIP for port 443 (SSL).

```
neutron lb-vip-create --name myvip --protocol-port 443 --protocol HTTP --subnet-id vipsubnet mypool
```

Configuring LBaaS using the Contrail Command UI

Create, edit, or delete load balancers using the Contrail Command UI. Use the following guidelines when creating load balancers:


- Each load balancer consists of one or more listeners, pools, pool members, and health monitors.
 - Listener: Port that listens for traffic from a particular load balancer. Multiple listeners can be associated with a single load balancer.
 - Pool: Group of hosts that serves traffic from the load balancer.
 - Pool Member: Server that is specified by the IP address and port for which it uses to serve the traffic it receives from the load balancer.
 - Health Monitor: Health monitors are associated with pools and help divert traffic away from pool members that are temporarily offline.
- Each load balancer can have multiple pools with one or more listeners for each pool.
- The native load balancer has a single pool that is shared among multiple listeners.

Creating a Load Balancer

Use the following steps to create a load balancer with the load balancer wizard.


- 1. Go to **Services > Load Balancers**.

Figure 67: Create Load Balancers



- 2. To create a load balancer, click **Create**.

Figure 68: Load Balancer Information



Add the load balancer information:

- **Name:** Name of the load balancer.
- **Description:** (Optional) Description of the load balancer.

- **Subnet:** Drop-down menu displays all subnets from list of all available networks. The subnet is the network on which to allocate the IP address of the load balancer.
- **Loadbalancer Provider:** Drop-down menu includes available options. Default is opencontrail.
- **Fixed IPs:** (Optional) IPv4 or IPv6 address.
- **Floating IP:** (Optional) IPv4 or IPv6 address.
- **Admin State:** Check the checkbox for UP or uncheck the checkbox for DOWN. Default is UP.

3. Click **Next**. The Listener fields are displayed.

Figure 69: Listener Information

The screenshot shows the 'Create Load Balancer' wizard in the Contrail Command Line interface. The wizard is currently at Step 2: Listener. The fields shown are:

- Name:** Listener-http
- Description:** Listener Description
- Protocol:** HTTP
- Port:** 80
- Connection Limit:** -1
- Admin State:** ☒

Navigation buttons at the bottom include 'Previous', 'Cancel', and 'Next'.

Add the listener information:

- **Name:** Name of the listener.
- **Description:** (Optional) Description of the listener.
- **Protocol:** Dropdown menu includes HTTP and TCP.
- **Port:** Must be an integer in the range of 1 to 65535.
- **Connection Limit:** (Optional). -1 indicates an infinite limit.
- **Admin State:** Check the checkbox for UP or uncheck the checkbox for DOWN. Default is UP.

4. Click **Next**. The Pool fields are displayed.

Figure 70: Pool Information

The screenshot shows the 'Create Load Balancer' wizard in the Contrail Command Line interface. The wizard is at Step 3: Pool. The previous steps are Load Balancer and Listener. The next steps are Pool Member and Monitor. The Pool step is active, showing fields for Name, Description, Method, Protocol, Session Persistence, and Admin State. The Admin State checkbox is checked. There are buttons for Previous, Cancel, and Next.

Add the pool information:

- **Name:** Name of the pool.
- **Description:** (Optional) Description of the pool.
- **Method:** Load balancing method used to distribute incoming requests. Dropdown menu includes LEAST_CONNECTIONS, ROUND_ROBIN, and SOURCE_IP.
- **Protocol:** The protocol used by the pool and its members for the load balancer traffic. Dropdown menu includes TCP and HTTP.
- **Session Persistence:** (Optional) Default value is an empty dictionary.
- **Admin State:** Check the checkbox for UP or uncheck the checkbox for DOWN. Default is UP.

5. Click **Next**. The list of available pool member instances are displayed. To add an external member, click the **+ Add** icon. Each pool member must have a unique IP address and port combination.

Figure 71: Pool Member Information

The screenshot shows the 'Create Load Balancer' wizard in the Contrail Command interface. The wizard is at Step 4 (optional) 'Pool Member'. The form contains the following fields:

- Name ***: Member-1
- Subnet ***: 2.0.0.0/24 (VN2)
- IP Address ***: xxxxx
- Port**: 80
- Weight**: 1
- Admin State**: ☒ Admin State

Below the form, there is a table with one member:

Name	Subnet	IP Address	Port	Weight	Admin State
Member-1	2.0.0.0/24 (VN2)	xxxxx	80	1	☒ Admin State

At the bottom of the form, there is an '+ Add' button and 'Previous', 'Cancel', and 'Next' buttons.

The pool member information includes:

- **Name:** Name of the pool member.
- **Subnet:** The subnet in which to access the member.
- **IP Address:** The IP address of the member that is used to receive traffic from the load balancer.
- **Port:** The port to which the member listens to receive traffic from the load balancer.
- **Weight:** The default value is 1.
- **Admin State:** Check the checkbox for UP or uncheck the checkbox for DOWN. Default is UP.

6. Click **Next**. The Monitor fields are displayed.

Figure 72: Health Monitor Information

The screenshot shows the 'Create Load Balancer' wizard in the Contrail Command interface. The wizard is at Step 5: Monitor. The left sidebar shows the navigation menu with 'SERVICES' selected. The main area displays the 'Monitor' configuration fields.

Monitor Configuration Fields:

- Monitor Type:** HTTP (dropdown menu)
- HTTP Method:** GET (dropdown menu)
- Expected HTTP Status Code:** 200 (text input)
- URL Path:** / (text input)
- Health check interval (sec):** 5 (text input)
- Retry count before markdown:** 3 (text input)
- Timeout (sec):** 5 (text input)
- Admin State:** ☒ Admin State (checkbox)

Navigation buttons: Previous, Cancel, Finish.

Add the health monitor information:

- **Monitor Type:** Dropdown menu includes HTTP, PING, and TCP.
- **HTTP Method:** Required if monitor type is HTTP. Dropdown menu includes GET and HEAD. The default value is GET.
- **Expected HTTP Status Code:** Required if monitor type is HTTP. The default value is 200.
- **URL Path:** Required if monitor type is HTTP. The default value is “/.”
- **Health check interval (sec):** The time interval, in seconds, between each health check. The default value is 5.
- **Retry count before markdown:** The maximum number of failed health checks before the state of a member is changed to OFFLINE. The default value is 3.
- **Timeout (sec):** The maximum number of seconds allowed for any given health check to complete. The timeout value should always be less than the health check interval. The default value is 5.
- **Admin State:** Check the checkbox for UP or uncheck the checkbox for DOWN. Default is UP.


Click **Finish**.

Viewing or Editing Load Balancers

Use the following steps to view or edit existing load balancers.


- 1. Go to **Services > Load Balancers**. A summary screen of the Load Balancers is displayed.

Figure 73: Summary Screen of Load Balancers



- 2. To view summary of a load balancer, click the drop down arrow next to a load balancer listed in the summary screen. The Load Balancer Info window is displayed.


Figure 74: Load Balancer Info Window



Deleting a Load Balancer

To delete the load balancer, click the trash can icon.

ERROR: Unresolved graphic fileref="" not found in
"/cmsxml/default/main/supplemental/STAGING/images/".



SEE ALSO

- Installing OpenStack Octavia LBaaS with RHOSP in Contrail Networking | 334
- Installing OpenStack Octavia LBaaS with Juju Charms in Contrail Networking | 455
- Using Load Balancers in Contrail | 368
- Support for OpenStack LBaaS | 381

Optimizing Contrail Networking

IN THIS SECTION

- [Multiqueue Virtio Interfaces in Virtual Machines | 399](#)

Multiqueue Virtio Interfaces in Virtual Machines

IN THIS SECTION

- [Multiqueue Virtio Overview | 399](#)
- [Requirements and Setup for Multiqueue Virtio Interfaces | 399](#)

Contrail 3.2 adds support for multiqueue for the DPDK-based vrouter.

Contrail 3.1 supports multiqueue virtio interfaces for Ubuntu kernel-based router, only.

Multiqueue Virtio Overview

OpenStack Liberty supports the ability to create VMs with multiple queues on their virtio interfaces. Virtio is a Linux platform for I/O virtualization, providing a common set of I/O virtualization drivers. Multiqueue virtio is an approach that enables the processing of packet sending and receiving to be scaled to the number of available virtual CPUs (vCPUs) of a guest, through the use of multiple queues.

Requirements and Setup for Multiqueue Virtio Interfaces

To use multiqueue virtio interfaces, ensure your system meets the following requirements:

- The OpenStack version must be Liberty or greater.
- The maximum number of queues in the VM interface is set to the same value as the number of vCPUs in the guest.
- The VM image metadata property is set to enable multiple queues inside the VM.

Setting Virtual Machine Metadata for Multiple Queues

Use the following command on the OpenStack node to enable multiple queues on a VM:

```
source /etc/contrail/openstackrc
nova image-meta <image_name> set hw_vif_multiqueue_enabled="true"
```

After the VM is spawned, use the following command on the virtio interface in the guest to enable multiple queues inside the VM:

ethtool -L <interface_name> combined <#queues>

Packets will now be forwarded on all queues in the VM to and from the vRouter running on the host.

NOTE: Multiple queues in the VM are only supported with the kernel mode vRouter in Contrail 3.1.

Contrail 3.2 adds support for multiple queues with the DPDK-based vrouter, using OpenStack Mitaka. The DPDK vrouter has the same setup requirements as the kernel mode vrouter. However, in the **ethtool -L** setup command, the number of queues cannot be higher than the number of CPU cores assigned to vrouter in the testbed file.

Contrail Networking OpenStack Analytics

IN THIS SECTION

- [Ceilometer Support in Contrail | 400](#)

Ceilometer Support in Contrail

IN THIS SECTION

- [Overview | 401](#)
- [Ceilometer Details | 401](#)
- [Verification of Ceilometer Operation | 401](#)
- [Contrail Ceilometer Plugin | 404](#)
- [Ceilometer Installation and Provisioning | 407](#)

Ceilometer is an OpenStack feature that provides an infrastructure for collecting SDN metrics from OpenStack projects. The metrics can be used by various rating engines to transform events into billable items. The Ceilometer collection process is sometimes referred to as “metering”. The Ceilometer service provides data that can be used by platforms that provide metering, tracking, billing, and similar services. This topic describes how to configure the Ceilometer service for Contrail.

Overview

Contrail Release 2.20 and later supports the OpenStack Ceilometer service, on the OpenStack Juno release on Ubuntu 14.04.1 LTS.

The prerequisites for installing Ceilometer are:

- Contrail Cloud installation
- Provisioned using **enable_ceilometer = True** in the **provisioning** file.

NOTE: Ceilometer services are only installed on the first OpenStack controller node and do not support high availability in Contrail Release 2.20.

Ceilometer Details

Ceilometer is used to reliably collect measurements of the utilization of the physical and virtual resources comprising deployed clouds, persist these data for subsequent retrieval and analysis, and trigger actions when defined criteria are met.

The Ceilometer architecture consists of:

Polling agent—Agent designed to poll OpenStack services and build meters. The polling agents are also run on the compute nodes in addition to the OpenStack controller.

Notification agent—Agent designed to listen to notifications on message queue and convert them to events and samples.

Collector —Gathers and records event and metering data created by the notification and polling agents.

API server—Provides a REST API to query and view data recorded by the collector service.

Alarms—Daemons to evaluate and notify based on defined alarming rules.

Database—Stores the metering data, notifications, and alarms. The supported databases are MongoDB, SQL-based databases compatible with SQLAlchemy, and HBase. The recommended database is MongoDB, which has been thoroughly tested with Contrail and deployed on a production scale.

Verification of Ceilometer Operation

The Ceilometer services are named slightly differently on the Ubuntu and RHEL Server 7.0.

On Ubuntu, the service names are:

Polling agent—`ceilometer-agent-central` and `ceilometer-agent-compute`

Notification agent—`ceilometer-agent-notification`

Collector —`ceilometer-collector`

API Server—`ceilometer-api`

Alarms—`ceilometer-alarm-evaluator` and `ceilometer-alarm-notifier`

On RHEL Server 7.0, the service names are:

Polling agent—`openstack-ceilometer-central` and `openstack-ceilometer-compute`

Notification agent—`openstack-ceilometer-notification`

Collector —`openstack-ceilometer-collector`

API server—`openstack-ceilometer-api`

Alarms—`openstack-ceilometer-alarm-evaluator` and `openstack-ceilometer-alarm-notifier`

To verify the Ceilometer installation, users can verify that the Ceilometer services are up and running by using the **openstack-status** command.

For example, using the **openstack-status** command on an all-in-one node running Ubuntu 14.04.1 LTS with release 2.2 of Contrail installed shows the following Ceilometer services as active:

```
== Ceilometer services ==
ceilometer-api:           active
ceilometer-agent-central: active
ceilometer-agent-compute: active
ceilometer-collector:     active
ceilometer-alarm-notifier: active
ceilometer-alarm-evaluator: active
ceilometer-agent-notification: active
```

You can issue the **ceilometer meter-list** command on the OpenStack controller node to verify that meters are being collected, stored, and reported via the REST API. The following is an example of the output:

```
user@host:~# (source /etc/contrail/openstackrc; ceilometer meter-list)
```

Name	User ID	Type	Unit	Resource ID	Project ID
ip.floating.receive.bytes		cumulative	B		

```

a726f93a-65fa-4cad-828b-54dbfcf4a119 | None | None
|
| ip.floating.receive.packets | cumulative | packet |
a726f93a-65fa-4cad-828b-54dbfcf4a119 | None | None
|
| ip.floating.transmit.bytes | cumulative | B |
a726f93a-65fa-4cad-828b-54dbfcf4a119 | None | None
|
| ip.floating.transmit.packets | cumulative | packet |
a726f93a-65fa-4cad-828b-54dbfcf4a119 | None | None
|
| network | gauge | network |
7fa6796b-756e-4320-9e73-87d4c52ecc83 | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| network | gauge | network |
9408e287-d3e7-41e2-89f0-5c691c9ca450 | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| network | gauge | network |
b3b72b98-f61e-4e1f-9a9b-84f4f3ddec0b | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| network | gauge | network |
cb829abd-e6a3-42e9-a82f-0742db55d329 | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| network.create | delta | network |
7fa6796b-756e-4320-9e73-87d4c52ecc83 | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| network.create | delta | network |
9408e287-d3e7-41e2-89f0-5c691c9ca450 | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| network.create | delta | network |
b3b72b98-f61e-4e1f-9a9b-84f4f3ddec0b | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| network.create | delta | network |
cb829abd-e6a3-42e9-a82f-0742db55d329 | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| port | gauge | port |
0d401d96-c2bf-4672-abf2-880eecf25ceb | 01edcedd989f43b3a2d6121d424b254d |
82ab961f88994e168217ddd746fdd826 |
| port | gauge | port |
211b94a4-581d-45d0-8710-c6c69df15709 | 01edcedd989f43b3a2d6121d424b254d |
82ab961f88994e168217ddd746fdd826 |
| port | gauge | port |
2287ce25-4eef-4212-b77f-3cf590943d36 | 01edcedd989f43b3a2d6121d424b254d |
82ab961f88994e168217ddd746fdd826 |

```

```

| port.create          | delta      | port      |
f62f3732-222e-4c40-8783-5bcbc1fd6a1c | 01edcedd989f43b3a2d6121d424b254d |
82ab961f88994e168217ddd746fdd826 |
| port.create          | delta      | port      |
f8c89218-3cad-48e2-8bd8-46c1bc33e752 | 01edcedd989f43b3a2d6121d424b254d |
82ab961f88994e168217ddd746fdd826 |
| port.update          | delta      | port      |
43ed422d-b073-489f-877f-515a3cc0b8c4 | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| subnet               | gauge      | subnet     |
09105ed1-1654-4b5f-8c12-f0f2666fa304 | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| subnet               | gauge      | subnet     |
4bf00aac-407c-4266-a048-6ff52721ad82 | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| subnet.create        | delta      | subnet     |
09105ed1-1654-4b5f-8c12-f0f2666fa304 | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
| subnet.create        | delta      | subnet     |
4bf00aac-407c-4266-a048-6ff52721ad82 | 15c0240142084d16b3127d6f844adbd9 |
ded208991de34fe4bb7dd725097f1c7e |
+-----+-----+-----+-----+-----+-----+

```

NOTE: The **ceilometer meter-list** command lists the meters only if images have been created, or instances have been launched, or if subnet, port, floating IP addresses have been created, otherwise the meter list is empty. You also need to source the **/etc/contrail/openstackrc** file when executing the command.

Contrail Ceilometer Plugin

The Contrail Ceilometer plugin adds the capability to meter the traffic statistics of floating IP addresses in Ceilometer. The following meters for each floating IP resource are added by the plugin in Ceilometer.

```

ip.floating.receive.bytes
ip.floating.receive.packets
ip.floating.transmit.bytes
ip.floating.transmit.packets

```

The Contrail Ceilometer plugin configuration is done in the **/etc/ceilometer/pipeline.yaml** file when Contrail is installed by the Fabric provisioning scripts.

The following example shows the configuration that is added to the file:

```
sources:
  - name: contrail_source
    interval: 600
    meters:
      - "ip.floating.receive.packets"
      - "ip.floating.transmit.packets"
      - "ip.floating.receive.bytes"
      - "ip.floating.transmit.bytes"
    resources:
      - contrail://<IP-address-of-Contrail-Analytics-Node>:8081
    sinks:
      - contrail_sink
sinks:
  - name: contrail_sink
    publishers:
      - rpc://
    transformers:
```

The following example shows the Ceilometer meter list output for the floating IP meters:

Name	Type	Unit	Resource ID	
Project ID		User ID		
ip.floating.receive.bytes	cumulative	B		
451c93eb-e728-4ba1-8665-6e7c7a8b49e2				None
				None
ip.floating.receive.bytes	cumulative	B		
9cf76844-8f09-4518-a09e-e2b8832bf894				None
				None
ip.floating.receive.packets	cumulative	packet		
451c93eb-e728-4ba1-8665-6e7c7a8b49e2				None
				None
ip.floating.receive.packets	cumulative	packet		
9cf76844-8f09-4518-a09e-e2b8832bf894				None
				None
ip.floating.transmit.bytes	cumulative	B		
451c93eb-e728-4ba1-8665-6e7c7a8b49e2				None
				None
ip.floating.transmit.bytes	cumulative	B		
9cf76844-8f09-4518-a09e-e2b8832bf894				None
				None
ip.floating.transmit.packets	cumulative	packet		

Period		Period Start			Period End			Count	Min
Max	Sum	Avg		Duration		Duration Start		Duration	
End									
0		2015-02-13T19:50:40.795000			2015-02-13T19:50:40.795000			2892	0.0
325.0	1066.0	0.368603042877		439069.674		2015-02-13T19:50:40.795000			
2015-02-18T21:48:30.469000									
</									

Ceilometer Installation and Provisioning

There are two scenarios possible for Contrail Ceilometer plugin installation.

- 1. If you install your own OpenStack distribution, you can install the Contrail Ceilometer plugin on the OpenStack controller node.
- 2. When using Contrail Cloud services, the Ceilometer controller services are installed and provisioned as part of the OpenStack controller node and the compute agent service is installed as part of the compute node when **enable_ceilometer** is set as **True** in the cluster **config** or **testbed** files.

Contrail OpenStack APIs

IN THIS SECTION

- [Working with Neutron | 408](#)

Working with Neutron

IN THIS SECTION

- [Data Structure | 408](#)
- [Network Sharing in Neutron | 409](#)
- [Commands for Neutron Network Sharing | 409](#)
- [Support for Neutron APIs | 410](#)
- [Contrail Neutron Plugin | 410](#)
- [DHCP Options | 411](#)
- [Incompatibilities | 411](#)

OpenStack's networking solution, Neutron, has representative elements for Contrail elements for Network (VirtualNetwork), Port (VirtualMachineInterface), Subnet (IpamSubnets), and Security-Group. The Neutron plugin translates the elements from one representation to another.

Data Structure

Although the actual data between Neutron and Contrail is similar, the listings of the elements differs significantly. In the Contrail API, the networking elements list is a summary, containing only the UUID, FQ name, and an href, however, in Neutron, all details of each resource are included in the list.

The Neutron plugin has an inefficient list retrieval operation, especially at scale, because it:

- reads a list of resources (for example. **GET /virtual-networks**), then
- iterates and reads in the details of the resource (**GET /virtual-network/<uuid>**).

As a result, the API server spends most of the time in this type of GET operation just waiting for results from the Cassandra database.

The following features in Contrail improve performance with Neutron:

- An optional detail query parameter is added in the GET of collections so that the API server returns details of all the resources in the list, instead of just a summary. This is accompanied by changes in the Contrail API library so that a caller gets returned a list of the objects.
- The existing Contrail list API takes in an optional **parent_id** query parameter to return information about the resource anchored by the parent.
- The Contrail API server reads objects from Cassandra in a multiget format into **obj_uuid_cf**, where object contents are stored, instead of reading in an xget/get format. This reduces the number of round-trips to and from the Cassandra database.

Network Sharing in Neutron

Using Neutron, a deployer can make a network accessible to other tenants or projects by using one of two attributes on a network:

- Set the **shared** attribute to allow sharing.
- Set the **router:external** attribute, when the plugin supports an **external_net** extension.

Using the Shared Attribute

When a network has the **shared** attribute set, users in other tenants or projects, including non-admin users, can access that network, using:

neutron net-list --shared

Users can also launch a virtual machine directly on that network, using:

nova boot <other-parameters> -nic net-id=<shared-net-id>

Using the Router:External Attribute

When a network has the **router:external** attribute set, users in other tenants or projects, including non-admin users, can use that network for allocating floating IPs, using:

neutron floatingip-create <router-external-net-id>

then associating the IP address pool with their instances.

NOTE: The VN hosting the FIP pool should be marked shared and external.

Commands for Neutron Network Sharing

The following table summarizes the most common Neutron commands used with Contrail.

Action	Command
List all shared networks.	neutron net-list --shared
Create a network that has the shared attribute.	neutron net-create <net-name> -shared
Set the shared attribute on an existing network.	neutron net-update <net-name> -shared
List all router:external networks.	neutron net-list --router:external
Create a network that has the router:external attribute.	neutron net-create <net-name> -router:external
Set the router:external attribute on an existing network.	neutron net-update <net-name> -router:external

Support for Neutron APIs

The OpenStack Neutron project provides virtual networking services among devices that are managed by the OpenStack compute service. Software developers create applications by using the OpenStack Networking API v2.0 (Neutron).

Contrail provides the following features to increase support for OpenStack Neutron:

- Create a port independently of a virtual machine.
- Support for more than one subnet on a virtual network.
- Support for allocation pools on a subnet.
- Per tenant quotas.
- Enabling DHCP on a subnet.
- External router can be used for floating IPs.

For more information about using OpenStack Networking API v2.0 (Neutron), refer to:

<http://docs.openstack.org/api/openstack-network/2.0/content/> and the OpenStack Neutron Wiki at: <http://wiki.openstack.org/wiki/Neutron>.

Contrail Neutron Plugin

The Contrail Neutron plugin provides an implementation for the following core resources:

- Network
- Subnet
- Port

It also implements the following standard and upstreamed Neutron extensions:

- Security group
- Router IP and floating IP
- Per-tenant quota
- Allowed address pair

The following Contrail-specific extensions are implemented:

- Network IPAM
- Network policy
- VPC table and route table
- Floating IP pools

The plugin does not implement native bulk, pagination, or sort operations and relies on emulation provided by the Neutron common code.

DHCP Options

In Neutron commands, DHCP options can be configured using `extra-dhcp-options` in `port-create`.

Example

```
neutron port-create net1 --extra-dhcp-opt  
opt_name=<dhcp_option_name>,opt_value=<value>
```

The `opt_name` and `opt_value` pairs that can be used are maintained in GitHub:

<https://github.com/Juniper/contrail-controller/wiki/Extra-DHCP-Options> .

Incompatibilities

In the Contrail architecture, the following are known incompatibilities with the Neutron API.

- Filtering based on any arbitrary key in the resource is not supported. The only supported filtering is by **id**, **name**, and **tenant_id**.
- To use a floating IP, it is not necessary to connect the public subnet and the private subnet to a Neutron router. Marking a public network with **router:external** is sufficient for a floating IP to be created and associated, and packet forwarding to it will work.
- The default values for quotas are sourced from **/etc/contrail/contrail-api.conf** and not from **/etc/neutron/neutron.conf**.

Using Contrail with Juju Charms

IN THIS CHAPTER

- [Installing Contrail with OpenStack by Using Juju Charms | 412](#)
- [Installing Contrail with Kubernetes by Using Juju Charms | 437](#)
- [Installing Contrail with Kubernetes in Nested Mode by Using Juju Charms | 451](#)
- [Installing OpenStack Octavia LBaaS with Juju Charms in Contrail Networking | 455](#)
- [Using Netronome SmartNIC vRouter with Contrail Networking and Juju Charms | 464](#)

Installing Contrail with OpenStack by Using Juju Charms

IN THIS SECTION

- [Preparing to Deploy Contrail by Using Juju Charms | 413](#)
- [Deploying Contrail Charms | 415](#)
- [Options for Juju Charms | 429](#)

You can deploy Contrail by using Juju Charms. Juju helps you deploy, configure, and efficiently manage applications on private clouds and public clouds. Juju accesses the cloud with the help of a Juju controller. A Charm is a module containing a collection of scripts and metadata and is used with Juju to deploy Contrail.

Starting in Contrail Networking Release 2011, Contrail Networking supports OpenStack Ussuri with Ubuntu version 18.04 (Bionic Beaver) and Ubuntu version 20.04 (Focal Fossa).

Contrail supports the following charms:

- `contrail-agent`
- `contrail-analytics`
- `contrail-analyticsdb`

- contrail-controller
- contrail-keystone-auth
- contrail-openstack

These topics describe how to deploy Contrail by using Juju Charms.

Preparing to Deploy Contrail by Using Juju Charms

Follow these steps to prepare for deployment:

1. Install Juju.

```
sudo apt-get update
sudo apt-get upgrade
sudo apt-get install juju
```

2. Configure Juju.

You can add a cloud to Juju, identify clouds supported by Juju, and also manage clouds already added to Juju.

- **Adding a cloud**—Juju recognizes a wide range of cloud types. You can use any one of the following methods to add a cloud to Juju:

- **Adding a Cloud by Using Interactive Command**

Example: Adding an MAAS cloud to Juju

```
juju add-cloud
```

```
Cloud Types
```

```
maas
manual
openstack
oracle
vsphere
```

```
Select cloud type: maas
```

```
Enter a name for your maas cloud: maas-cloud
```

```
Enter the API endpoint url: http://<ip-address>:<node>/MAAS
```

```
Cloud "maas-cloud" successfully added
You may bootstrap with 'juju bootstrap maas-cloud'
```

NOTE: Juju 2.x is compatible with MAAS series 1.x and 2.x.

• Adding a Cloud Manually

You use a YAML configuration file to add a cloud manually. Enter the following command:

```
juju add-cloud <cloud-name>
juju add-credential <cloud name>
```

For an example, to add the cloud *junmaas*, assuming that the name of the configuration file in the directory is **maas-clouds.yaml**, you run the following command:

```
juju add-cloud junmaas maas-clouds.yaml
```

The following is the format of the YAML configuration file:

```
clouds:
  <cloud_name>:
    type: <type_of_cloud>
    auth-types: [<authentication_types>]
    regions:
      <region-name>:
        endpoint: <http://<ip-address>:<node>/MAAS>
```

NOTE: The **auth-types** for a MAAS cloud type is **oauth1**.

• Identifying a supported cloud

Juju recognizes the cloud types given below. You use the **juju clouds** command to list cloud types that are supported by Juju.

```
$ juju clouds
```

Cloud	Regions	Default	Type	Description

aws	15	us-east-1	ec2	Amazon Web Services
aws-china	1	cn-north-1	ec2	Amazon China
aws-gov	1	us-gov-west-1	ec2	Amazon (USA Government)
azure	26	centralus	azure	Microsoft Azure
azure-china	2	chinaeast	azure	Microsoft Azure China
cloudsigma	5	hnl	cloudsigma	CloudSigma Cloud
google	13	us-east1	gce	Google Cloud Platform
joyent	6	eu-ams-1	joyent	Joyent Cloud
oracle	5	uscom-central-1	oracle	Oracle Cloud
rackspace	6	dfw	rackspace	Rackspace Cloud
localhost	1	localhost	lxd	LXD Container Hypervisor

3. Create a Juju controller.

```
juju bootstrap --bootstrap-series=xenial <cloud name> <controller name>
```

NOTE: A Juju controller manages and keeps track of applications in the Juju cloud environment.

Deploying Contrail Charms

IN THIS SECTION

- [Deploy Contrail Charms in a Bundle | 415](#)
- [Deploying Juju Charms with OpenStack Manually | 423](#)

You can deploy Contrail Charms in a bundle or manually.

Deploy Contrail Charms in a Bundle

Follow these steps to deploy Contrail Charms in a bundle.

1. Deploy Contrail Charms.

To deploy Contrail Charms in a bundle, use the **juju deploy <bundle_yaml_file>** command.

The following example shows you how to use **bundle_yaml_file** to deploy Contrail on Amazon Web Services (AWS) Cloud.

```

series: bionic

variables:
  openstack-origin:      &openstack-origin      distro
  #vhost-gateway:        &vhost-gateway          "192.x.40.254"
  data-network:          &data-network            "192.x.40.0/24"

  control-network:       &control-network         "192.x.30.0/24"

  virtioforwarder-coremask: &virtioforwarder-coremask    "1,2"
  agilio-registry:        &agilio-registry        "netronomesystems"

  agilio-image-tag:       &agilio-image-tag
"latest-ubuntu-queens"
  agilio-user:            &agilio-user
"<agilio-username>"
  agilio-password:        &agilio-password
"<agilio-password>"
  agilio-insecure:        &agilio-insecure        false
  agilio-phy:             &agilio-phy              "nfp_p0"
  docker-registry:        &docker-registry
"<registry-directory>"
  #docker-user:           &docker-user
"<docker_username>"
  #docker-password:       &docker-password
"<docker_password>"
  image-tag:              &image-tag              "2008.121"
  docker-registry-insecure: &docker-registry-insecure    "true"
  dockerhub-registry:     &dockerhub-registry
"https://index.docker.io/v1/"

machines:
  "1":
    constraints: tags=controller
    series: bionic
  "2":
    constraints: tags=compute
    series: bionic
  "3":
    constraints: tags=neutron
    series: bionic
services:
  ubuntu:
    charm: cs:ubuntu

```



```

    num_units: 1
    to: [ "1" ]
ntp:
  charm: cs:ntp
  num_units: 0
  options:
    #source: ntp.ubuntu.com
    source: 10.204.217.158
mysql:
  charm: cs:percona-cluster
  num_units: 1
  options:
    dataset-size: 15%
    max-connections: 10000
    root-password: <password>
    sst-password: <password>
    min-cluster-size: 1
  to: [ "lxd:1" ]
rabbitmq-server:
  num_units: 1
  options:
    min-cluster-size: 1
  to: [ "lxd:1" ]
heat:
  charm: cs:heat
  num_units: 1
  expose: true
  options:
    debug: true
    openstack-origin: *openstack-origin
  to: [ "lxd:1" ]
keystone:
  charm: cs:keystone
  expose: true
  num_units: 1
  options:
    admin-password: <password>
    admin-role: admin
    openstack-origin: *openstack-origin
    preferred-api-version: 3
nova-cloud-controller:
  charm: cs:nova-cloud-controller
  num_units: 1
  expose: true

```

```

    options:
      network-manager: Neutron
      openstack-origin: *openstack-origin
    to: [ "lxd:1" ]
neutron-api:
  charm: cs:neutron-api
  expose: true
  num_units: 1
  series: bionic
  options:
    manage-neutron-plugin-legacy-mode: false
    openstack-origin: *openstack-origin
  to: [ "3" ]
glance:
  charm: cs:glance
  expose: true
  num_units: 1
  options:
    openstack-origin: *openstack-origin
  to: [ "lxd:1" ]
openstack-dashboard:
  charm: cs:openstack-dashboard
  expose: true
  num_units: 1
  options:
    openstack-origin: *openstack-origin
  to: [ "lxd:1" ]
nova-compute:
  charm: cs:nova-compute
  num_units: 0
  expose: true
  options:
    openstack-origin: *openstack-origin
nova-compute-dpdk:
  charm: cs:nova-compute
  num_units: 0
  expose: true
  options:
    openstack-origin: *openstack-origin
nova-compute-accel:
  charm: cs:nova-compute
  num_units: 2
  expose: true
  options:

```

```

    openstack-origin: *openstack-origin
  to: [ "2" ]
contrail-openstack:
  charm: ./tf-charms/contrail-openstack
  series: bionic
  expose: true
  num_units: 0
  options:
    docker-registry: *docker-registry
    #docker-user: *docker-user
    #docker-password: *docker-password
    image-tag: *image-tag
    docker-registry-insecure: *docker-registry-insecure
contrail-agent:
  charm: ./tf-charms/contrail-agent
  num_units: 0
  series: bionic
  expose: true
  options:
    log-level: "SYS_DEBUG"
    docker-registry: *docker-registry
    #docker-user: *docker-user
    #docker-password: *docker-password
    image-tag: *image-tag
    docker-registry-insecure: *docker-registry-insecure
    #vhost-gateway: *vhost-gateway
    physical-interface: *agilio-phy
contrail-agent-dpdk:
  charm: ./tf-charms/contrail-agent
  num_units: 0
  series: bionic
  expose: true
  options:
    log-level: "SYS_DEBUG"
    docker-registry: *docker-registry
    #docker-user: *docker-user
    #docker-password: *docker-password
    image-tag: *image-tag
    docker-registry-insecure: *docker-registry-insecure
    dpdk: true
    dpdk-main-mempool-size: "65536"
    dpdk-pmd-txd-size: "2048"
    dpdk-pmd-rxd-size: "2048"
    dpdk-driver: ""

```

```

    dpdk-coremask: "1-4"
    #vhost-gateway: *vhost-gateway
    physical-interface: "nfp_p0"
contrail-analytics:
  charm: ./tf-charms/contrail-analytics
  num_units: 1
  series: bionic
  expose: true
  options:
    log-level: "SYS_DEBUG"
    docker-registry: *docker-registry
    #docker-user: *docker-user
    #docker-password: *docker-password
    image-tag: *image-tag
    control-network: *control-network
    docker-registry-insecure: *docker-registry-insecure
  to: [ "1" ]
contrail-analyticsdb:
  charm: ./tf-charms/contrail-analyticsdb
  num_units: 1
  series: bionic
  expose: true
  options:
    log-level: "SYS_DEBUG"
    cassandra-minimum-diskgb: "4"
    cassandra-jvm-extra-opts: "-Xms8g -Xmx8g"
    docker-registry: *docker-registry
    #docker-user: *docker-user
    #docker-password: *docker-password
    image-tag: *image-tag
    control-network: *control-network
    docker-registry-insecure: *docker-registry-insecure
  to: [ "1" ]
contrail-controller:
  charm: ./tf-charms/contrail-controller
  series: bionic
  expose: true
  num_units: 1
  options:
    log-level: "SYS_DEBUG"
    cassandra-minimum-diskgb: "4"
    cassandra-jvm-extra-opts: "-Xms8g -Xmx8g"
    docker-registry: *docker-registry
    #docker-user: *docker-user

```

```

    #docker-password: *docker-password
    image-tag: *image-tag
    docker-registry-insecure: *docker-registry-insecure
    control-network: *control-network
    data-network: *data-network
    auth-mode: no-auth
  to: [ "1" ]
contrail-keystone-auth:
  charm: ./tf-charms/contrail-keystone-auth
  series: bionic
  expose: true
  num_units: 1
  to: [ "lxd:1" ]
agilio-vrouter5:
  charm: ./charm-agilio-vrt-5-37
  expose: true
  options:
    virtioforwarder-coremask: *virtioforwarder-coremask
    agilio-registry: *agilio-registry
    agilio-insecure: *agilio-insecure
    agilio-image-tag: *agilio-image-tag
    agilio-user: *agilio-user
    agilio-password: *agilio-password
relations:
  - [ "ubuntu", "ntp" ]
  - [ "neutron-api", "ntp" ]
  - [ "keystone", "mysql" ]
  - [ "glance", "mysql" ]
  - [ "glance", "keystone" ]
  - [ "nova-cloud-controller:shared-db", "mysql:shared-db" ]
  - [ "nova-cloud-controller:amqp", "rabbitmq-server:amqp" ]
  - [ "nova-cloud-controller", "keystone" ]
  - [ "nova-cloud-controller", "glance" ]
  - [ "neutron-api", "mysql" ]
  - [ "neutron-api", "rabbitmq-server" ]
  - [ "neutron-api", "nova-cloud-controller" ]
  - [ "neutron-api", "keystone" ]
  - [ "nova-compute:amqp", "rabbitmq-server:amqp" ]
  - [ "nova-compute", "glance" ]
  - [ "nova-compute", "nova-cloud-controller" ]
  - [ "nova-compute", "ntp" ]
  - [ "openstack-dashboard:identity-service", "keystone" ]
  - [ "contrail-keystone-auth", "keystone" ]
  - [ "contrail-controller", "contrail-keystone-auth" ]

```

```

- [ "contrail-analytics", "contrail-analyticsdb" ]
- [ "contrail-controller", "contrail-analytics" ]
- [ "contrail-controller", "contrail-analyticsdb" ]
- [ "contrail-openstack", "nova-compute" ]
- [ "contrail-openstack", "neutron-api" ]
- [ "contrail-openstack", "contrail-controller" ]
- [ "contrail-agent:juju-info", "nova-compute:juju-info" ]
- [ "contrail-agent", "contrail-controller" ]
- [ "contrail-agent-dpdk:juju-info", "nova-compute-dpdk:juju-info" ]
- [ "contrail-agent-dpdk", "contrail-controller" ]
- [ "nova-compute-dpdk:amqp", "rabbitmq-server:amqp" ]
- [ "nova-compute-dpdk", "glance" ]
- [ "nova-compute-dpdk", "nova-cloud-controller" ]
- [ "nova-compute-dpdk", "ntp" ]
- [ "contrail-openstack", "nova-compute-dpdk" ]
- [ "contrail-agent:juju-info", "nova-compute-accel:juju-info" ]
- [ "nova-compute-accel:amqp", "rabbitmq-server:amqp" ]
- [ "nova-compute-accel", "glance" ]
- [ "nova-compute-accel", "nova-cloud-controller" ]
- [ "nova-compute-accel", "ntp" ]
- [ "contrail-openstack", "nova-compute-accel" ]
- [ "agilio-vrouter5:juju-info", "nova-compute-accel:juju-info" ]

```

You can create or modify the Contrail Charm deployment bundle YAML file to:

- Point to machines or instances where the Contrail Charms must be deployed.
- Include the options you need.

Each Contrail Charm has a specific set of options. The options you choose depend on the charms you select. For more information on the options that are available, see [“Options for Juju Charms” on page 429](#).

2. (Optional) Check the status of deployment.

You can check the status of the deployment by using the **juju status** command.

3. Enable configuration statements.

Based on your deployment requirements, you can enable the following configuration statements:

- **contrail-agent**

For more information, see <https://jaas.ai/u/juniper-os-software/contrail-agent/>.

- **contrail-analytics**

For more information, see <https://jaas.ai/u/juniper-os-software/contrail-analytics/>.

- **contrail-analyticsdb**

For more information, see <https://jaas.ai/u/juniper-os-software/contrail-analyticsdb>.

- **contrail-controller**

For more information, see <https://jaas.ai/u/juniper-os-software/contrail-controller>.

- **contrail-keystone-auth**

For more information, see <https://jaas.ai/u/juniper-os-software/contrail-keystone-auth>.

- **contrail-openstack**

For more information see, <https://jaas.ai/u/juniper-os-software/contrail-openstack>.

Deploying Juju Charms with OpenStack Manually

Before you begin deployment, ensure that you have:

- Installed and configured Juju
- Created a Juju controller
- Ubuntu 16.04 or Ubuntu 18.04 installed

Follow these steps to deploy Juju Charms manually:

1. Create machine instances for OpenStack, compute, and Contrail.

```
juju add-machine --constraints mem=8G cores=2 root-disk=40G --series=xenial
#for openstack machine(s) 0
```

```
juju add-machine --constraints mem=7G cores=4 root-disk=40G --series=xenial
#for compute machine(s) 1,(3)
```

```
juju add-machine --constraints mem=15G cores=2 root-disk=300G --series=xenial
#for contrail machine 2
```

2. Deploy OpenStack services.

You can deploy OpenStack services by using any one of the following methods:

- **By specifying the OpenStack parameters in a YAML file**

The following is an example of a YAML-formatted (**nova-compute-config.yaml**) file.

```
nova-compute:
  openstack-origin: cloud:xenial-ocata
```

```
virt-type: qemu
enable-resize: True
enable-live-migration: True
migration-auth-type: ssh
```

Use this command to deploy OpenStack services by using a YAML-formatted file:

```
juju deploy cs:xenial/nova-compute --config ./nova-compute-config.yaml
```

- **By using CLI**

To deploy OpenStack services through the CLI:

```
juju deploy cs:xenial/nova-cloud-controller --config
console-access-protocol=novnc --config openstack-origin=cloud:xenial-ocata
```

- **By using a combination of YAML-formatted file and CLI**

To deploy OpenStack services by using a combination of YAML-formatted file and CLI:

NOTE: Use the `--to <machine number>` command to point to a machine or container where you want the application to be deployed.

```
juju deploy cs:xenial/ntp
juju deploy cs:xenial/rabbitmq-server --to lxd:0
juju deploy cs:xenial/percona-cluster mysql --config
root-password=<root-password> --config max-connections=1500 --to lxd:0
juju deploy cs:xenial/openstack-dashboard --config
openstack-origin=cloud:xenial-ocata --to lxd:0
juju deploy cs:xenial/nova-cloud-controller --config
console-access-protocol=novnc --config openstack-origin=cloud:xenial-ocata
--config network-manager=Neutron --to lxd:0
juju deploy cs:xenial/neutron-api --config
manage-neutron-plugin-legacy-mode=false --config
openstack-origin=cloud:xenial-ocata --config neutron-security-groups=true --to
lxd:0
juju deploy cs:xenial/glance --config openstack-origin=cloud:xenial-ocata --to
lxd:0
```



```
juju deploy cs:xenial/keystone --config admin-password=<admin-password> --config
admin-role=admin --config openstack-origin=cloud:xenial-ocata --to lxd:0
```

NOTE: You set OpenStack services on different machines or on different containers to prevent HAProxy conflicts from applications.

3. Deploy and configure nova-compute.

```
juju deploy cs:xenial/nova-compute --config ./nova-compute-config.yaml --to 1
```

NOTE: You can deploy nova-compute to more than one compute machine.

(Optional) To add additional computes:

```
juju add-unit nova-compute --to 3 # Add one more unit
```

4. Deploy and configure Contrail services.

```
juju deploy --series=xenial
$CHARMS_DIRECTORY/contrail-charms/contrail-keystone-auth --to 2
juju deploy --series=xenial $CHARMS_DIRECTORY/contrail-charms/contrail-controller
--config auth-mode=rbac --config cassandra-minimum-diskgb=4 --config
cassandra-jvm-extra-opts="-Xms1g -Xmx2g" --to 2
juju deploy --series=xenial $CHARMS_DIRECTORY/contrail-charms/contrail-analyticsdb
cassandra-minimum-diskgb=4 --config cassandra-jvm-extra-opts="-Xms1g -Xmx2g"
--to 2
juju deploy --series=xenial $CHARMS_DIRECTORY/contrail-charms/contrail-analytics
--to 2
juju deploy --series=xenial $CHARMS_DIRECTORY/contrail-charms/contrail-openstack
juju deploy --series=xenial $CHARMS_DIRECTORY/contrail-charms/contrail-agent
```

5. Enable applications to be available to external traffic:

```
juju expose openstack-dashboard
juju expose nova-cloud-controller
```

```
juju expose neutron-api
juju expose glance
juju expose keystone
```

6. Enable contrail-controller and contrail-analytics services to be available to external traffic if you do not use HAProxy.

```
juju expose contrail-controller
juju expose contrail-analytics
```

7. Apply SSL.

You can apply SSL if needed. To use SSL with Contrail services, deploy easy-rsa service and **add-relation** command to create relations to contrail-controller service and contrail-agent services.

```
juju deploy cs:~containers/xenial/easyrsa --to 0
juju add-relation easyrsa contrail-controller
juju add-relation easyrsa contrail-agent
```

8. (Optional) HA configuration.

If you use more than one controller, follow the HA solution given below:

- a. Deploy HAProxy and Keepalived services.

HAProxy charm is deployed on machines with Contrail controllers. HAProxy charm must have **peering_mode** set to **active-active**. If **peering_mode** is set to **active-passive**, HAProxy creates additional listeners on the same ports as other Contrail services. This leads to port conflicts.

Keepalived charm does not require **to** option.

```
juju deploy cs:xenial/haproxy --to <first contrail-controller machine> --config
  peering_mode=active-active
juju add-unit haproxy --to <another contrail-controller machine>
juju deploy cs:~boucherv29/keepalived-19 --config virtual_ip=<vip>
```

- b. Enable HAProxy to be available to external traffic.

```
juju expose haproxy
```

NOTE: If you enable HAProxy to be available to external traffic, do not follow step 6.

- c. Add HAProxy and Keepalived relations.

```
juju add-relation haproxy:juju-info keepalived:juju-info
juju add-relation contrail-analytics:http-services haproxy
juju add-relation contrail-controller:http-services haproxy
juju add-relation contrail-controller:https-services haproxy
```

- d. Configure contrail-controller service with VIP.

```
juju set contrail-controller vip=<vip>
```

9. Add other necessary relations.

```
juju add-relation keystone:shared-db mysql:shared-db
juju add-relation glance:shared-db mysql:shared-db
juju add-relation keystone:identity-service glance:identity-service
juju add-relation nova-cloud-controller:image-service glance:image-service
juju add-relation nova-cloud-controller:identity-service keystone:identity-service
juju add-relation nova-cloud-controller:cloud-compute nova-compute:cloud-compute
juju add-relation nova-compute:image-service glance:image-service
juju add-relation nova-compute:amqp rabbitmq-server:amqp
juju add-relation nova-cloud-controller:shared-db mysql:shared-db
juju add-relation nova-cloud-controller:amqp rabbitmq-server:amqp
juju add-relation openstack-dashboard:identity-service keystone

juju add-relation neutron-api:shared-db mysql:shared-db
juju add-relation neutron-api:neutron-api nova-cloud-controller:neutron-api
juju add-relation neutron-api:identity-service keystone:identity-service
juju add-relation neutron-api:amqp rabbitmq-server:amqp

juju add-relation contrail-controller ntp
juju add-relation nova-compute:juju info ntp:juju info

juju add-relation contrail-controller contrail-keystone-auth
juju add-relation contrail-keystone-auth keystone
juju add-relation contrail-controller contrail-analytics
juju add-relation contrail-controller contrail-analyticsdb
```

```
juju add-relation contrail-analytics contrail-analyticsdb

juju add-relation contrail-openstack neutron-api
juju add-relation contrail-openstack nova-compute
juju add-relation contrail-openstack contrail-controller

juju add-relation contrail-agent:juju info nova-compute:juju info
juju add-relation contrail-agent contrail-controller
```

Options for Juju Charms

Each Contrail Charm has a specific set of options. The options you choose depend on the charms you select. The following tables list the various options you can choose:

- Options for **contrail-agent** Charms.

Table 24: Options for contrail-agent

Option	Default option	Description
physical-interface		Specify the interface where you want to install vhost0 on. If you do not specify an interface, vhost0 is installed on the default gateway interface.
vhost-gateway	auto	Specify the gateway for vhost0. You can enter either an IP address or the keyword (auto) to automatically set a gateway based on the existing vhost routes.
remove-juju-bridge	true	To install vhost0 directly on the interface, enable this option to remove any bridge created to deploy LXD/LXC and KVM workloads.
dppdk	false	Specify DPDK vRouter.
dppdk-driver	uio_pci_generic	Specify DPDK driver for the physical interface.
dppdk-hugepages	70%	Specify the percentage of huge pages reserved for DPDK vRouter and OpenStack instances.
dppdk-coremask	1	Specify the vRouter CPU affinity mask to determine on which CPU the DPDK vRouter will run.
dppdk-main-mempool-size		Specify the main packet pool size.
dppdk-pmd-txd-size		Specify the DPDK PMD Tx Descriptor size.
dppdk-pmd-rxd-size		Specify the DPDK PMD Rx Descriptor size.
docker-registry	opencontrailnightly	Specify the URL of the docker-registry.
docker-registry-insecure	false	Specify if the docker-registry should be configured.
docker-user		Log in to the docker registry.
docker-password		Specify the docker-registry password.
image-tag	latest	Specify the docker image tag.

Table 24: Options for contrail-agent (continued)

Option	Default option	Description
log-level	SYS_NOTICE	Specify the log level for Contrail services. Options: SYS_EMERG , SYS_ALERT , SYS_CRIT , SYS_ERR , SYS_WARN , SYS_NOTICE , SYS_INFO , SYS_DEBUG
http_proxy		Specify URL.
https_proxy		Specify URL.
kernel-hugepages-1g	Parameter not enabled by default. NOTE: 2MB huge pages for kernel-mode vRouters are enabled by default.	Specify the number of 1G huge pages for use with vRouters in kernel mode. You can enable huge pages to avoid compute node reboots during software upgrades. This parameter must be specified at initial deployment. It cannot be modified in an active deployment. If you need to migrate to huge page usage in an active deployment, use 2MB huge pages if suitable for your environment. We recommend allotting 2GB of memory—either using the default 1024x2MB huge page size setting or the 2x1GB size setting—for huge pages. Other huge page size settings should only be set by expert users in specialized circumstances. 1GB and 2MB huge pages cannot be enabled simultaneously in environments using Juju. If you are using this command parameter to enable 1GB huge pages, you must also disable 2MB huge pages. 2MB huge pages can be disabled by entering the juju config contrail-agent kernel-hugepages-2m="" command with an empty value. A compute node reboot is required to enable a huge page setting configuration change. After this initial reboot, compute nodes can complete software upgrades without a reboot. Huge pages are disabled for kernel-mode vRouters if the kernel-hugepages-1g and the kernel-hugepages-2m options are not set. This parameter was introduced in Contrail Networking Release 2005.

Table 24: Options for contrail-agent (continued)

Option	Default option	Description
kernel-hugepages-2m	1024	Specify the number of 2MB huge pages for use with vRouters in kernel mode. Huge pages in Contrail Networking are used primarily to allocate flow and bridge table memory within the vRouter. Huge pages for kernel-mode vRouters provide enough flow and bridge table memory to avoid compute node reboots to complete future Contrail Networking software upgrades. 1024x2MB huge pages are configured by default starting in Contrail Networking Release 2005. A compute node reboot is required to enable a kernel-mode vRouter huge page setting configuration change, however, so this huge page setting is not enabled on a compute node until the compute node is rebooted. After a compute node is rebooted to enable a vRouter huge page setting, compute nodes can complete software upgrades without a reboot. We recommend allotting 2GB of memory—either using the default 1024x2MB huge page size setting or the 2x1GB size setting—for kernel-mode vRouter huge pages. Other huge page size settings should only be set by expert users in specialized circumstances. 1GB and 2MB huge pages cannot be enabled simultaneously in environments using Juju. If you are using this command parameter to enable 2MB huge pages, you must also disable 1GB huge pages. 1GB huge pages are disabled by default and can also be disabled by entering the juju config contrail-agent kernel-hugepages-1g="" command with an empty value. 1GB huge pages can only be enabled at initial deployment; you cannot initially enable 1GB huge pages in an active deployment. Huge pages are disabled for kernel-mode vRouters if the kernel-hugepages-1g and the kernel-hugepages-2m options are not set.

Table 24: Options for contrail-agent (continued)

Option	Default option	Description
no_proxy		Specify the list of destinations that must be directly accessed.

- Options for **contrail-analytics** Charms.

Table 25: Options for contrail-analytics

Option	Default option	Description
control-network		Specify the IP address and network mask of the control network.
docker-registry		Specify the URL of the docker-registry.
docker-registry-insecure	false	Specify if the docker-registry should be configured.
docker-user		Log in to the docker registry.
docker-password		Specify the docker-registry password.
image-tag		Specify the docker image tag.
log-level	SYS_NOTICE	Specify the log level for Contrail services. Options: SYS_EMERG , SYS_ALERT , SYS_CRIT , SYS_ERR , SYS_WARN , SYS_NOTICE , SYS_INFO , SYS_DEBUG
http_proxy		Specify URL.
https_proxy		Specify URL.
no_proxy		Specify the list of destinations that must be directly accessed.

- Options for **contrail-analyticsdb** Charms.

Table 26: Options for contrail-analyticsdb

Option	Default option	Description
control-network		Specify the IP address and network mask of the control network.

Table 26: Options for contrail-analyticsdb (continued)

Option	Default option	Description
<code>cassandra-minimum-diskgb</code>	256	Specify the minimum disk requirement.
<code>cassandra-jvm-extra-opts</code>		Specify the memory limit.
<code>docker-registry</code>		Specify the URL of the docker-registry.
<code>docker-registry-insecure</code>	false	Specify if the docker-registry should be configured.
<code>docker-user</code>		Log in to the docker registry.
<code>docker-password</code>		Specify the docker-registry password.
<code>image-tag</code>		Specify the docker image tag.
<code>log-level</code>	SYS_NOTICE	Specify the log level for Contrail services. Options: SYS_EMERG , SYS_ALERT , SYS_CRIT , SYS_ERR , SYS_WARN , SYS_NOTICE , SYS_INFO , SYS_DEBUG
<code>http_proxy</code>		Specify URL.
<code>https_proxy</code>		Specify URL.
<code>no_proxy</code>		Specify the list of destinations that must be directly accessed.

- Options for **contrail-controller** Charms.

Table 27: Options for contrail-controller

Option	Default option	Description
<code>control-network</code>		Specify the IP address and network mask of the control network.
<code>auth-mode</code>	rbac	Specify the authentication mode. Options: rbac , cloud-admin , no-auth . For more information, see https://github.com/Juniper/contrail-controller/wiki/RBAC .
<code>cassandra-minimum-diskgb</code>	20	Specify the minimum disk requirement.

Table 27: Options for contrail-controller (*continued*)

Option	Default option	Description
<code>cassandra-jvm-extra-opts</code>		Specify the memory limit.
<code>cloud-admin-role</code>	<code>admin</code>	Specify the role name in keystone for users who have admin-level access.
<code>global-read-only-role</code>		Specify the role name in keystone for users who have read-only access.
<code>vip</code>		Specify if the Contrail API VIP is used for configuring client-side software. If not specified, private IP of the first Contrail API VIP unit will be used.
<code>use-external-rabbitmq</code>	<code>false</code>	To enable the Charm to use the internal RabbitMQ server, set <code>use-external-rabbitmq</code> to <code>false</code>. To use an external AMQP server, set <code>use-external-rabbitmq</code> to <code>true</code>. NOTE: Do not change the flag after deployment.
<code>flow-export-rate</code>	<code>0</code>	Specify how many flow records are exported by vRouter agent to the Contrail Collector when a flow is created or deleted.
<code>docker-registry</code>		Specify the URL of the docker-registry.
<code>docker-registry-insecure</code>	<code>false</code>	Specify if the docker-registry should be configured.
<code>docker-user</code>		Log in to the docker registry.
<code>docker-password</code>		Specify the docker-registry password.
<code>image-tag</code>		Specify the docker image tag.
<code>log-level</code>	<code>SYS_NOTICE</code>	Specify the log level for Contrail services. Options: <code>SYS_EMERG</code>, <code>SYS_ALERT</code>, <code>SYS_CRIT</code>, <code>SYS_ERR</code>, <code>SYS_WARN</code>, <code>SYS_NOTICE</code>, <code>SYS_INFO</code>, <code>SYS_DEBUG</code>
<code>http_proxy</code>		Specify URL.
<code>https_proxy</code>		Specify URL.

Table 27: Options for contrail-controller (*continued*)

Option	Default option	Description
no_proxy		Specify the list of destinations that must be directly accessed.

- Options for **contrail-keystone-auth** Charms.

Table 28: Options for contrail-keystone-auth

Option	Default option	Description
ssl_ca		Specify if the base64-encoded SSL CA certificate is provided to Contrail keystone clients. NOTE: This certificate is required if you use a privately signed <code>ssl_cert</code> and <code>ssl_key</code> .

- Options for **contrail-openstack** Charms.

Table 29: Options for contrail-controller

Option	Default option	Description
enable-metadata-server	true	Set enable-metadata-server to true to configure metadata and enable nova to run a local instance of nova-api-metadata for virtual machines
use-internal-endpoints	false	Set use-internal-endpoints to true for OpenStack to configure services to use internal endpoints.
heat-plugin-dirs	<code>/usr/lib64/heat,/usr/lib/heat/usr/lib/python2.7/dist-packages/vnc_api/gen/heat/resources</code>	Specify the heat plugin directories.
docker-registry		Specify the URL of the docker-registry.
docker-registry-insecure	false	Specify if the docker-registry should be configured.
docker-user		Log in to the docker registry.
docker-password		Specify the docker-registry password.
image-tag		Specify the docker image tag.

Table 29: Options for contrail-controller (*continued*)

Option	Default option	Description
log-level	SYS_NOTICE	Specify the log level for Contrail services. Options: SYS_EMERG, SYS_ALERT, SYS_CRIT, SYS_ERR, SYS_WARN, SYS_NOTICE, SYS_INFO, SYS_DEBUG
http_proxy		Specify URL.
https_proxy		Specify URL.
no_proxy		Specify the list of destinations that must be directly accessed.

SEE ALSO

[Bundle Reference File](#)
[Installing Contrail with OpenStack by Using Juju Charms | 412](#)

Installing Contrail with Kubernetes by Using Juju Charms

IN THIS SECTION

- [Understanding Juju Charms with Kubernetes | 438](#)
- [Preparing to Deploy Contrail with Kubernetes by Using Juju Charms | 438](#)
- [Deploying Contrail Charms with Kubernetes | 440](#)

You can deploy Contrail Networking using Juju Charms. Juju helps you deploy, configure, and efficiently manage applications on private clouds and public clouds. Juju accesses the cloud with the help of a Juju controller. A Charm is a module containing a collection of scripts and metadata and is used with Juju to deploy Contrail.

A Juju Charm helps you deploy Docker containers to the cloud. For more information on containerized Contrail, see [“Understanding Contrail Containers” on page 5](#). Juju Charms simplifies Contrail deployment by providing a simple way to deploy, configure, scale, and manage Contrail operations.

Understanding Juju Charms with Kubernetes

Contrail supports the following charms:

- contrail-agent
- contrail-analytics
- contrail-analyticsdb
- contrail-controller
- contrail-kubernetes-master
- contrail-kubernetes-node

Preparing to Deploy Contrail with Kubernetes by Using Juju Charms

You can deploy Contrail Networking by using Juju bundle.

Follow these steps to prepare for deployment:

1. Install Juju.

```
apt install bridge-utils -y
apt install snapd -y
snap install juju --classic
```

2. Configure Juju.

You can add a cloud to Juju, identify clouds supported by Juju, and manage clouds already added to Juju.

Adding a cloud

Juju already has knowledge of the AWS cloud, which means adding your AWS account to Juju is quick and easy.

```
juju show-cloud --local aws
```

NOTE: In versions prior to Juju v.2.6.0 the **show-cloud** command only operates locally. There is no **--local** option.

You must ensure that Juju's information is up to date (e.g. new region support). Run the following command to update Juju's public cloud data:

```
juju update-public-clouds
```

Juju recognizes a wide range of cloud types. You can use any one of the following methods to add a cloud credentials to Juju:

- **Adding a Cloud Credentials by Using Interactive Command**

Example: Adding AWS cloud credentials to Juju

```
juju add-credential aws
Enter credential name: jlaurin

Using auth-type "access-key".

Enter access-key: AKIAIFII5EH5F0CYZJMA

Enter secret-key: *****

Credential "jlaurin" added locally for cloud "aws".
```

- **Adding a Cloud Credentials Manually**

You can use a YAML configuration file to add AWS cloud credentials. Run the following command:

```
juju add-credential aws -f <mycreds.yaml>
```

For details, refer to [Juju Adding Credentials from a File](#).

Identifying a supported cloud

Use the **juju clouds** command to list cloud types that are supported by Juju.

```
$ juju clouds
```

Cloud	Regions	Default	Type	Description
aws	15	us-east-1	ec2	Amazon Web Services

aws-china	1	cn-north-1	ec2	Amazon China
aws-gov	1	us-gov-west-1	ec2	Amazon (USA Government)
azure	26	centralus	azure	Microsoft Azure
azure-china	2	chinaeast	azure	Microsoft Azure China
cloudsigma	5	hnl	cloudsigma	CloudSigma Cloud
google	13	us-east1	gce	Google Cloud Platform
joyent	6	eu-ams-1	joyent	Joyent Cloud
oracle	5	uscom-central-1	oracle	Oracle Cloud
rackspace	6	dfw	rackspace	Rackspace Cloud
localhost	1	localhost	lxd	LXD Container Hypervisor

3. Create a Juju controller.

```
juju bootstrap --bootstrap-series=xenial <cloud name> <controller name>
```

A Juju controller manages and keeps track of applications in the Juju cloud environment.

4. Download the Contrail bundle from [JAAS - Contrail Kubernetes](#).

Deploying Contrail Charms with Kubernetes

IN THIS SECTION

- [Deploying Contrail Charms in a Bundle | 440](#)
- [Deploying Juju Charms with Kubernetes Manually | 446](#)

Juju Charms simplifies Contrail deployment by providing a simple way to deploy, configure, scale, and manage Contrail operations.

You can deploy Contrail Charms in a bundle or manually.

Deploying Contrail Charms in a Bundle

Follow these steps to deploy Contrail Charms in a bundle.

1. Deploy Contrail Charms.

To deploy Contrail Charms in a bundle, use the **juju deploy <bundle_yaml_file>** command.

The following example shows you how to use a bundle YAML file to deploy Contrail on Amazon Web Services (AWS) Cloud.

```

series: "bionic"

machines:

  # kubernetes pods
  0:
    series: "bionic"
    constraints: mem=8G cores=2 root-disk=60G

  # kubernetes master
  2:
    series: "bionic"
    constraints: mem=8G cores=2 root-disk=60G

  # contrail components
  5:
    series: "bionic"
    constraints: mem=16G cores=4 root-disk=60G

services:

  # kubernetes

  easyrsa:
    series: "bionic"
    charm: cs:~containers/easyrsa
    num_units: 1
    annotations:
      gui-x: '1168.1039428710938'
      gui-y: '-59.11077045466004'
    to:
      - lxd:2

  etcd:
    series: "bionic"
    charm: cs:~containers/etcd
    annotations:
      gui-x: '1157.2041015625'
      gui-y: '719.1614406201691'
    num_units: 1
    options:
      channel: 3.2/stable

```

```

to: [2]

kubernetes-master:
  series: "bionic"
  charm: cs:~containers/kubernetes-master-696
  annotations:
    gui-x: '877.1133422851562'
    gui-y: '325.6035540382413'
  expose: true
  num_units: 1
  options:
    channel: '1.14/stable'
    service-cidr: '10.96.0.0/12'
    docker_runtime: 'custom'
    docker_runtime_repo: 'deb [arch={ARCH}]
https://download.docker.com/linux/ubuntu {CODE} stable'
    docker_runtime_key_url: 'https://download.docker.com/linux/ubuntu/gpg'
    docker_runtime_package: 'docker-ce'
  to: [2]

kubernetes-worker:
  series: "bionic"
  charm: cs:~containers/kubernetes-worker-550
  annotations:
    gui-x: '745.8510131835938'
    gui-y: '-57.369691124215706'
  num_units: 1
  options:
    channel: '1.14/stable'
    docker_runtime: 'custom'
    docker_runtime_repo: 'deb [arch={ARCH}]
https://download.docker.com/linux/ubuntu {CODE} stable'
    docker_runtime_key_url: 'https://download.docker.com/linux/ubuntu/gpg'
    docker_runtime_package: 'docker-ce'
  to: [0]

# contrail-kubernetes

contrail-kubernetes-master:
  series: "bionic"
  charm: cs:~juniper-os-software/contrail-kubernetes-master
  annotations:
    gui-x: '586.8027801513672'
    gui-y: '753.914497641757'

```

```

options:
  log-level: 'SYS_DEBUG'
  service_subnets: '10.96.0.0/12'
  docker-registry: "opencontrailnightly"
  image-tag: "master-latest"

contrail-kubernetes-node:
  series: "bionic"
  charm: cs:~juniper-os-software/contrail-kubernetes-node
  annotations:
    gui-x: '429.1971130371094'
    gui-y: '216.05209087397168'
  options:
    log-level: 'SYS_DEBUG'
    docker-registry: "opencontrailnightly"
    image-tag: "master-latest"

# contrail

contrail-agent:
  series: "bionic"
  charm: cs:~juniper-os-software/contrail-agent
  annotations:
    gui-x: '307.5467224121094'
    gui-y: '-24.150856522753656'
  options:
    log-level: 'SYS_DEBUG'
    docker-registry: "opencontrailnightly"
    image-tag: "master-latest"

contrail-analytics:
  series: "bionic"
  charm: cs:~juniper-os-software/contrail-analytics
  annotations:
    gui-x: '15.948270797729492'
    gui-y: '705.2326686475128'
  expose: true
  num_units: 1
  options:
    log-level: 'SYS_DEBUG'
    docker-registry: "opencontrailnightly"
    image-tag: "master-latest"
  to: [5]

```

```

contrail-analyticsdb:
  series: "bionic"
  charm: cs:~juniper-os-software/contrail-analyticsdb
  annotations:
    gui-x: '24.427139282226562'
    gui-y: '283.9550754931123'
  num_units: 1
  options:
    cassandra-minimum-diskgb: '4'
    cassandra-jvm-extra-opts: '-Xms1g -Xmx2g'
    log-level: 'SYS_DEBUG'
    docker-registry: "opencontrailnightly"
    image-tag: "master-latest"
  to: [5]

contrail-controller:
  series: "bionic"
  charm: cs:~juniper-os-software/contrail-controller
  annotations:
    gui-x: '212.01282501220703'
    gui-y: '480.69961284662793'
  expose: true
  num_units: 1
  options:
    auth-mode: 'no-auth'
    cassandra-minimum-diskgb: '4'
    cassandra-jvm-extra-opts: '-Xms1g -Xmx2g'
    log-level: 'SYS_DEBUG'
    docker-registry: "opencontrailnightly"
    image-tag: "master-latest"
  to: [5]

# misc

ntp:
  charm: "cs:bionic/ntp"
  annotations:
    gui-x: '678.6017761230469'
    gui-y: '415.27124759750086'

relations:

- [ kubernetes-master:kube-api-endpoint, kubernetes-worker:kube-api-endpoint ]

```

```

- [ kubernetes-master:kube-control, kubernetes-worker:kube-control ]
- [ kubernetes-master:certificates, easyrsa:client ]
- [ kubernetes-master:etcd, etcd:db ]
- [ kubernetes-worker:certificates, easyrsa:client ]
- [ etcd:certificates, easyrsa:client ]

# contrail
- [ kubernetes-master, ntp ]
- [ kubernetes-worker, ntp ]
- [ contrail-controller, ntp ]

- [ contrail-controller, contrail-analytics ]
- [ contrail-controller, contrail-analyticsdb ]
- [ contrail-analytics, contrail-analyticsdb ]
- [ contrail-agent, contrail-controller ]

# contrail-kubernetes
- [ contrail-kubernetes-node:cni, kubernetes-master:cni ]
- [ contrail-kubernetes-node:cni, kubernetes-worker:cni ]
- [ contrail-kubernetes-master:contrail-controller,
contrail-controller:contrail-controller ]
- [ contrail-kubernetes-master:kube-api-endpoint,
kubernetes-master:kube-api-endpoint ]
- [ contrail-agent:juju-info, kubernetes-worker:juju-info ]
- [ contrail-agent:juju-info, kubernetes-master:juju-info ]
- [ contrail-kubernetes-master:contrail-kubernetes-config,
contrail-kubernetes-node:contrail-kubernetes-config ]

```

You can create or modify the Contrail Charm deployment bundle YAML file to:

- Point to machines or instances where the Contrail Charms must be deployed.
- Include the options you need.

Each Contrail Charm has a specific set of options. The options you choose depend on the charms you select. For more information on the options that are available, see *config.yaml* file for each charm located at [Contrail Charms](#).

2. (Optional) Check the status of deployment.

You can check the status of the deployment by using the **juju status** command.

3. Enable configuration statements.

Based on your deployment requirements, you can enable the following configuration statements:

- **contrail-agent**

For more information, see

<https://github.com/tungstenfabric/tf-charms/blob/master/contrail-agent/README.md>.

- **contrail-analytics**

For more information, see

<https://github.com/tungstenfabric/tf-charms/blob/master/contrail-analytics/README.md>.

- **contrail-analyticsdb**

For more information, see

<https://github.com/tungstenfabric/tf-charms/blob/master/contrail-analyticsdb/README.md>.

- **contrail-controller**

For more information, see

<https://github.com/tungstenfabric/tf-charms/blob/master/contrail-controller/README.md>.

- **contrail-kubernetes-master**

For more information, see

<https://github.com/tungstenfabric/tf-charms/blob/master/contrail-kubernetes-master/README.md>.

- **contrail-kubernetes-node**

For more information, see

<https://github.com/tungstenfabric/tf-charms/blob/master/contrail-kubernetes-node/README.md>.

Deploying Juju Charms with Kubernetes Manually

Before you begin deployment, ensure that you have:

- Installed and configured Juju
- Created a Juju controller
- Installed Ubuntu 16.04 or Ubuntu 18.04

Follow these steps to deploy Juju Charms with Kubernetes manually:

1. Create machine instances for Kubernetes master, Kubernetes workers, and Contrail.

```
juju add-machine ssh:<sshusername>@<IP> --constraints mem=8G cores=2 root-disk=32G
--series=xenial #for Kubernetes worker machine
```

```
juju add-machine ssh:<sshusername>@<IP> --constraints mem=18G cores=2
root-disk=32G --series=xenial #for Kubernetes master machine
```

```
juju add-machine ssh:<sshusername>@<IP> --constraints mem=16G cores=4
root-disk=32G --series=xenial #for Contrail machine
```

2. Deploy the Kubernetes services.

Some of the applications may need an additional configuration.

You can deploy Kubernetes services using any one of the following methods:

- By specifying the Kubernetes parameters in a YAML file
- By using CLI
- By using a combination of YAML-formatted file and CLI

NOTE: You must use the same docker version for Contrail and Kubernetes.

For more details, refer to [Juju Application Configuration](#).

3. Deploy and configure **ntp**, **easysrsa**, **etcd**, **kubernetes-master**, **kubernetes-worker**.

```
juju deploy cs:xenial/ntp ntp

juju deploy cs:~containers/easysrsa easysrsa --to lxd:0

juju deploy cs:~containers/etcd etcd \
    --resource etcd=3 \
    --resource snapshot=0
juju set etcd channel="3.2/stable"

juju deploy cs:~containers/kubernetes-master kubernetes-master \
    --resource cdk-addons=0 \
```

```

--resource kube-apiserver=0 \
--resource kube-controller-manager=0 \
--resource kube-proxy=0 \
--resource kube-scheduler=0 \
--resource kubectl=0
juju set kubernetes-master channel="1.14/stable" \
enable-dashboard-addons="false" \
enable-metrics="false" \
dns-provider="none" \
docker_runtime="custom" \
docker_runtime_repo="deb [arch={ARCH}]
https://download.docker.com/linux/ubuntu {CODE} stable" \
docker_runtime_key_url="https://download.docker.com/linux/ubuntu/gpg" \
docker_runtime_package="docker-ce"

juju deploy cs:~containers/kubernetes-worker kubernetes-worker \
--resource kube-proxy="0" \
--resource kubectl="0" \
--resource kubelet="0"
juju set kubernetes-worker channel="1.14/stable" \
ingress="false" \
docker_runtime="custom" \
docker_runtime_repo="deb [arch={ARCH}]
https://download.docker.com/linux/ubuntu {CODE} stable" \
docker_runtime_key_url="https://download.docker.com/linux/ubuntu/gpg" \
docker_runtime_package="docker-ce"

```

4. Deploy and configure Contrail services.

Deploy **contrail-analyticsdb**, **contrail-analytics**, **contrail-controller**, **contrail-kubernetes-master**, **contrail-kubernetes-node**, **contrail-agent** from the directory where you have downloaded the charms.

NOTE: You must set the *auth-mode* parameter of the **contrail-controller** charm to **no-auth** if Contrail is deployed without a keystone.

```

juju deploy contrail-analytics contrail-analytics

juju deploy contrail-analyticsdb contrail-analyticsdb
juju set contrail-analyticsdb cassandra-minimum-diskgb="4"
cassandra-jvm-extra-opts="-Xms1g -Xmx2g"

```



```
juju deploy contrail-controller contrail-controller
juju set contrail-controller cassandra-minimum-diskgb="4"
cassandra-jvm-extra-opts="-Xms1g -Xmx2g" auth-mode="no-auth"

juju deploy contrail-kubernetes-master contrail-kubernetes-master

juju deploy contrail-kubernetes-node contrail-kubernetes-node

juju deploy contrail-agent contrail-agent
```

5. Enable applications to be available to external traffic:

```
juju expose kubernetes-master
juju expose kubernetes-worker
```

6. Enable contrail-controller and contrail-analytics services to be available to external traffic if you do not use HAProxy.

```
juju expose contrail-controller
juju expose contrail-analytics
```

7. Apply SSL.

You can apply SSL if needed. To use SSL with Contrail services, deploy easy-rsa service and **add-relation** command to create relations to contrail-controller service and contrail-agent services.

```
juju add-relation easyrsa contrail-controller
juju add-relation easyrsa contrail-analytics
juju add-relation easyrsa contrail-analyticsdb
juju add-relation easyrsa contrail-kubernetes-master
juju add-relation easyrsa contrail-agent
```

8. Add other necessary relations.

```
juju add-relation "contrail-controller" "contrail-analytics"
juju add-relation "contrail-controller" "contrail-analyticsdb"
juju add-relation "contrail-analytics" "contrail-analyticsdb"
juju add-relation "contrail-agent" "contrail-controller"
juju add-relation "contrail-controller" "ntp"
juju add-relation "kubernetes-worker", "ntp"
```

```

juju add-relation "kubernetes-master", "ntp"

juju add-relation "kubernetes-master:kube-api-endpoint"
"kubernetes-worker:kube-api-endpoint"
juju add-relation "kubernetes-master:kube-control"
"kubernetes-worker:kube-control"
juju add-relation "kubernetes-master:certificates" "easyrsa:client"
juju add-relation "kubernetes-master:etcd" "etcd:db"
juju add-relation "kubernetes-worker:certificates" "easyrsa:client"
juju add-relation "etcd:certificates" "easyrsa:client"

juju add-relation contrail-agent:juju-info, kubernetes-master:juju-info

juju add-relation "contrail-kubernetes-node:cni" "kubernetes-master:cni"
juju add-relation "contrail-kubernetes-node:cni" "kubernetes-worker:cni"
juju add-relation "contrail-kubernetes-master:contrail-controller"
"contrail-controller:contrail-controller"
juju add-relation "contrail-kubernetes-master:kube-api-endpoint"
"kubernetes-master:kube-api-endpoint"
juju add-relation "contrail-agent:juju-info" "kubernetes-worker:juju-info"
juju add-relation "contrail-agent:juju-info" "kubernetes-master:juju-info"
juju add-relation "contrail-kubernetes-master:contrail-kubernetes-config"
"contrail-kubernetes-node:contrail-kubernetes-config"

```

RELATED DOCUMENTATION

| <https://juju.is/docs/installing>

Installing Contrail with Kubernetes in Nested Mode by Using Juju Charms

Contrail Networking Release 1909 and later support provisioning of a Kubernetes cluster inside an OpenStack cluster. Contrail Networking offers a nested control and data plane where a single Contrail control plane and a single network stack can manage and service both the OpenStack and Kubernetes clusters.

In nested mode, a Kubernetes cluster is provisioned in virtual machines of an OpenStack cluster. The CNI plugin and the Contrail-Kubernetes manager of the Kubernetes cluster interface directly with Contrail components that manage the OpenStack cluster.

All Kubernetes features, functions and specifications are supported when used in nested mode.

NOTE: Nested mode deployment is only supported for Contrail with OpenStack cluster.

Before you begin:

- Deploy Contrail with OpenStack either on bare metal server or virtual machines.

BEST PRACTICE: Public cloud deployment is not recommended because of slow nested virtualization.

- The VMs must have internet connectivity.
- Contrail in underlay network must be configured to support nested mode.

You must select an unused IP in the cluster to configure *link-local*.

For example:

10.10.10.5 is the selected service IP.

LL Service Name	Service IP	Service Port	Fabric IP	Fabric Port
K8s-cni-to-agent	10.10.10.5	9091	127.0.0.1	9091

Follow these steps to deploy Juju Charms with Kubernetes in nested mode using bundle deployment:

Use this method if you want to use the existing machines.

1. Create a Juju controller.

juju bootstrap --bootstrap-series=xenial <cloud name> <controller name>

You can use OpenStack Cloud provider or manually spun-up VMs. For details, refer to [Preparing to Deploy Contrail with Kubernetes by Using Juju Charms](#).

2. Deploy bundle.

```
juju deploy --series xenial cs:~containers/kubernetes-worker-550 --to:0 \
--config channel="1.14/stable" \
--config docker_runtime="custom" \
```

If the machines for the setup are already provisioned, run the following command to deploy bundle:

```
juju deploy --map-machines=existing,0=0,5=1 ./bundle.yaml
where bundle-id=existing-id
```

For details, refer to <https://jaas.ai/u/juniper-os-software/contrail-k8s-nested/bundle>.

or

Follow these steps to deploy Juju Charms with Kubernetes in nested mode manually:

1. Create a Juju controller.

```
juju bootstrap --bootstrap-series=xenial <cloud name> <controller name>
```

You can use OpenStack Cloud provider or manually spun-up VMs. For details, refer to [Preparing to Deploy Contrail with Kubernetes by Using Juju Charms](#).

2. Create machine instances for Contrail components, Kubernetes master and Kubernetes workers.

Sample constraints for minimal deployment:

All-In-One deployment:

```
juju add-machine --constraints mem=32G cores=8 root-disk=150G --series=xenial # for all-in-one machine
```

or

Multinode deployment:

```
juju add-machine --constraints mem=8G cores=2 root-disk=50G --series=xenial # kubernetes workers
juju add-machine --constraints mem=8G cores=2 root-disk=50G --series=xenial # kubernetes masters
juju add-machine --constraints mem=4G cores=4 root-disk=50G --series=xenial # contrail components
```

You can use any series—*xenial* or *bionic*.

3. Add machines to the cloud.

For details, refer to [Using Constraints-Juju](#).

4. Deploy the Kubernetes services.

Some of the applications may need additional configuration.

You can deploy Kubernetes services using any one of the following methods:

- By specifying the Kubernetes parameters in a YAML file.
- By passing options/values directly on the command line.

NOTE: You must use the same docker version for Contrail and Kubernetes.

For more details, refer to [Juju Application Configuration](#).

5. Deploy and configure **ntp**, **easysrsa**, **etcd**, **kubernetes-master**, **kubernetes-worker**.

```
juju deploy --series xenial cs:ntp ntp

juju deploy --series xenial cs:~containers/easysrsa --to lxd:0

juju deploy --series xenial cs:~containers/etcd --to:0 --config
channel="3.2/stable"

juju deploy --series xenial cs:~containers/kubernetes-master-696 --to:0 \
  --config channel="1.14/stable" \
  --config docker_runtime="custom" \
  --config docker_runtime_repo="deb [arch={ARCH}]
https://download.docker.com/linux/ubuntu {CODE} stable" \
  --config docker_runtime_key_url="https://download.docker.com/linux/ubuntu/gpg"
\
  --config docker_runtime_package="docker-ce"

juju deploy --series xenial cs:~containers/kubernetes-worker-550 --to:0 \
  --config channel="1.14/stable" \
  --config ingress="false" \
  --config docker_runtime="custom" \
  --config docker_runtime_repo="deb [arch={ARCH}]
https://download.docker.com/linux/ubuntu {CODE} stable" \
  --config docker_runtime_key_url="https://download.docker.com/linux/ubuntu/gpg"
\
  --config docker_runtime_package="docker-ce"
```

6. Deploy and configure Contrail services.

Deploy **contrail-kubernetes-master**, **contrail-kubernetes-node**, **contrail-agent** from the directory where you have downloaded the charms.

```
contrail-kubernetes-master:
  nested_mode: true
  cluster_project: '{"domain': 'default-domain', 'project': 'admin'}'
  cluster_network:
    '{"domain': 'default-domain', 'project': 'admin', 'name': 'juju-net'}'
  service_subnets: '10.96.0.0/12'
  nested_mode_config: |
    {
      "CONTROLLER_NODES": "10.0.12.20",
      "AUTH_MODE": "keystone",
      "KEYSTONE_AUTH_ADMIN_TENANT": "admin",
      "KEYSTONE_AUTH_ADMIN_USER": "admin",
      "KEYSTONE_AUTH_ADMIN_PASSWORD": "password",
      "KEYSTONE_AUTH_URL_VERSION": "/v2.0",
      "KEYSTONE_AUTH_HOST": "10.0.12.122",
      "KEYSTONE_AUTH_PROTO": "http",
      "KEYSTONE_AUTH_PUBLIC_PORT": "5000",
      "KEYSTONE_AUTH_REGION_NAME": "RegionOne",
      "KEYSTONE_AUTH_INSECURE": "True",
      "KUBERNETES_NESTED_VROUTER_VIP": "10.10.10.5"
    }
```

```
juju deploy --series xenial cs:~juniper-os-software/contrail-kubernetes-master
\
  --config ./path-to-config.yaml

juju deploy --series xenial cs:~juniper-os-software/contrail-kubernetes-node
```

7. Add the necessary relations.

```
juju add-relation "kubernetes-master:juju-info" "ntp:juju-info"
juju add-relation "kubernetes-worker:juju-info" "ntp:juju-info"

juju add-relation "kubernetes-master:kube-api-endpoint"
"kubernetes-worker:kube-api-endpoint"
juju add-relation "kubernetes-master:kube-control"
"kubernetes-worker:kube-control"
juju add-relation "kubernetes-master:certificates" "easyrsa:client"
juju add-relation "kubernetes-master:etcd" "etcd:db"
juju add-relation "kubernetes-worker:certificates" "easyrsa:client"
```

```
juju add-relation "etcd:certificates" "easyrsa:client"

juju add-relation "contrail-kubernetes-node:cni" "kubernetes-master:cni"
juju add-relation "contrail-kubernetes-node:cni" "kubernetes-worker:cni"
juju add-relation "contrail-kubernetes-master:kube-api-endpoint"
"kubernetes-master:kube-api-endpoint"
juju add-relation "contrail-kubernetes-master:contrail-kubernetes-config"
"contrail-kubernetes-node:contrail-kubernetes-config"
```

8. Apply SSL, if needed.

You must provide the same certificates to the *contrail-kubernetes-master* node if Contrail in underlay cluster has SSL enabled.

Release History Table

Release	Description
1909	Contrail Networking Release 1909 and later support provisioning of a Kubernetes cluster inside an OpenStack cluster. Contrail Networking offers a nested control and data plane where a single Contrail control plane and a single network stack can manage and service both the OpenStack and Kubernetes clusters.

RELATED DOCUMENTATION

[Installing Contrail with Kubernetes by Using Juju Charms](#) | [437](#)

[Installing Contrail with OpenStack by Using Juju Charms](#) | [412](#)

Installing OpenStack Octavia LBaaS with Juju Charms in Contrail Networking

Contrail Networking Release 2005 supports Octavia as LBaaS. The deployment supports RHOSP and Juju platforms.

With Octavia as LBaaS, Contrail Networking is only maintaining network connectivity and is not involved in any load balancing functions.

For each OpenStack load balancer creation, Octavia launches a VM known as *amphora* VM. The VM starts the HAPROXY when listener is created for the load balancer in OpenStack. Whenever the load balancer gets updated in OpenStack, *amphora* VM updates the running HAPROXY configuration. The *amphora* VM is deleted on deleting the load balancer.

Contrail Networking provides connectivity to *amphora* VM interfaces. *Amphora* VM has two interfaces; one for management and the other for data. The management interface is used by the Octavia services for the management communication. Since, Octavia services are running in the underlay network and *amphora* VM is running in the overlay network, SDN gateway is needed to reach the overlay network. The data interface is used for load balancing.

Follow the procedure to install OpenStack Octavia LBaaS in Canonical deployment:

1. Prepare Juju setup with OpenStack Train version and Octavia overlay bundle.

Refer to [Sample octavia-bundle.yaml file on page 461](#) output.

```
juju deploy --overlay=./octavia-bundle.yaml ./contrail-bundle.yaml
```

or

Add Octavia service after deploying the main bundle on the existing cluster.

```
juju deploy --overlay=./octavia-bundle.yaml --map-machines=existing
./contrail-bundle.yaml
```

2. Prepare ssh key for amphora VM. Add the options in the **octavia-bundle.yaml** file.

```
ssh-keygen -f octavia # generate the key base64 octavia.pub # print public
key data
```

Add the following options to Octavia options.

```
amp-ssh-pub-key: # paste public key data here amp-ssh-key-name: octavia
```

3. Generate certificates.

```
rm -rf demoCA/
mkdir -p demoCA/newcerts
touch demoCA/index.txt
touch demoCA/index.txt.attr
openssl genrsa -passout pass:foobar -des3 -out issuing_ca_key.pem 2048
openssl req -x509 -passin pass:foobar -new -nodes -key issuing_ca_key.pem \
    -config /etc/ssl/openssl.cnf \
    -subj "/C=US/ST=Somestate/O=Org/CN=www.example.com" \
    -days 30 \
    -out issuing_ca.pem
openssl genrsa -passout pass:foobar -des3 -out controller_ca_key.pem 2048
openssl req -x509 -passin pass:foobar -new -nodes \
    -key controller_ca_key.pem \
    -config /etc/ssl/openssl.cnf \
    -subj "/C=US/ST=Somestate/O=Org/CN=www.example.com" \
```



```

    -days 30 \
    -out controller_ca.pem
openssl req \
    -newkey rsa:2048 -nodes -keyout controller_key.pem \
    -subj "/C=US/ST=Somestate/O=Org/CN=www.example.com" \
    -out controller.csr
openssl ca -passin pass:foobar -config /etc/ssl/openssl.cnf \
    -cert controller_ca.pem -keyfile controller_ca_key.pem \
    -create_serial -batch \
    -in controller.csr -days 30 -out controller_cert.pem
cat controller_cert.pem controller_key.pem > controller_cert_bundle.pem
juju config octavia \
    lb-mgmt-issuing-cacert="$(base64 controller_ca.pem)" \
    lb-mgmt-issuing-ca-private-key="$(base64 controller_ca_key.pem)" \
    lb-mgmt-issuing-ca-key-passphrase=foobar \
    lb-mgmt-controller-cacert="$(base64 controller_ca.pem)" \
    lb-mgmt-controller-cert="$(base64 controller_cert_bundle.pem)"

```

Make sure all the units are in *active* or *blocked* state.

4. Configure vault service.

- a. SSH into the machine where vault service is installed.

```
juju ssh vault/0
```

- b. Export vault address and run **init**.

```
export VAULT_ADDR='http://localhost:8200'
/snap/bin/vault operator init -key-shares=5 -key-threshold=3
```

It will print 5 unseal keys and initial root token.

- c. Call unseal command by using any three of the five printed unseal keys.

```
/snap/bin/vault operator unseal Key1
/snap/bin/vault operator unseal Key2
/snap/bin/vault operator unseal Key3
```

- d. Export initial root token.

```
export VAULT_TOKEN="..."
```

- e. Create user token.

```
/snap/bin/vault token create -ttl=10m
```

- f. Exit from vault's machine and initialize vault's charm with the user token.

```
juju run-action --wait vault/leader authorize-charm token="..."
```

5. Create amphora image.

```
juju run-action --wait octavia-diskimage-retrofit/leader retrofit-image
```

For more details, refer to

<https://docs.openstack.org/project-deploy-guide/charm-deployment-guide/latest/app-octavia.html#amphora-image>.

6. Install *python-openstackclient* and *python-octaviaclient* and create management network for Octavia.

You must create these objects in *services* project.

```
project=$(openstack project list --domain service_domain | awk
'/services/{print $2}')
openstack network create octavia --tag charm-octavia --project $project
openstack subnet create --subnet-range 172.x.0.0/24 --network octavia --tag
charm-octavia octavia
# security group for octavia
openstack security group create octavia --tag charm-octavia --project
$project
openstack security group rule create --ingress --ethertype IPv4 --protocol
icmp octavia
openstack security group rule create --ingress --ethertype IPv6 --protocol
icmp octavia
openstack security group rule create --ingress --ethertype IPv4 --protocol
tcp --dst-port 22:22 octavia
openstack security group rule create --ingress --ethertype IPv6 --protocol
tcp --dst-port 22:22 octavia
openstack security group rule create --ingress --ethertype IPv6 --protocol
tcp --dst-port 9443:9443 octavia
openstack security group rule create --ingress --ethertype IPv4 --protocol
tcp --dst-port 9443:9443 octavia
# security group for octavia-health
openstack security group create octavia-health --tag charm-octavia-health
--project $project
openstack security group rule create --ingress --ethertype IPv4 --protocol
icmp octavia-health
openstack security group rule create --ingress --ethertype IPv6 --protocol
icmp octavia-health
```

```
openstack security group rule create --ingress --ethertype IPv4 --protocol
udp --dst-port 5555:5555 octavia-health
openstack security group rule create --ingress --ethertype IPv6 --protocol
udp --dst-port 5555:5555 octavia-health
```

7. The management network created in step 6 is in overlay network and Octavia services are running in the underlay network. Verify network connectivity between overlay and underlay network via SDN gateway.
8. Configure Octavia with the created network.

```
juju run-action --wait octavia/leader configure-resources
```

Make sure the juju cluster is functional and all units have *active* status.

If you want to run amphora instances on DPDK computes, you have to create your own flavor with the required options and set the ID to configuration of Octavia charm via *custom-amp-flavor-id* option before call *configure-resources*.

Or

Set the required options to created flavor with name *charm-octavia* by charm

```
openstack flavor set charm-octavia --property hw:mem_page_size=any
```

Here is an example for creating and testing load balancer:

Prerequisites:

- You must have connectivity between Octavia controller and amphora instances,
- You must have OpenStack services into LXD containers.
- You must have separate interfaces for control plane and data plane.

1. Create private network.

```
openstack network create private
openstack subnet create private --network private --subnet-range
10.10.10.0/24 --allocation-pool
start=10.10.10.50,end=10.10.10.70 --gateway none
```

2. Create security group.

```
openstack security group create allow_all
openstack security group rule create --ingress --protocol any --prefix
'0.0.0.0/0' allow_all
```

3. Check available flavors and images. You can create them, if needed.

```
openstack flavor list
openstack image list
```

4. Create two servers for load balancer.

```
openstack server create --flavor test_flavor --image cirros --security-group
allow_all --network private cirros1
openstack server create --flavor test_flavor --image cirros --security-group
allow_all --network private cirros2
```

5. Create additional server to test load balancer.

```
openstack server create --flavor test_flavor --image cirros --security-group
allow_all --network private cirros-test
```

6. Check status and IP addresses.

```
openstack server list --long
```

7. Create simple HTTP server on every cirros. Login on both the cirros instances and run following commands:

```
MYIP=$(ifconfig eth0|grep 'inet addr'|awk -F: '{print $2}'| awk '{print
$1}') while true; do echo -e "HTTP/1.0 200 OK\r\n\r\nWelcome to $MYIP" |
sudo nc -l -p 80 ; done&
```

8. Create load balancer

```
openstack loadbalancer create --name lb1 --vip-subnet-id private
```

Make sure *provisioning_status* is Active.

```
openstack loadbalancer show lb1
```

9. Setup load balancer

```
openstack loadbalancer listener create --protocol HTTP --protocol-port 80
--name listener1 lb1
openstack loadbalancer show lb1 # Wait for the provisioning_status to be
ACTIVE.
openstack loadbalancer pool create --lb-algorithm ROUND_ROBIN --listener
listener1 --protocol HTTP --name pool1
openstack loadbalancer healthmonitor create --delay 5 --timeout 2
--max-retries 1 --type HTTP pool1
openstack loadbalancer member create --subnet-id private --address
10.10.10.50 --protocol-port 80 pool1
```

```
openstack loadbalancer member create --subnet-id private --address
10.10.10.51 --protocol-port 80 pool1
```

IP addresses 10.10.10.50 and 10.10.10.51 belong to VMs created with test http server in step 7.

10. Check the status of load balancer.

```
openstack loadbalancer show lb1 # Wait for the provisioning_status to be
ACTIVE.
openstack loadbalancer pool list
openstack loadbalancer pool show pool1
openstack loadbalancer member list pool1
openstack loadbalancer listener list
```

11. Login to load balancer client and verify if round robin works.

```
ubuntu@comp-1:~$ ssh cirros@169.x.0.9
The authenticity of host '169.x.0.9 (169.x.0.9)' can't be established.
RSA key fingerprint is SHA256:jv0qgZkorxxxxxxxxmyk0SVQV3fFl0.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '169.x.0.9' (RSA) to the list of known hosts.
cirros@169.x.0.9's password:
$ curl 10.10.10.50
Welcome to 10.10.10.52
$ curl 10.10.10.50
Welcome to 10.10.10.53
$ curl 10.10.10.50
Welcome to 10.10.10.52
$ curl 10.10.10.50
Welcome to 10.10.10.53
$ curl 10.10.10.50
Welcome to 10.10.10.52
$ curl 10.10.10.50
Welcome to 10.10.10.53
```

Sample octavia-bundle.yaml file

```
# Loadbalancer (LBAASv2) with Octavia - requires Rocky or later
---
applications:
  barbican:
    charm: cs:barbican
    num_units: 1
    options:
```

```

    openstack-origin: cloud:bionic-train
  to:
  - lxd:4
barbican-vault:
  charm: cs:barbican-vault-12
octavia:
  series: bionic
  charm: cs:~apavlov-e/octavia
  num_units: 1
  options:
    openstack-origin: cloud:bionic-train
    create-mgmt-network: false
  to:
  - lxd:4
octavia-dashboard:
  charm: cs:octavia-dashboard
vault:
  charm: cs:vault
  num_units: 1
  to:
  - lxd:4
glance-simplestreams-sync:
  charm: cs:glance-simplestreams-sync
  num_units: 1
  options:
    source: ppa:simplestreams-dev/trunk
    use_swift: false
  to:
  - lxd:4
octavia-diskimage-retrofit:
  charm: cs:octavia-diskimage-retrofit
  options:
    amp-image-tag: 'octavia-amphora'
    retrofit-uca-pocket: train
relations:
- - mysql:shared-db
  - octavia:shared-db
- - mysql:shared-db
  - barbican:shared-db
- - mysql:shared-db
  - vault:shared-db
- - keystone:identity-service
  - octavia:identity-service
- - keystone:identity-service

```

- barbican:identity-service
- - rabbitmq-server:amqp
- octavia:amqp
- - rabbitmq-server:amqp
- barbican:amqp
- - neutron-api:neutron-load-balancer
- octavia:neutron-api
- - openstack-dashboard:dashboard-plugin
- octavia-dashboard:dashboard
- - barbican-vault:secrets
- barbican:secrets
- - vault:secrets
- barbican-vault:secrets-storage
- - glance-simplestreams-sync:juju-info
- octavia-diskimage-retrofit:juju-info
- - keystone:identity-service
- glance-simplestreams-sync:identity-service
- - rabbitmq-server:amqp
- glance-simplestreams-sync:amqp
- - keystone:identity-credentials
- octavia-diskimage-retrofit:identity-credentials
- - contrail-openstack
- octavia

Release History Table

Release	Description
2005	Contrail Networking Release 2005 supports Octavia as LBaaS.

RELATED DOCUMENTATION

Support for OpenStack LBaaS 381
Using Load Balancers in Contrail 368
Installing OpenStack Octavia LBaaS with RHOSP in Contrail Networking 334

Using Netronome SmartNIC vRouter with Contrail Networking and Juju Charms

IN THIS SECTION

- [Prepare to Install Contrail Networking by Using Juju Charms | 465](#)
- [Deploy Contrail Charms in a Bundle | 466](#)

NOTE: The Netronome SmartNIC vRouter technology covered in this document is available for evaluation purposes only. It is not intended for deployment in production networks.

You can deploy Contrail Networking by using Juju charms. Juju helps you deploy, configure, and efficiently manage applications on private clouds and public clouds. Juju accesses the cloud with the help of a Juju controller. A charm is a module containing a collection of scripts and metadata and is used with Juju to deploy Contrail.

Starting in Contrail Networking Release 2011, Contrail Networking supports Netronome Agilio CX SmartNICs for Contrail Networking deployment with Juju charms. This feature enables service providers to improve the forwarding performance which includes packets per second (PPS) of vRouter. This optimizes server CPU usage and you can deploy more Virtual Network Functions (VNFs) per server.

Before you begin:

- Equip compute nodes with Netronome Agilio CX SmartNIC. For details, see [Agilio CX SmartNICs documentation](#).
- Retrieve Agilio charm.

Register on Netronome support site at <https://help.netronome.com> and provide Docker Hub credentials.

Netronome will provide the Agilio charm for SmartNIC vRouter deployment on compute nodes. Add the charm version as *charm* variable in the [Bundle yaml file on page ?](#). Also, Netronome will authorize Docker Hub registry access.

- Note the *Container Tags* for your Contrail image to customize the *image-tag* variable in the [Bundle yaml file on page ?](#). See [README Access to Contrail Registry 20XX](#).
- Note the following version tags:
agilio-image-tag: 2.48-ubuntu-queens

maas version: 2.6.2

Linux kernel: bionic (ga-18.04)

Contrail supports the following charms:

- contrail-agent
- contrail-analytics
- contrail-analyticsdb
- contrail-controller
- contrail-keystone-auth
- contrail-openstack

The following topics describe how to use Netronome SmartNIC vRouter with Contrail Networking and Juju charms.

Prepare to Install Contrail Networking by Using Juju Charms

Follow these steps to prepare for deployment:

1. Install Juju.

```
sudo apt-get update
sudo apt-get upgrade
apt install snapd -y
snap install juju --classic
```

2. Configure Juju.

You can add a cloud to Juju, and manage clouds already added to Juju. Juju recognizes a wide range of cloud types for adding a cloud.

This is an example for adding a cloud by using interactive command.

Example: Adding an MAAS cloud to Juju

```
juju add-cloud
```

```
Cloud Types
  maas
  manual
  openstack
  oracle
  vsphere
```

```
Select cloud type: maas
```

```
Enter a name for your maas cloud: maas-cloud
```

```
Enter the API endpoint url: http://<ip-address>:<node>/MAAS
```

```
Cloud "maas-cloud" successfully added
```

```
You may bootstrap with 'juju bootstrap maas-cloud'
```

NOTE: Juju 2.x is compatible with MAAS series 1.x and 2.x.

3. Create a Juju controller.

```
juju bootstrap --bootstrap-series=xenial <cloud name> <controller name>
```

NOTE: A Juju controller manages and keeps track of applications in the Juju cloud environment.

Deploy Contrail Charms in a Bundle

Follow these steps to deploy Contrail charms in a bundle.

1. Deploy Contrail charms.

To deploy Contrail charms in a bundle, use the **juju deploy <bundle_yaml_file>** command.

The following example shows you how to use **bundle_yaml_file** to deploy Contrail Networking with Netronome SmartNIC vRouter on MAAS based deployment.

```
series: bionic
variables:
  openstack-origin:      &openstack-origin      distro
  #vhost-gateway:        &vhost-gateway          "192.x.40.254"
  data-network:          &data-network            "192.x.40.0/24"

  control-network:       &control-network        "192.x.30.0/24"
```

```

virtioforwarder-coremask:    &virtioforwarder-coremask    "1,2"
agilio-registry:            &agilio-registry            "netronomesystems"

    agilio-image-tag:        &agilio-image-tag
"2.48-ubuntu-queens"
    agilio-user:            &agilio-user
"<agilio-username>"
    agilio-password:        &agilio-password
"<agilio-password>"
    agilio-insecure:        &agilio-insecure            false
    agilio-phy:            &agilio-phy            "nfp_p0"
    docker-registry:        &docker-registry
"<registry-directory>"
    #docker-user:            &docker-user
"<docker_username>"
    #docker-password:        &docker-password
"<docker_password>"
    image-tag:            &image-tag            "2011.61"
    docker-registry-insecure: &docker-registry-insecure    "true"
    dockerhub-registry:    &dockerhub-registry
"https://index.docker.io/v1/"
machines:
    "1":
        constraints: tags=controller
        series: bionic
    "2":
        constraints: tags=compute
        series: bionic
    "3":
        constraints: tags=neutron
        series: bionic
services:
    ubuntu:
        charm: cs:ubuntu
        num_units: 1
        to: [ "1" ]
    ntp:
        charm: cs:ntp
        num_units: 0
        options:
            #source: ntp.ubuntu.com
            source: 10.204.217.158
    mysql:
        charm: cs:percona-cluster

```

```

num_units: 1
options:
  dataset-size: 15%
  max-connections: 10000
  root-password: <password>
  sst-password: <password>
  min-cluster-size: 1
to: [ "lxd:1" ]
rabbitmq-server:
  charm: cs:rabbitmq-server
  num_units: 1
  options:
    min-cluster-size: 1
  to: [ "lxd:1" ]
heat:
  charm: cs:heat
  num_units: 1
  expose: true
  options:
    debug: true
    openstack-origin: *openstack-origin
  to: [ "lxd:1" ]
keystone:
  charm: cs:keystone
  expose: true
  num_units: 1
  options:
    admin-password: <password>
    admin-role: admin
    openstack-origin: *openstack-origin
    preferred-api-version: 3
nova-cloud-controller:
  charm: cs:nova-cloud-controller
  num_units: 1
  expose: true
  options:
    network-manager: Neutron
    openstack-origin: *openstack-origin
  to: [ "lxd:1" ]
neutron-api:
  charm: cs:neutron-api
  expose: true
  num_units: 1
  series: bionic

```

```

    options:
      manage-neutron-plugin-legacy-mode: false
      openstack-origin: *openstack-origin
    to: [ "3" ]
glance:
  charm: cs:glance
  expose: true
  num_units: 1
  options:
    openstack-origin: *openstack-origin
    to: [ "lxd:1" ]
openstack-dashboard:
  charm: cs:openstack-dashboard
  expose: true
  num_units: 1
  options:
    openstack-origin: *openstack-origin
    to: [ "lxd:1" ]
nova-compute:
  charm: cs:nova-compute
  num_units: 0
  expose: true
  options:
    openstack-origin: *openstack-origin
nova-compute-dpdk:
  charm: cs:nova-compute
  num_units: 0
  expose: true
  options:
    openstack-origin: *openstack-origin
nova-compute-accel:
  charm: cs:nova-compute
  num_units: 2
  expose: true
  options:
    openstack-origin: *openstack-origin
    to: [ "2" ]
contrail-openstack:
  charm: ./tf-charms/contrail-openstack
  series: bionic
  expose: true
  num_units: 0
  options:
    docker-registry: *docker-registry

```

```

    #docker-user: *docker-user
    #docker-password: *docker-password
    image-tag: *image-tag
    docker-registry-insecure: *docker-registry-insecure
contrail-agent:
  charm: ./tf-charms/contrail-agent
  num_units: 0
  series: bionic
  expose: true
  options:
    log-level: "SYS_DEBUG"
    docker-registry: *docker-registry
    #docker-user: *docker-user
    #docker-password: *docker-password
    image-tag: *image-tag
    docker-registry-insecure: *docker-registry-insecure
    #vhost-gateway: *vhost-gateway
    physical-interface: *agilio-phy
contrail-agent-dpdk:
  charm: ./tf-charms/contrail-agent
  num_units: 0
  series: bionic
  expose: true
  options:
    log-level: "SYS_DEBUG"
    docker-registry: *docker-registry
    #docker-user: *docker-user
    #docker-password: *docker-password
    image-tag: *image-tag
    docker-registry-insecure: *docker-registry-insecure
    dpdk: true
    dpdk-main-mempool-size: "65536"
    dpdk-pmd-txd-size: "2048"
    dpdk-pmd-rxd-size: "2048"
    dpdk-driver: ""
    dpdk-coremask: "1-4"
    #vhost-gateway: *vhost-gateway
    physical-interface: "nfp_p0"
contrail-analytics:
  charm: ./tf-charms/contrail-analytics
  num_units: 1
  series: bionic
  expose: true
  options:

```

```

    log-level: "SYS_DEBUG"
    docker-registry: *docker-registry
    #docker-user: *docker-user
    #docker-password: *docker-password
    image-tag: *image-tag
    control-network: *control-network
    docker-registry-insecure: *docker-registry-insecure
  to: [ "1" ]
contrail-analyticsdb:
  charm: ./tf-charms/contrail-analyticsdb
  num_units: 1
  series: bionic
  expose: true
  options:
    log-level: "SYS_DEBUG"
    cassandra-minimum-diskgb: "4"
    cassandra-jvm-extra-opts: "-Xms8g -Xmx8g"
    docker-registry: *docker-registry
    #docker-user: *docker-user
    #docker-password: *docker-password
    image-tag: *image-tag
    control-network: *control-network
    docker-registry-insecure: *docker-registry-insecure
  to: [ "1" ]
contrail-controller:
  charm: ./tf-charms/contrail-controller
  series: bionic
  expose: true
  num_units: 1
  options:
    log-level: "SYS_DEBUG"
    cassandra-minimum-diskgb: "4"
    cassandra-jvm-extra-opts: "-Xms8g -Xmx8g"
    docker-registry: *docker-registry
    #docker-user: *docker-user
    #docker-password: *docker-password
    image-tag: *image-tag
    docker-registry-insecure: *docker-registry-insecure
    control-network: *control-network
    data-network: *data-network
    auth-mode: no-auth
  to: [ "1" ]
contrail-keystone-auth:
  charm: ./tf-charms/contrail-keystone-auth

```

```

    series: bionic
    expose: true
    num_units: 1
    to: [ "lxd:1" ]
agilio-vrouter5:
    charm: ./charm-agilio-vrt-5-37
    expose: true
    options:
        virtioforwarder-coremask: *virtioforwarder-coremask
        agilio-registry: *agilio-registry
        agilio-insecure: *agilio-insecure
        agilio-image-tag: *agilio-image-tag
        agilio-user: *agilio-user
        agilio-password: *agilio-password
relations:
    - [ "ubuntu", "ntp" ]
    - [ "neutron-api", "ntp" ]
    - [ "keystone", "mysql" ]
    - [ "glance", "mysql" ]
    - [ "glance", "keystone" ]
    - [ "nova-cloud-controller:shared-db", "mysql:shared-db" ]
    - [ "nova-cloud-controller:amqp", "rabbitmq-server:amqp" ]
    - [ "nova-cloud-controller", "keystone" ]
    - [ "nova-cloud-controller", "glance" ]
    - [ "neutron-api", "mysql" ]
    - [ "neutron-api", "rabbitmq-server" ]
    - [ "neutron-api", "nova-cloud-controller" ]
    - [ "neutron-api", "keystone" ]
    - [ "nova-compute:amqp", "rabbitmq-server:amqp" ]
    - [ "nova-compute", "glance" ]
    - [ "nova-compute", "nova-cloud-controller" ]
    - [ "nova-compute", "ntp" ]
    - [ "openstack-dashboard:identity-service", "keystone" ]
    - [ "contrail-keystone-auth", "keystone" ]
    - [ "contrail-controller", "contrail-keystone-auth" ]
    - [ "contrail-analytics", "contrail-analyticsdb" ]
    - [ "contrail-controller", "contrail-analytics" ]
    - [ "contrail-controller", "contrail-analyticsdb" ]
    - [ "contrail-openstack", "nova-compute" ]
    - [ "contrail-openstack", "neutron-api" ]
    - [ "contrail-openstack", "contrail-controller" ]
    - [ "contrail-agent:juju-info", "nova-compute:juju-info" ]
    - [ "contrail-agent", "contrail-controller" ]
    - [ "contrail-agent-dpdk:juju-info", "nova-compute-dpdk:juju-info" ]

```



```

- [ "contrail-agent-dpdk", "contrail-controller" ]
- [ "nova-compute-dpdk:amqp", "rabbitmq-server:amqp" ]
- [ "nova-compute-dpdk", "glance" ]
- [ "nova-compute-dpdk", "nova-cloud-controller" ]
- [ "nova-compute-dpdk", "ntp" ]
- [ "contrail-openstack", "nova-compute-dpdk" ]
- [ "contrail-agent:juju-info", "nova-compute-accel:juju-info" ]
- [ "nova-compute-accel:amqp", "rabbitmq-server:amqp" ]
- [ "nova-compute-accel", "glance" ]
- [ "nova-compute-accel", "nova-cloud-controller" ]
- [ "nova-compute-accel", "ntp" ]
- [ "contrail-openstack", "nova-compute-accel" ]
- [ "agilio-vrouter5:juju-info", "nova-compute-accel:juju-info" ]
- [ "heat:shared-db" , "mysql:shared-db" ]
- [ "heat:amqp" , "rabbitmq-server:amqp" ]
- [ "heat:identity-service" , "keystone:identity-service" ]
- [ "contrail-openstack:heat-plugin" , "heat:heat-plugin-subordinate" ]

```

You can create or modify the Contrail charm deployment bundle YAML file to:

- Point to machines or instances where the Contrail charms must be deployed.
- Include the options you need.

Each Contrail charm has a specific set of options. The options you choose depend on the charms you select. For more information on the options that are available, see [“Options for Juju Charms” on page 429](#).

2. (Optional) Check the status of deployment.

You can check the status of the deployment by using the **juju status** command.

3. Enable configuration statements.

Based on your deployment requirements, you can enable the following configuration statements:

- **contrail-agent**

For more information, see <https://jaas.ai/u/juniper-os-software/contrail-agent/>.

- **contrail-analytics**

For more information, see <https://jaas.ai/u/juniper-os-software/contrail-analytics>.

- **contrail-analyticsdb**

For more information, see <https://jaas.ai/u/juniper-os-software/contrail-analyticsdb>.

- **contrail-controller**

For more information, see <https://jaas.ai/u/juniper-os-software/contrail-controller>.

- **contrail-keystone-auth**
For more information, see <https://jaas.ai/u/juniper-os-software/contrail-keystone-auth>.
- **contrail-openstack**
For more information see, <https://jaas.ai/u/juniper-os-software/contrail-openstack>.

Release History Table

Release	Description
2011	Starting in Contrail Networking Release 2011, Contrail Networking supports Netronome Agilio CX SmartNICs for Contrail Networking deployment with Juju charms.

RELATED DOCUMENTATION

Installing Contrail with OpenStack by Using Juju Charms 412
MAAS Installation Documents

Using Contrail and Contrail Insights with Kolla/Ocata OpenStack

IN THIS CHAPTER

- [Contrail, Contrail Insights, and OpenStack Kolla/Ocata Deployment Requirements | 475](#)
- [Preparing for the Installation | 476](#)
- [Run the Playbooks | 479](#)
- [Accessing Contrail in Contrail Insights Management Infrastructure in UI | 480](#)
- [Notes and Caveats | 481](#)
- [Example Instances.yml for Contrail and Contrail Insights OpenStack Deployment | 481](#)
- [Contrail Insights Installation and Configuration for OpenStack | 485](#)
- [Contrail Insights Installation for OpenStack in HA | 499](#)

Contrail, Contrail Insights, and OpenStack Kolla/Ocata Deployment Requirements

Starting with Contrail Release 5.0.1, the combined installation of Contrail and Contrail Insights allows Contrail monitoring by Contrail Insights. The following topics are referenced for the deployment.

- *Installing Contrail with OpenStack and Kolla Ansible*
- [Contrail Insights Installation and Configuration for OpenStack on page 485](#)
- [Contrail Insights Installation for OpenStack in HA on page 499](#)

The following software and hardware requirements apply to the combined Contrail, Contrail Insights, and Kolla/Ocata deployment.

Software Requirements

- Contrail Release 5.0.x Targets: Centos 7.5 with kernel 3.10.0-862.3.2.el7.x86_64.
- Contrail Insights Targets: Refer to “Software Requirements” in *Contrail Insights General Requirements*.

- Targets running both Contrail and Contrail Insights: CentOS 7.5 Ansible 2.4.2 for the installer.
- Contrail Insights 2.18.x and later.

Hardware Requirements

- It is strongly recommended that the Contrail Insights controller and Contrail services be installed on separate targets.
- See [“Installing a Contrail Cluster using Contrail Command and instances.yml” on page 74](#) and [“Contrail Insights Installation and Configuration for OpenStack” on page 485](#) for specifics about requirements for installation.

Preparing for the Installation

In Contrail Release 5.1, nodes on which Contrail, Contrail Insights (formerly AppFormix), or both are installed are referred to as *targets*. The host from which Ansible is run is referred to as the *base host*. A *base host* can also be a *target*, meaning you can install either Contrail, Contrail Insights, or both on a *base host*.

Preparing the Targets

Workflow for preparing the targets consists of the following steps:

1. Image all the Contrail targets with CentOS 7.5 kernel 3.10.0-862.3.2.el7.x86_64.
2. Install the necessary platform software on the targets on which Contrail Insights Controller or Contrail Insights Agent is going to be installed. See the instructions in [“Contrail Insights Installation and Configuration for OpenStack” on page 485](#).

Preparing the Base Host using Ansible Installer

Workflow for preparing the base host consists of the following steps:

1. Install Ansible 2.4.2 on the base host. See “Set Up the Bare Host” in *Installing Contrail with OpenStack and Kolla Ansible*.
2. Set-up the base host. See “Set Up the Base Host” in *Installing Contrail with OpenStack and Kolla Ansible*. This section includes information about creating the Ansible **instances.yml** file.

3. On the base host, create a single Ansible **instances.yaml** file that lists inventory for both Contrail and Contrail Insights deployments. An example of the single **instances.yaml** file is provided later in this section.
 - The Contrail inventory section of the **instances.yaml** file is configured according to guidelines in the section “Set Up the Base Host” in *Installing Contrail with OpenStack and Kolla Ansible* .
 - The Contrail Insights inventory section of the **instances.yaml** file is configured according to guidelines in [“Contrail Insights Installation and Configuration for OpenStack” on page 485](#).

TCP/IP Port Conflicts Between Contrail and Contrail Insights

It is strongly recommended that Contrail Insights Controller and Contrail services be installed on separate target nodes. However, if Contrail Insights Controller and Contrail services are installed on the same target, the following configuration is required to resolve port conflicts.

The following Contrail Insights ports must be reconfigured in the Contrail Insights **group-vars** section of the **instances.yaml** file.

- **appformix_datamanager_port_http**
- **appformix_datamanager_port_https**
- **appformix_haproxy_datamanager_port_http**
- **appformix_haproxy_datamanager_port_https**
- **appformix_datamanager_port_http:8200**

Plugins to Enable for Contrail and Contrail Insights Deployment

Enable the following plugins by including them in the Contrail Insights **group-vars** section of the **instances.yaml** file.

```
appformix_plugins: '{{ appformix_contrail_factory_plugins }}'
appformix_openstack_log_plugins: '{{ appformix_openstack_log_factory_plugins }}'
```

Configuring Contrail Monitoring in Contrail Insights

Connections to Contrail are configured by providing complete URLs by which to access the analytics and configuration API services.

- **contrail_cluster_name: Contrail_Clusterxxx**

A name by which the Contrail instance will be displayed in the Dashboard. If not specified, this variable has a default value of **default_contrail_cluster**.

- **contrail_analytics_url:** **http://analytics-api-node-ip-address:8081**

URL for the Contrail analytics API. The URL should only specify the protocol, address, and optionally port.

- **contrail_config_url:** **http://contrail-config-api-server-api-address:8082**

URL for the Contrail configuration API. The URL should only specify the protocol, address, and optionally port.

NOTE: The IP address specified for contrail monitoring corresponds to one of the IPs listed in the Contrail roles for *config* and *analytics*. Typically, the first active IP address is selected.

Compute Monitoring: Listing IP Addresses to Monitor

The IP addresses to monitor can be added in the **compute** section of Contrail Insights in the **instances.yaml** file. A list of IP addresses with a *vrouter* role in the **instances.yaml** file.

Configuring Openstack_Controller Hosts for Contrail Insights

The Openstack_controller hosts section must be configured with at least one host. An example section is shown.

```
openstack_controller:
  hosts:
    <ip-address>:
      ansible_connection: ssh
      ansible_ssh_user: <root user>
      ansible_sudo_pass: <contrail password>
```

Other Contrail Insights group_vars That Must be Enabled in instances.yaml

The following **group_vars** must be enabled in **instances.yaml**:

- **openstack_platform_enabled:** true
- **appformix_remote_host_monitoring_enabled:** true

Contrail Insights License

(Required for Contrail Insights and Contrail Insights Flows installations in Release 2003 and earlier.) You must have an appropriate license that supports the combined deployment of Contrail with Contrail Insights for OpenStack. To obtain a license, send an email to APPFORMIX-KEY-REQUEST@juniper.net. Also, the following `group_vars` Contrail Insights in the `instances.yaml` file must point to this license.

- `appformix_license: /path/appformix-contrail-license-file.sig`

This is the path where the license is placed on the *bare host* so that the license can be deployed on the target.

RELATED DOCUMENTATION

Installing Contrail with OpenStack and Kolla Ansible

[Installing a Contrail Cluster using Contrail Command and instances.yaml](#) | 74

[Contrail Insights Installation and Configuration for OpenStack](#) | 485

[Example Instances.yaml for Contrail and Contrail Insights OpenStack Deployment](#) | 481

Run the Playbooks

Refer to section “Install Contrail and Kolla requirements” and section “Deploying contrail and Kolla containers” in *Installing Contrail with OpenStack and Kolla Ansible* and execute the **ansible-playbook**.

Following are examples listing the Contrail play-book invocation from the **contrail-ansible-deployer** directory:

- Configure Contrail OpenStack instances:

```
ansible-playbook -i inventory/ -e config_file=/path/instances.yaml -e
orchestrator=openstack playbooks/configure_instances.yaml (-vvv for debug)
```

- Install OpenStack:

```
ansible-playbook -i inventory/ -e config_file=/path/instances.yaml
playbooks/install_openstack.yaml
```

- Install Contrail:

```
ansible-playbook -i inventory/ -e config_file=/path/instances.yaml -e
orchestrator=openstack playbooks/install_contrail.yml
```

Source the `/etc/kolla/kolla-toolbox/admin-openrc.sh` file from the OpenStack controller node (`/etc/kolla/kolla-toolbox/ admin-openrc.sh`) to the Contrail Insights Controller to authenticate the OpenStack adapter to access admin privileges over controller services. If the OpenStack control node is different from the base host, either Secure Copy Protocol (SCP) the file over and source it (for example, **execute** `source /path/admin-openrc.sh`) or manually export the environment enumerated in `/etc/kolla/kolla-toolbox/ admin-openrc.sh` by invoking **export** `OS_USERNAME=admin` etc. and the remainder as listed in `admin-openrc.sh`

Also at this point, obtain a list of IP addresses to include in the **compute** section of Contrail Insights in the `instances.yaml` file. Refer to **Compute monitoring: Listing IP addresses to monitor** in the **compute** section of Contrail Insights in the `instances.yaml` file.

Refer to *Installing AppFormix for OpenStack* and validate target configuration requirements and inventory parameters for Contrail Insights Controller and Agent. In place of `-i inventory/use -i /absolute-file-path/instances.yaml`.

Following is an example listing the Contrail Insights playbook invocation from the **appformix-2.18.x** directory where `appformix_openstack.yml` is located:

- Install Contrail Insights:

```
ansible-playbook -i /path/instances.yaml appformix_openstack.yml (-vvv for debug)
```

RELATED DOCUMENTATION

Installing Contrail with OpenStack and Kolla Ansible

[Installing a Contrail Cluster using Contrail Command and instances.yaml](#) | 74

[Contrail Insights Installation and Configuration for OpenStack](#) | 485

Accessing Contrail in Contrail Insights Management Infrastructure in UI

Contrail Insights service monitoring Dashboard for a Contrail cluster displays the overall state of the cluster and its components. For more information, see “Dashboard” in “Contrail Monitoring” in the [Contrail Insights User Guide](#).

Open the Dashboard in a Web browser and log in.

`http://<controller-IP-address>:9000`

RELATED DOCUMENTATION

| [Contrail Insights User Guide](#)

Notes and Caveats

- Versions of Contrail Insights-2.17 and earlier are not supported with Ansible-2.4.2. The combined Contrail and Contrail Insights installation is not validated on these earlier releases.
- The installation was validated with Contrail Insights-2.18 Agent.
- To view and monitor Contrail in the Contrail Insights Management Infrastructure dashboard, the license used in the deployment must include support for Contrail.
- Verify the datamanager port (re)definitions in the inventory file.
- For Contrail Insights OpenStack HA installation steps, see [“Contrail Insights Installation for OpenStack in HA” on page 499](#).

RELATED DOCUMENTATION

| [Contrail Insights Installation for OpenStack in HA | 499](#)

Example Instances.yml for Contrail and Contrail Insights OpenStack Deployment

See *Installing Contrail with OpenStack and Kolla Ansible* and [“Contrail Insights Installation and Configuration for OpenStack” on page 485](#) for specific inventory file details:

The following items are part of the **all** section in the **instances.yml** file for Contrail Insights:

```
all:
  children:
    openstack_controller:
      hosts:
        <ip-address>:
```

```

ansible_connection: ssh
ansible_ssh_user: <ssh-user>
ansible_sudo_pass: <sudo-password>

```

The following items are part of the **vars** section in the **instances.yaml** file for Contrail Insights:

```

openstack_platform_enabled: true
##License must support Contrail and Openstack
appformix_license: /path/license-file.sig
contrail_cluster_name: 'Contrail_Cluster'
contrail_analytics_url: 'http://<contrail-analytics-api-server-ip-address>:8081'
contrail_config_url: 'http://<contrail-config-api-server-ip-address>:8082'
# Defaults from roles/appformix_defaults/defaults/main.yml are overwritten below
appformix_datamanager_port_http: "{{ (appformix_scale_setup_flag|bool) |
ternary(28200, 8200) }}"
appformix_datamanager_port_https: "{{ (appformix_scale_setup_flag|bool) |
ternary(28201, 8201) }}"
appformix_haproxy_datamanager_port_http: 8200
appformix_haproxy_datamanager_port_https: 8201
appformix_plugins: '{{ appformix_contrail_factory_plugins }} + {{
appformix_network_device_factory_plugins }}'

```

Following is an example listing of the **instances.yaml**:

There is one **instances.yaml** file for the Contrail and Contrail Insights combined installation.

```

#Contrail inventory section
provider_config:
  bms:
    ssh_pwd: <ssh-password>
    ssh_user: <ssh-user>
    ntpserver: <ntp-server-ip-address>
    domainsuffix: local
instances:
  bms1:
    provider: bms
    ip: <ip-address>
    roles:
      config_database:
        config:
      control:
      analytics_database:
      analytics:
      webui:

```

```

    vrouter:
    openstack:
    openstack_compute:
global_configuration:
    CONTAINER_REGISTRY: <ci-repository-URL>:5000
    REGISTRY_PRIVATE_INSECURE: True
contrail_configuration:
    #UPGRADE_KERNEL: true
    CONTRAIL_VERSION: <contrail-version>
    #CONTRAIL_VERSION: latest
    CLOUD_ORCHESTRATOR: openstack
    VROUTER_GATEWAY: <gateway-ip-address>
    RABBITMQ_NODE_PORT: 5673
    PHYSICAL_INTERFACE: <interface-name>
    AUTH_MODE: keystone
    KEYSTONE_AUTH_HOST: <keystone-ip-address>
    KEYSTONE_AUTH_URL_VERSION: /v3
    CONFIG_NODEMGR_DEFAULTS__minimum_diskGB: 2
    DATABASE_NODEMGR_DEFAULTS__minimum_diskGB: 2
kolla_config:
    kolla_globals:
        network_interface: <interface-name>
        kolla_internal_vip_address: <ip-address>
        contrail_api_interface_address: <ip-address>
        enable_haproxy: no
        enable_swift: no
    kolla_passwords:
        keystone_admin_password: <password>

# Contrail Insights inventory section
all:
    children:
        appformix_controller:
            hosts:
                <ip-address>:
                    ansible_connection: ssh
                    ansible_ssh_user: <ssh-user>
                    ansible_sudo_pass: <sudo-password>
        openstack_controller:
            hosts:
                <ip-address>:
                    ansible_connection: ssh
                    ansible_ssh_user: <ssh-user>
                    ansible_sudo_pass: <sudo-password>

```

```

compute:
  hosts:
    #List IP addresses of Contrail roles to be monitored here
    <<IP-addresses>>:
      ansible_connection: ssh
      ansible_ssh_user: <ssh-user>
      ansible_sudo_pass: <sudo-password>
bare_host:
  hosts:
    <ip-address>:
      ansible_connection: ssh
      ansible_ssh_user: <ssh-user>
      ansible_sudo_pass: <sudo-password>
    #If host is local
    <ip-address>:
      ansible_connection: local
vars:
  appformix_docker_images:
    - /opt/software/appformix/contrail-insights-platform-images-<version>.tar.gz
    - /opt/software/appformix/contrail-insights-dependencies-images-<version>.tar.gz
    -
    /opt/software/appformix/contrail-insights-network_device-images-<version>.tar.gz
    - /opt/software/appformix/contrail-insights-openstack-images-<version>.tar.gz

  openstack_platform_enabled: true
  # appformix_license:
  /opt/software/openstack_appformix/<appformix-contrail-license-file>.sig
  appformix_license: /opt/software/configs/contrail.sig
  appformix_docker_registry: registry.appformix.com/
  appformix_version: <version>      #Must be 2.18.x or above
  appformix_plugins: '{{ appformix_contrail_factory_plugins }} + {{
appformix_network_device_factory_plugins }} + {{ appformix_openstack_factory_plugins
}}'
  appformix_kvm_instance_discovery: true
  # For enabling pre-requisites for package installation
  appformix_network_device_monitoring_enabled: true
  # For running the appformix-network-device-adapter
  network_device_discovery_enabled: true
  appformix_remote_host_monitoring_enabled: true
  appformix_jti_network_device_monitoring_enabled: true
  contrail_cluster_name: 'Contrail_Cluster'
  contrail_analytics_url: 'http://<contrail-analytics-api-server-IP-address>:8081'

```

```

    contrail_config_url: 'http://<contrail-config-api-server-IP-address>:8082'
    # Defaults overwritten below were defined in
roles/appformix_defaults/defaults/main.yml
    appformix_datamanager_port_http: "{{ (appformix_scale_setup_flag|bool) |
ternary(28200, 8200) }}"
    appformix_datamanager_port_https: "{{ (appformix_scale_setup_flag|bool) |
ternary(28201, 8201) }}"
    appformix_haproxy_datamanager_port_http: 8200
    appformix_haproxy_datamanager_port_https: 8201

```

NOTE: Replace `<contrail_version>` with the correct **contrail_container_tag** value for your Contrail release. The respective **contrail_container_tag** values are listed in [README Access to Contrail Networking Registry 20xx](#).

RELATED DOCUMENTATION

Installing Contrail with OpenStack and Kolla Ansible

[Installing a Contrail Cluster using Contrail Command and instances.yml](#) | 74

[Contrail Insights Installation and Configuration for OpenStack](#) | 485

Contrail Insights Installation and Configuration for OpenStack

IN THIS SECTION


- [Getting Started with OpenStack](#) | 486
- [Requirements](#) | 487
- [Workflow in Three Steps](#) | 487
- [Configure OpenStack](#) | 488
- [Create Ansible Inventory](#) | 488
- [Install Contrail Insights](#) | 490
- [Additional Ansible Configuration](#) | 493
- [Upgrade Notices](#) | 496
- [Remove a Node from Contrail Insights](#) | 498

- Upgrade Contrail Insights for an OpenStack Cluster | 498
- Uninstall Contrail Insights from an OpenStack Cluster | 498

Getting Started with OpenStack

Contrail Insights provides resource control and visibility for hosts and virtual machines in an [OpenStack](#) cluster. This topic explains how to install Contrail Insights for an OpenStack cluster. Contrail Insights Agent runs on a host to monitor resource consumption of the host itself and the virtual machines executing on that host. See the *Contrail Insights General Requirements* and *Contrail Insights Agent Requirements* before reading this section. [Figure 75 on page 486](#) shows the Contrail Insights architecture with OpenStack.

Figure 75: Contrail Insights Architecture with OpenStack



- Contrail Insights Agent monitors resource usage on the compute nodes.
- Contrail Insights Platform offers REST APIs to configure the system.
- Contrail Insights DataManager stores data from multiple Agents.
- Contrail Insights Dashboard provides a Web-based user interface.

- An adapter discovers platform-specific resources and configures the Contrail Insights Platform.
- Adapters exist for OpenStack and Kubernetes.

Requirements

The following are the requirements for installing Contrail Insights for OpenStack.

- Supported OpenStack versions: Icehouse, Juno, Kilo, Liberty, Mitaka, Newton, Ocata.
- See *Contrail Insights General Requirements* for software and hardware requirements.
- An administrator account for OpenStack.
- API access to OpenStack services: Cinder, Glance, Heat, Keystone, Neutron, Nova, and Swift. Contrail Insights reads information from these services. The administrator account must provide sufficient permission for read-only API calls. Further, Contrail Insights Platform must be able to open connections to the host and port on which these services listen. Contrail Insights can be configured to use the **admin**, **internal**, or **public** service endpoints. See **OS_ENDPOINT_TYPE** in OpenStack environment variables in the section Installing Contrail Insights, [10](#).

-

NOTE: Upgrade notice: Starting with Contrail Insights 3.2.6, the requirement for a license file is removed. If you are installing a version earlier than 3.2.6, a license is required prior to installation.

You can obtain a license key from APPFORMIX-KEY-REQUEST@juniper.net. Provide the following information in your request:

```
Group name:
Target customers or use:
Cluster type: OpenStack
Number of hosts:
Number of instances:
```

Workflow in Three Steps

Installation consists of the following steps:

1. Configure OpenStack.
2. Create Ansible inventory.
3. Install Contrail Insights.

Configure OpenStack

To create an administrator account for Contrail Insights, perform the following steps in the OpenStack Horizon dashboard:

1. Create a user account and name it **appformix**.
2. Select a new project for the user account.
3. Select role as **admin**.

Create Ansible Inventory

Ansible is used to deploy the software to the compute node(s) and the Platform Host. An Ansible inventory file describes groups of hosts in your cluster. Define the inventory in a separate location than the release files, so that the inventory may be reused for an upgrade.

Contrail Insights requires two groups *compute* and *appformix_controller*. Each group lists the hosts in that group. Only the agent is installed on the *compute* hosts. The agent and the Contrail Insights Platform services are installed on the *appformix_controller* host.

Optionally, an *openstack_controller* group can be defined. The agent is installed on hosts in this group to monitor the hosts that execute OpenStack controller services. (New in v2.3.0)

Create a directory **inventory** (or name of your choice) that contains a **hosts** file and a **group_vars/all** file. For example:

```
inventory/
  hosts          # inventory file
  group_vars/
    all          # configuration variables
```

The **inventory/hosts** file contains the list of hosts in each group. For example:

```
[appformix_controller]
appformix01

[compute]
compute01
compute02
compute03
```



```
[openstack_controller]
openstack_infra01
openstack_infra02
```

See [Ansible inventory documentation](#).

Ansible Configuration Variables

The Contrail Insights software includes a number of Ansible roles to perform the configuration of Contrail Insights settings. Define the values of variables in the **inventory/group_vars/all** file, in order to be able to use the settings and inventory for future upgrades.

In the **inventory/group_vars/all** file, configure the following variables for installation of Contrail Insights for OpenStack.

```
appformix_docker_images:
  - /path/to/contrail-insights-platform-images-<version>.tar.gz
  - /path/to/contrail-insights-dependencies-images<version>.tar.gz
  - /path/to/contrail-insights-openstack-images-<version>.tar.gz
```

Refer to *Platform Dependencies* for steps to install dependencies on a Platform Host that cannot fetch files from the Internet.

Configure an HTTP Proxy for Fetching Prerequisites

The Ansible playbook will fetch files from the Internet to install prerequisites on the Platform Host. If the Platform Host requires an HTTP proxy to access the Internet, configure the following variables in **inventory/group_vars/all**:

```
http_proxy_url: 'http://proxy.example.com:1234'
prerequisites_env:

    http_proxy: '{{ http_proxy_url }}'
    https_proxy: '{{ http_proxy_url }}'
```

The **prerequisites_env** is a dictionary that defines environment variables that will be used when invoking commands to install prerequisites. In the above example, the same proxy URL (<http://proxy.example.com:1234>) is used for both the **http_proxy** and **https_proxy** environment variables because the single proxy can be used to access HTTP and HTTPS URLs. As a convenience, the proxy URL is defined once in the **http_proxy_url** variable. Adjust **prerequisites_env** as necessary for the proxy requirements of your network.

Install Contrail Insights

To install Contrail Insights:

1. Install Ansible on the Contrail Insights Platform node. Ansible will install docker and docker-py on the platform.

```
#Ubuntu
apt-get install python-pip python-dev #Installs Pip
pip install ansible==2.3.0 #Installs Ansible
2.3
sudo apt-get install build-essential libssl-dev libffi-dev #Dependencies
pip install markupsafe httplib2 requests #Dependencies

#RHEL/CentOS
yum install epel-release #Enable EPEL
repository
In case the above command does not work, manually download and install the
epel-release
package with one of the below commands, depending on your system's version.
yum install
https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
yum install
https://dl.fedoraproject.org/pub/epel/epel-release-latest-6.noarch.rpm

yum groupinstall 'Development Tools' #Install
development tools
yum install openssl-devel libffi libffi-devel #Dependencies
yum install python-pip python-devel #Install Pip
pip install ansible==2.3.0 #Install Ansible
2.3
pip install markupsafe httplib2 requests #Dependencies
```

NOTE: For RHEL, the following IPtables rule is needed to access port 9000.

```
iptables -t filter -A IN_public_allow -p tcp --dport 9000 -j ACCEPT
```

2. Install **libvirt** on the compute nodes using the following command:

```
yum -y install libvirt
```

This installs the KVM (Kernel-based Virtual Machine) and associated packages for collecting data from virtual machines running on the compute nodes.

3. On the compute nodes where Contrail Insights Agent runs, verify that python **virtualenv** is installed.

```
#Ubuntu
apt-get install -y python-pip
pip install virtualenv

#RHEL/CentOS
yum install -y python-pip
pip install virtualenv
```

4. Enable passwordless login to facilitate Contrail Insights Platform node with Ansible to install agents on the nodes. Create a SSH public key on the node where Ansible playbooks are run, and then copy the key to the **appformix_controller** node.

```
ssh-keygen -t rsa    #Creates Keys

ssh-copy-id -i ~/.ssh/id_rsa.pub <target_host>    #Copies key from the node to
other hosts
```

5. Use the **Sample_Inventory** file as a template to create a host file.

```
# Example naming schemes are as below:
#   hostname ansible_ssh_user='username' ansible_sudo_pass='password'

# List all Compute Nodes
[compute]
172.16.70.5
172.16.70.17

# appformix_controller host
#
# Host variables can be defined to control Contrail Insights configuration
parameters
# for particular host.  For example, to specify the directory in which MongoDB
# data is stored on hostname1 (the default is /opt/appformix/mongo/data):
#
#   hostname1 appformix_mongo_data_dir=/var/lib/appformix/mongo
#
# For variables with same value for all appformix_controller hosts, set group
```

```
# variables below.
#
[appformix_controller]
172.16.70.119

[openstack_controller]
172.16.70.120
```

6. Verify that all the hosts listed in the inventory file are reachable from the Contrail Insights Platform.

```
export ANSIBLE_HOST_KEY_CHECKING=False # Eliminates interactive experience
prompting for Known_Hosts

ansible -i inventory -m ping all      # Pings all the hosts in the inventory
file
```

7. At the top-level of the distribution, create a directory named **group_vars**.

```
mkdir group_vars
```

8. Download the Contrail Insights installation packages from [software downloads](#) to the Contrail Insights Platform node. Get the following files:

```
contrail-insights-<version>.tar.gz
contrail-insights-dependencies-images-<version>.tar.gz
contrail-insights-openstack-images-<version>.tar.gz
contrail-insights-platform-images-<version>.tar.gz
contrail-insights-network_device-images-<version>.tar.gz
```

If you are installing a version earlier than 3.2.6, copy the Contrail Insights license file to the Contrail Insights Platform node.

9. In **group_vars** directory, create a file named **all**. Add the following:

```
openstack_platform_enabled: true

appformix_manager_version: <version>
appformix_docker_images:
  - /path/to/contrail-insights-platform-images-<version>.tar.gz
```

- /path/to/contrail-insights-dependencies-images-<version>.tar.gz
- /path/to/contrail-insights-openstack-images-<version>.tar.gz

If you are installing a version earlier than 3.2.6, include the path to the Contrail Insights license file in **group_vars/all**:

```
appformix_license: path/to/<contrail-insights-license-file>.sig
```

10. Contrail Insights must be configured to communicate with the OpenStack cluster. The Ansible playbooks use OpenStack environment variables to configure Contrail Insights with details of the OpenStack environment.

OS_AUTH_URL	Keystone URL (e.g., https://host:5000/)
OS_ENDPOINT_TYPE	endpoint type to communicate with service (default: publicURL)
OS_USERNAME	admin account to be used by Contrail Insights
OS_PASSWORD	password for admin account
OS_PROJECT_NAME	admin project created for Contrail Insights account
OS_PROJECT_DOMAIN_NAME	domain for admin project
OS_USER_DOMAIN_NAME	domain for admin user
OS_DOMAIN_NAME	(optional) use domain-scoped token for admin account

11. Source the **openrc** file that contains the [10](#) environment variables and ensure the variables are in the environment of the shell from which the Ansible-playbooks are going to be executed. Then, install Contrail Insights by executing the **appformix_openstack.yml** playbook. Specify the path to the **inventory** directory that you created earlier.

12. Open the Contrail Insights Dashboard in a Web browser. For example:

```
http://<contrail-insights-platform-node-ip>:9000
```

Select **Skip Installation** because the initial configuration was performed by Ansible using the OpenStack environment variables in [10](#). Log in to Contrail Insights Dashboard using OpenStack Keystone credentials.

Additional Ansible Configuration

To set up additional Ansible configurations:

1. To install in a Keystone SSL-enabled cluster, include the following variables in the **group_vars/all** file:

```
appformix_keystone_ssl_ca: '/path/to/keystone_ca.crt'
```

Contrail Insights Ansible will distribute this Keystone CA to all of the Contrail Insights Platform nodes and ask Contrail Insights components to talk to Keystone using this CA file with SSL enabled.

NOTE: Deprecation Notice: The **appformix_mongo_cache_size_gb** parameter previously available starting in Contrail Insights 2.19.5 is now deprecated and no longer supported from Contrail Insights 3.2.0 and going forward. Starting with Contrail Insights version 3.2.0, Mongo will be configured to use a maximum of 40 percent of the available memory on the Contrail Insights Platform nodes.

2. To enable network device monitoring in the cluster, include the following in the **group_vars/all** file:

```
# For enabling pre-requisites for packdge installation.
appformix_install_snmp_dependencies: true
appformix_install_jti_dependencies: true
# For running the appformix-network-device-adapter
network_device_discovery_enabled: true
appformix_plugins: '{{ appformix_network_device_factory_plugins }}'
# After 3.1, SNMP Traps can be enabled also so appformix_plugins can be specified
as below:
# appformix_plugins: '{{ appformix_network_device_factory_plugins }} + {{
appformix_snmp_trap_factory_plugins }}'
```

3. To install Contrail Insights certified plug-ins on the cluster, include the following variables in the **group_vars/all** file:

```
appformix_plugins: <list of certified plugins to be installed>
appformix_openstack_log_plugins: <list of OpenStack log plugins to be installed>
```

For example:

```
appformix_plugins:
- { plugin_info: 'certified_plugins/cassandra_node_usage.json' }
- { plugin_info: 'certified_plugins/contrail_vrouter.json' }
- { plugin_info: 'certified_plugins/zookeeper_usage.json' }
- { plugin_info: 'certified_plugins/heavy_hitters.json' }

appformix_openstack_log_plugins:
- { plugin_info: 'certified_plugins/cinder_api_logparser.json',
```

```

    log_file_path: '/var/log/cinder/cinder-api.log' }
- { plugin_info: 'certified_plugins/glance_logparser.json',
    log_file_path: '/var/log/glance/glance-api.log' }
- { plugin_info: 'certified_plugins/keystone_logparser.json',
    log_file_path:
'/var/log/apache2/keystone_access.log,/var/log/httpd/keystone_wsgi_admin_access.log,/var/log/keystone/keystone.log'
}

```

For a list of all Contrail Insights certified plug-ins that can be installed, look for the entries starting with **plugin_info** in the file **roles/appformix_defaults/defaults/main.yml**.

The OpenStack log parser plug-ins parse the API log files of each OpenStack service to collect metrics about API calls and response status codes. To install these plug-ins, add them to the variable **appformix_openstack_log_plugins** in **group_vars/all**, as shown above. Each plug-in entry in this list requires a parameter called **log_file_path** to be specified. This parameter should be set to the complete path to the service's API log file on the OpenStack Controller node(s). Multiple comma-separated paths may be specified.

To identify the right log file to be specified in **log_file_path**, look for entries like the following, containing a client IP address, REST call type, and response status code:

```

2019-04-02 06:50:13.103 3465 INFO nova.osapi_compute.wsgi.server
[req-d07e953a-6921-4224-a056-afb6ff69adde 953ea56a96b944b3b170a299af9e87bd
10c9e8809feb4bd1b55955d9c2ed5aba - - -] 172.18.0.6 "GET
/v2/10c9e8809feb4bd1b55955d9c2ed5aba/os-hypervisors/detail HTTP/1.1" status: 200
len: 1427 time: 0.0208740
2019-04-02 06:50:13.183 3465 INFO nova.osapi_compute.wsgi.server
[req-34b2f686-9eb5-4112-b3fc-e0b37798a302 953ea56a96b944b3b170a299af9e87bd
10c9e8809feb4bd1b55955d9c2ed5aba - - -] 172.18.0.6 "GET
/v2/10c9e8809feb4bd1b55955d9c2ed5aba/servers/detail?all_tenants=1&status=SHELVED_OFFLOADED
HTTP/1.1" status: 200 len: 211 time: 0.0754580

```

Default locations for these files are listed in the variable **appformix_openstack_log_factory_plugins** in **roles/appformix_defaults/defaults/main.yml**.

4. In Contrail Insights version 2.19.8, a timeout value can be configured for connecting to OpenStack services. The default value of this timeout is **10** seconds and can be changed to a value between **5** and **20** seconds (both inclusive). To change the value, add the following variable to **group_vars/all**:

```
appformix_openstack_session_timeout: <number of seconds>
```

To modify the timeout value after the Contrail Insights Platform has been installed, add the variable to the **group_vars/all** file as described above and re-run the Contrail Insights installation playbook. Restart the **appformix-openstack-adapter** Docker container after the playbook has completed:

```
docker restart appformix-openstack-adapter
```

Upgrade Notices

NOTE: In Contrail Insights version 3.2.0, support for discovering OpenStack Octavia Load Balancer services is added. Contrail Insights only collects load balancer state information, such as **provisioning_status** and **operating_status**, as well as flavor information. To enable this service discovery, provide Octavia service's endpoint as variable **appformix_octavia_endpoint_url** in the **group_vars/all** file. For example:

```
appformix_octavia_endpoint_url: http://10.1.1.1:9876
```

Chargeback costs can also be configured for the Octavia Load Balancer services. See [Configure Load Balancer Costs](#).

Run Ansible with the created inventory file.

```
ansible-playbook -i inventory appformix-openstack.yml
```

NOTE: In Contrail Insights version 3.0, the variable **appformix_openstack_factory_plugins** is deprecated. All OpenStack log parser plug-ins have to be specified in the variable **appformix_openstack_log_plugins**. When upgrading from an older version to version 3.0, make sure to move all OpenStack log parser plug-ins defined in **appformix_openstack_factory_plugins** or **appformix_plugins** to **appformix_openstack_log_plugins**. Also, in Contrail Insights version 3.0, all entries in this list have to be specified with a **log_file_path** value, as described in [Step 3](#) above.

When upgrading from version 2.18.x to version 3.0, make the following changes in the **group_vars/all** file:

In version 2.18.x:

```
appformix_openstack_factory_plugins:
  - { plugin_info: 'certified_plugins/cinder_api_logparser.json', log_file_path:
    '/var/log/cinder/cinder-api.log' }
  - { plugin_info: 'certified_plugins/glance_logparser.json', log_file_path:
    '/var/log/glance/api.log' }
```



```

- { plugin_info: 'certified_plugins/heavy_hitters.json' }
- { plugin_info: 'certified_plugins/keystone_logparser.json', log_file_path:
'/var/log/keystone/keystone.log'}
- { plugin_info: 'certified_plugins/neutron_logparser.json', log_file_path:
'/var/log/neutron/server.log'}
- { plugin_info: 'certified_plugins/nova_logparser.json', log_file_path:
'/var/log/nova/nova-api.log'}

appformix_plugins: {{ appformix_openstack_factory_plugins }} + ...

```

In version 3.0.x:

```

appformix_plugins:
- { plugin_info: 'certified_plugins/heavy_hitters.json' }

appformix_openstack_log_plugins:
- { plugin_info: 'certified_plugins/cinder_api_logparser.json', log_file_path:
'/var/log/cinder/cinder-api.log'}
- { plugin_info: 'certified_plugins/glance_logparser.json', log_file_path:
'/var/log/glance/api.log'}
- { plugin_info: 'certified_plugins/keystone_logparser.json', log_file_path:
'/var/log/keystone/keystone.log'}
- { plugin_info: 'certified_plugins/neutron_logparser.json', log_file_path:
'/var/log/neutron/server.log'}
- { plugin_info: 'certified_plugins/nova_logparser.json', log_file_path:
'/var/log/nova/nova-api.log'}

```

Source the **openrc** file from the OpenStack controller node (**/etc/contrail/openstackrc**) to the Contrail Insights Platform node to authenticate the adapter to access administrator privileges over the controller services. The file should look like the following:

```

export OS_USERNAME=admin
export OS_PASSWORD=<admin-password>
export OS_AUTH_URL=http://172.16.80.2:5000/v2.0/
export OS_NO_CACHE=1
export OS_PROJECT_DOMAIN_NAME=Default
export OS_USER_DOMAIN_NAME=Default
export OS_PROJECT_NAME=admin
export OS_IDENTITY_API_VERSION=3
export OS_IMAGE_API_VERSION=2

```

Run Ansible with the created inventory file.

```
ansible-playbook -i inventory appformix_openstack.yml
```

Remove a Node from Contrail Insights

To remove a node from Contrail Insights:

1. Edit the inventory file and add **appformix_state=absent** to each node that you want to remove from Contrail Insights.

```
# Example naming schemes are as below:
#   hostname ansible_ssh_user='username' ansible_sudo_pass='password'

# List all Compute Nodes
[compute]
172.16.70.5 appformix_state=absent
172.16.70.17
```

2. Run Ansible with the edited inventory file. This will remove the node and all its resources from Contrail Insights.

```
ansible-playbook -i inventory appformix_openstack.yml
```

Upgrade Contrail Insights for an OpenStack Cluster

Contrail Insights can be easily upgraded by running the **appformix_openstack.yml** playbook of the new release. Follow the same procedure as the installation.

Uninstall Contrail Insights from an OpenStack Cluster

To uninstall Contrail Insights and destroy all data, execute the following command:

```
ansible-playbook -i <inventory_file> clean_appformix_openstack.yml
```

RELATED DOCUMENTATION

Contrail Insights General Requirements

Aggregate Discovery and Alarms with OpenStack Heat Services

Contrail Insights Installation for OpenStack in HA

IN THIS SECTION

- [HA Design Overview | 499](#)
- [Requirements | 499](#)
- [Install Contrail Insights for High Availability | 500](#)

HA Design Overview

Contrail Insights Platform can be deployed to multiple hosts for high availability (HA). Platform services continue to communicate using an API proxy that listens on a virtual IP address. Only one host will have the virtual IP at a time, and so only one API proxy will be the “active” API proxy at a time.

The API proxy is implemented by HAProxy. HAProxy is configured to use services in active-standby or load-balanced active-active mode, depending on the service.

At most, one host will be assigned the virtual IP at any given time. This host is considered the “active” HAProxy. The virtual IP address is assigned to a host by keepalived, which uses VRRP protocol for election.

Services are replicated in different modes of operation. In the “active-passive” mode, HAProxy sends all requests to a single “active” instance of a service. If the service fails, then HAProxy will select a new “active” from the other hosts, and begin to send requests to the new “active” service. In the “active-active” mode, HAProxy load balances requests across hosts on which a service is operational.

Contrail Insights Platform can be deployed in a 3-node, 5-node, or 7-node configuration for high availability.

Requirements

- For each host, on which Contrail Insights Platform is installed, see *Contrail Insights General Requirements* for hardware and software requirements. For a list of Contrail Insights Agent supported platforms, see *Contrail Insights Agent Requirements*.

- **NOTE:** Upgrade notice: Starting with Contrail Insights 3.2.6, the requirement for a license file is removed. If you are installing a version earlier than 3.2.6, a license is required prior to installation.

You can obtain a license key from APPFORMIX-KEY-REQUEST@juniper.net. Provide the following information in your request:

```
Group name:
Target customers or use:
Cluster type: OpenStack
Number of hosts:
Number of instances:
```

Connectivity

- One virtual IP address to be shared among all the Platform Hosts. This IP address should not be used by any host before installation. It should have reachability from all the Platform Hosts after installation.
- Dashboard client (in browser) must have IP connectivity to the virtual IP.
- IP addresses for each Platform Host for installation and for services running on these hosts to communicate.
- `keepalived_vrrp_interface` for each Platform Host which would be used for assigning virtual IP address. Details on how to configure this interface is described in the `sample_inventory` section.

Install Contrail Insights for High Availability

To install Contrail Insights to multiple hosts for high availability:

1. Download the Contrail Insights installation packages from [software downloads](#) to the Contrail Insights Platform node. Get the following files:

```
contrail-insights-<version>.tar.gz
contrail-insights-dependencies-images-<version>.tar.gz
contrail-insights-openstack-images-<version>.tar.gz
contrail-insights-platform-images-<version>.tar.gz
contrail-insights-network_device-images-<version>.tar.gz
```

If you are installing a version earlier than 3.2.6, copy the Contrail Insights license file to the Contrail Insights Platform node.

2. Install Ansible on the installer node. Ansible will install docker and the docker Python package on the **appformix_controller**.

```
# sudo apt-get install python-pip python-dev build-essential libssl-dev libffi-dev
# sudo pip install ansible==2.7.6 markupsafe httpplib2
```

For Ansible 2.3:

```
# sudo pip install ansible==2.3 markupsafe httpplib2 cryptography==1.5
```

3. Install Python and **python-pip** on all the Platform hosts so that Ansible can run between the installer node and the **appformix_controller** node.

```
# sudo apt-get install -y python python-pip
```

4. Install **python pip** package on the hosts where Contrail Insights Agents run.

```
# apt-get install -y python-pip
```

5. To enable passwordless login to all Platform hosts by Ansible, create an SSH public key on the node where Ansible playbooks are run and then copy the key to all the Platform hosts.

```
# ssh-keygen -t rsa                                #Creates Keys
# ssh-copy-id -i ~/.ssh/id_rsa.pub <platform_host_1>.....#Copies key from the
node to all platform hosts
# ssh-copy-id -i ~/.ssh/id_rsa.pub <platform_host_2>.....#Copies key from the
node to all platform hosts
# ssh-copy-id -i ~/.ssh/id_rsa.pub <platform_host_3>.....#Copies key from the
node to all platform hosts
```

6. Use the sample_inventory file as a template to create a host file. Add all the Platform hosts and compute hosts details.

```
# List all compute hosts which needs to be monitored by Contrail Insights
[compute]
203.0.113.5
203.0.113.17
# Contrail Insights controller hosts
[appformix_controller]
```

```
203.0.113.119 keepalived_vrrp_interface=eth0
203.0.113.120 keepalived_vrrp_interface=eth0
203.0.113.121 keepalived_vrrp_interface=eth0
```

NOTE: Note: In the case of 5-node or 7-node deployment, list all the nodes under **appformix_controller**.

7. At top-level of the distribution, create a directory named **group_vars** and then create a file named **all** inside this directory.

```
# mkdir group_vars
# touch group_vars/all
```

Add the following entries to the newly created **all** file:

```
appformix_vip: <ip-address>
appformix_docker_images:
- /path/to/contrail-insights-platform-images-<version>.tar.gz
- /path/to/contrail-insights-dependencies-images-<version>.tar.gz
- /path/to/contrail-insights-openstack-images-<version>.tar.gz
```

If you are installing a version earlier than 3.2.6, include the path to the Contrail Insights license file in **group_vars/all**:

```
appformix_license: path/to/<contrail-insights-license-file>.sig
```

8. Copy and source the **openrc** file from the OpenStack controller node (**/etc/contrail/openrc**) to the **appformix_controller** to authenticate the adapter to access admin privileges over the controller services.

```
root@installer_node:~# cat /etc/contrail/openrc
export OS_USERNAME=<admin user>
export OS_PASSWORD=<password>
export OS_TENANT_NAME=admin
export OS_AUTH_URL=http://<openstack-auth-URL>/v2.0/
export OS_NO_CACHE=1
root@installer_node:~# source /etc/contrail/openrc
```

NOTE: In Contrail Insights version 3.2.0, support for discovering OpenStack Octavia Load Balancer services is added. Contrail Insights only collects load balancer state information, such as **provisioning_status** and **operating_status**, as well as flavor information. To enable this service discovery, provide Octavia service's endpoint as variable **appformix_octavia_endpoint_url** in the **group_vars/all** file. For example:

```
appformix_octavia_endpoint_url: http://10.1.1.1:9876
```

Chargeback costs can also be configured for the Octavia Load Balancer services. See [Configure Load Balancer Costs](#).

9. Run Ansible with the created inventory file.

```
ansible-playbook -i inventory appformix_openstack_ha.yml
```

10. If running the playbooks as root user then this step can be skipped. As a non-root user (for example, "ubuntu"), the user "ubuntu" needs access to the **docker** user group. The following command adds the user to the docker group.

```
sudo usermod -aG docker ubuntu
```

RELATED DOCUMENTATION

Contrail Insights General Requirements

Contrail Insights Agent Requirements

Platform Dependencies

Post Installation Tasks

IN THIS CHAPTER

- [Configuring Role and Resource-Based Access Control | 504](#)
- [Configuring Role-Based Access Control for Analytics | 513](#)
- [Configuring the Control Node with BGP | 514](#)
- [Configuring MD5 Authentication for BGP Sessions | 524](#)
- [Configuring Transport Layer Security-Based XMPP in Contrail | 526](#)
- [Configuring Graceful Restart and Long-lived Graceful Restart | 528](#)
- [Scaling Up Contrail Networking Configuration API Server Instances | 534](#)
- [Scaling Up Contrail Networking Configuration API | 538](#)

Configuring Role and Resource-Based Access Control

IN THIS SECTION

- [Contrail Role and Resource-Based Access \(RBAC\) Overview | 504](#)
- [API-Level Access Control | 505](#)
- [Object Level Access Control | 506](#)
- [Configuration | 507](#)
- [Upgrading from Previous Releases | 509](#)
- [Configuring RBAC Using the Contrail User Interface | 509](#)
- [RBAC Resources | 513](#)

Contrail Role and Resource-Based Access (RBAC) Overview

Contrail Networking supports role and resource-based access control (RBAC) with API operation-level access control.

The RBAC implementation relies on user credentials obtained from Keystone from a token present in an API request. Credentials include user, role, tenant, and domain information.

API-level access is controlled by a list of rules. The attachment points for the rules include **global-system-config**, domain, and project. Resource-level access is controlled by permissions embedded in the object.

API-Level Access Control

If the RBAC feature is enabled, the API server requires a valid token to be present in the **X-Auth-Token** of any incoming request. The API server trades the token for user credentials (role, domain, project, and so on) from Keystone.

If a token is missing or is invalid, an HTTP error 401 is returned.

The **api-access-list** object holds access rules of the following form:

<object, field> => list of <role:CRUD>

Where:

object—An API resource such as network or subnet.

field—Any property or reference within the resource. The **field** option can be multilevel, for example, **network.ipam.host-routes** can be used to identify multiple levels. The **field** is optional, so in its absence, the create, read, update, and delete (CRUD) operation refers to the entire resource.

role—The Keystone role name.

Each rule also specifies the list of roles and their corresponding permissions as a subset of the CRUD operations.

Example: ACL RBAC Object

The following is an example access control list (ACL) object for a project in which the admin and any users with the **Development** role can perform CRUD operations on the network in a project. However, only the **admin** role can perform CRUD operations for policy and IP address management (IPAM) inside a network.

```
<virtual-network, network-policy> => admin:CRUD

<virtual-network, network-ipam> => admin:CRUD

<virtual-network, *>      => admin:CRUD, Development:CRUD
```

Rule Sets and ACL Objects

The following are the features of rule sets for access control objects in Contrail.

- The rule set for validation is the union of rules from the ACL attached to:

- User project
- User domain
- Default domain

It is possible for the project or domain access object to be empty.

- Access is only granted if a rule in the combined rule set allows access.
- There is no explicit deny rule.
- An ACL object can be shared within a domain. Therefore, multiple projects can point to the same ACL object. You can make an ACL object the default.

Object Level Access Control

The **perms2** permission property of an object allows fine-grained access control per resource.

The **perms2** property has the following fields:

owner —This field is populated at the time of creation with the tenant UUID value extracted from the token.

share list —The share list gets built when the object is selected for sharing with other users. It is a list of tuples with which the object is shared.

The **permission** field has the following options:

- **R**—Read object
- **W**—Create or update object
- **X**—Link (refer to) object

Access is allowed as follows:

- If the user is the owner and permissions allow (rwx)
- Or if the user tenant is in a shared list and permissions allow
- Or if world access is allowed

Configuration

IN THIS SECTION

- [Parameter: aaa-mode | 507](#)
- [Parameter: cloud_admin_role | 508](#)
- [Global Read-Only Role | 508](#)
- [Parameter Changes in /etc/neutron/api-paste.ini | 509](#)

This section describes the parameters used in Contrail RBAC.

Parameter: aaa-mode

RBAC is controlled by a parameter named **aaa-mode**. This parameter is used in place of the multi-tenancy parameter of previous releases.

The **aaa-mode** can be set to the following values:

- **no-auth**—No authentication is performed and full access is granted to all.
- **cloud-admin**—Authentication is performed and only the admin role has access.
- **rbac**—Authentication is performed and access is granted based on role.

If you are using Contrail Ansible Deployer to provision Contrail Networking, set the value for **AAA_MODE** to **rbac** to enable RBAC by default.

```
contrail_configuration:
  .
  .
  .
  AAA_MODE: rbac
```

If you are installing Contrail Networking from Contrail Command, specify the key and value as **AAA_MODE** and **rbac**, respectively, under the section Contrail Configuration on the **Step 2 Provisioning Options** page.

After enabling RBAC, you must restart the neutron server by running the **service neutron-server restart** command for the changes to take effect.

NOTE: The **multi_tenancy** parameter is deprecated, starting with Contrail 3.0. The parameter should be removed from the configuration. Instead, use the **aaa_mode** parameter for RBAC to take effect.

If the **multi_tenancy** parameter is not removed, the **aaa-mode** setting is ignored.

Parameter: cloud_admin_role

A user who is assigned the **cloud_admin_role** has full access to everything.

This role name is configured with the **cloud_admin_role** parameter in the API server. The default setting for the parameter is **admin**. This role must be configured in Keystone to change the default value.

If a user has the **cloud_admin_role** in one tenant, and the user has a role in other tenants, then the **cloud_admin_role** role must be included in the other tenants. A user with the **cloud_admin_role** doesn't need to have a role in all tenants, however, if that user has any role in another tenant, that tenant must include the **cloud_admin_role**.

Configuration Files with Cloud Admin Credentials

The following configuration files contain **cloud_admin_role** credentials:

- **/etc/contrail/contrail-keystone-auth.conf**
- **/etc/neutron/plugins/opencontrail/ContrailPlugin.ini**
- **/etc/contrail/contrail-webui-userauth.js**

Changing Cloud Admin Configuration Files

Modify the cloud admin credential files if the **cloud_admin_role** role is changed.

1. Change the configuration files with the new information.

2. Restart the following:

- API server
service supervisor-config restart
- Neutron server
service neutron-server restart
- WebUI
service supervisor-webui restart

Global Read-Only Role

You can configure a global read-only role (**global_read_only_role**).

A **global_read_only_role** allows read-only access to all Contrail resources. The **global_read_only_role** must be configured in Keystone. The default **global_read_only_role** is not set to any value.

A **global_read_only_role** user can use the Contrail Web Ui to view the global configuration of Contrail default settings.

Setting the Global Read-Only Role

To set the global read-only role:

1. The **cloud_admin** user sets the **global_read_only_role** in the Contrail API:

```
/etc/contrail/contrail-api.conf

global_read_only_role = <new-admin-read-role>
```

2. Restart the **contrail-api** service:

```
service contrail-api restart
```

Parameter Changes in /etc/neutron/api-paste.ini

Contrail RBAC operation is based upon a user token received in the **X-Auth-Token** header in API requests. The following change must be made in **/etc/neutron/api-paste.ini** to force Neutron to pass the user token in requests to the Contrail API server:

```
keystone = user_token request_id catch_errors ....
...
...
[filter:user_token]
paste.filter_factory =
neutron_plugin_contrail.plugins.opencontrail.neutron_middleware:token_factory
```

Upgrading from Previous Releases

The **multi_tenancy** parameter is deprecated.. The parameter should be removed from the configuration. Instead, use the **aaa_mode** parameter for RBAC to take effect.

If the **multi_tenancy** parameter is not removed, the **aaa-mode** setting is ignored.

Configuring RBAC Using the Contrail User Interface

To use the Contrail UI with RBAC:

1. Set the **aaa_mode** to **no_auth**.

```
/etc/contrail/contrail-analytics-api.conf
```

`aaa_mode = no-auth`

- Restart the **analytics-api** service.
`service contrail-analytics-api restart`


- Restart services by restarting the container.

You can use the Contrail UI to configure RBAC at both the API level and the object level. API level access control can be configured at the global, domain, and project levels. Object level access is available from most of the create or edit screens in the Contrail UI.

Configuring RBAC at the Global Level

To configure RBAC at the global level, navigate to **Configure > Infrastructure > Global Config > RBAC**, see [Figure 76 on page 510](#).


Figure 76: RBAC Global Level



Configuring RBAC at the Domain Level

To configure RBAC at the domain level, navigate to **Configure > RBAC > Domain**, see [Figure 77 on page 510](#).


Figure 77: RBAC Domain Level



Configuring RBAC at the Project Level

To configure RBAC at the project level, navigate to **Configure > RBAC > Project**, see [Figure 78 on page 511](#).

Figure 78: RBAC Project Level



Configuring RBAC Details

Configuring RBAC is similar at all of the levels. To add or edit an API access list, navigate to the global, domain, or project page, then click the plus (+) icon to add a list, or click the gear icon to select from Edit, Insert After, or Delete, see [Figure 79 on page 511](#).


Figure 79: RBAC Details API Access



Creating or Editing API Level Access

Clicking create, edit, or insert after activates the Edit API Access popup window, where you enter the details for the API Access Rules. Enter the user type in the Role field, and use the + icon in the Access field to enter the types of access allowed for the role, including, Create, Read, Update, Delete, and so on, see [Figure 80 on page 512](#).


Figure 80: Edit API Access



Creating or Editing Object Level Access

You can configure fine-grained access control by resource. A **Permissions** tab is available on all create or edit popups for resources. Use the **Permissions** popup to configure owner permissions and global share permissions. You can also share the resource to other tenants by configuring it in the **Share List**, see [Figure 81 on page 512](#).

Figure 81: Edit Object Level Access



RBAC Resources

Refer to the *OpenStack Administrator Guide* for additional information about RBAC:

- [Identity API protection with role-based access control \(RBAC\)](#)

Configuring Role-Based Access Control for Analytics

The analytics API uses role-based access control (RBAC) to provide the ability to access UVE and query information based on the permissions of the user for the UVE or queried object.

Contrail Networking extends authenticated access so that tenants can view network monitoring information about the networks for which they have read permissions.

The analytics API can map query and UVE objects to configuration objects on which RBAC rules are applied, so that read permissions can be verified using the VNC API.

RBAC is applied to analytics in the following ways:

- For statistics queries, annotations are added to the Sandesh file so that indices and tags on statistics queries can be associated with objects and UVEs. These are used by the contrail-analytics-api to determine the object level read permissions.
- For flow and log queries, the object read permissions are evaluated for each AND term in the where query.
- For UVEs list queries (e.g. analytics/uve/virtual-networks/), the contrail-analytics-api gets a list of UVEs that have read permissions for a given token. For a UVE query for a specific resource (e.g. analytics/uves/virtual-network/vn1), contrail-analytics-api checks the object level read permissions using VNC API.

Tenants cannot view system logs and flow logs, those logs are displayed for cloud-admin roles only.

A non-admin user can see only non-global UVEs, including:

- virtual_network
- virtual_machine
- virtual_machine_interface
- service_instance
- service_chain
- tag
- firewall_policy

- firewall_rule
- address_group
- service_group
- application_policy_set

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

Configuring the Control Node with BGP

IN THIS SECTION

- [Configuring the Control Node from Contrail Web UI | 515](#)
- [Configuring the Control Node with BGP from Contrail Command | 521](#)

An important task after a successful installation is to configure the control node with BGP. This procedure shows how to configure basic BGP peering between one or more virtual network controller control nodes and any external BGP speakers. External BGP speakers, such as Juniper Networks MX80 routers, are needed for connectivity to instances on the virtual network from an external infrastructure or a public network.

Before you begin, ensure that the following tasks are completed:

- The Contrail Controller base system image has been installed on all servers.
- The role-based services have been assigned and provisioned.
- IP connectivity has been verified between all nodes of the Contrail Controller.
- You have access to Contrail Web User Interface (UI) or Contrail Command User Interface (UI). You can access the user interface at <http://nn.nn.nn.nn:8143>, where *nn.nn.nn.nn* is the IP address of the configuration node server that is running the contrail service.


These topics provide instructions to configure the Control Node with BGP.

Configuring the Control Node from Contrail Web UI

To configure BGP peering in the control node:

1. From the Contrail Controller module control node (<http://nn.nn.nn.nn:8143>), select **Configure > Infrastructure > BGP Routers**; see [Figure 82 on page 516](#).

Figure 82: Configure> Infrastructure > BGP Routers



A summary screen of the control nodes and BGP routers is displayed; see [Figure 83 on page 516](#).

Figure 83: BGP Routers Summary

Configure > Infrastructure > BGP Routers

BGP Routers + 🗑️ 📄 🔍 ⬆️

☐ IP Address	Type	Vendor	HostName	
▶ ☐ 10.84.25.31	Control Node	contrail	b5s31	⚙️
▶ ☐ 10.84.11.252	BGP Router	mx	a3-mx80-1	⚙️
▶ ☐ 10.84.25.30	Control Node	contrail	b5s30	⚙️
▶ ☐ 10.84.25.29	Control Node	contrail	b5s29	⚙️
▶ ☐ 10.84.25.28	Control Node	contrail	b5s28	⚙️
▶ ☐ 10.84.25.27	Control Node	contrail	b5s27	⚙️
▶ ☐ 10.84.11.253	BGP Router	mx	mx1	⚙️

Total: 7 records | 50 Records ▼ | Page 1 ▼ of 1

- (Optional) The global AS number is 64512 by default. To change the AS number, on the **BGP Router** summary screen click the gear wheel and select **Edit**. In the Edit BGP Router window enter the new number.
- To create control nodes and BGP routers, on the **BGP Routers** summary screen, click the **+** icon. The **Create BGP Router** window is displayed; see [Figure 84 on page 518](#).

Figure 84: Create BGP Router

4. In the **Create BGP Router** window, click **BGP Router** to add a new BGP router or click **Control Node** to add control nodes.

For each node you want to add, populate the fields with values for your system. See [Table 30 on page 518](#).

Table 30: Create BGP Router Fields


Field	Description
Hostname	Enter a name for the node being added.
Vendor ID	Required for external peers. Populate with a text identifier, for example, "MX-0". (BGP peer only)
IP Address	The IP address of the node.
Router ID	Enter the router ID.
Autonomous System	Enter the AS number in the range 1-65535 for the node. (BGP peer only)

Table 30: Create BGP Router Fields (*continued*)

Field	Description
Address Families	Enter the address family, for example, inet-vpn
Hold Time	BGP session hold time. The default is 90 seconds; change if needed.
BGP Port	The default is 179; change if needed.
Authentication Mode	Enable MD5 authentication if desired.
Authentication key	Enter the Authentication Key value.
Physical Router	The type of the physical router.
Available Peers	Displays peers currently available.
Configured Peers	Displays peers currently configured.


- Click **Save** to add each node that you create.
- To configure an existing node as a peer, select it from the list in the **Available Peers** box, then click >> to move it into the **Configured Peers** box.
Click << to remove a node from the **Configured Peers** box.
- You can check for peers by selecting **Monitor > Infrastructure > Control Nodes**; see [Figure 85 on page 520](#).

Figure 85: Control Nodes



In the **Control Nodes** window, click any hostname in the memory map to view its details; see [Figure 86 on page 520](#).

Figure 86: Control Node Details



8. Click the **Peers** tab to view the peers of a control node; see [Figure 87 on page 521](#).

Figure 87: Control Node Peers Tab

Monitor > Infrastructure > Control Nodes > b5s29

DetailsPeersRoutesConsole

Peers

Download

Search

Refresh

Peer	Peer Type	Peer ASN	Status	Last flap	Messages (Recv/Sent)
▶ 10.84.21.1	XMPP	-	Established, in sync	-	35497 / 138229
▶ 10.84.21.10	XMPP	-	Established, in sync	-	35511 / 137011
▶ 10.84.21.11	XMPP	-	Established, in sync	-	37045 / 141735
▶ 10.84.21.12	XMPP	-	Established, in sync	-	37493 / 140054
▶ 10.84.21.13	XMPP	-	Established, in sync	-	35540 / 137864
▶ 10.84.21.14	XMPP	-	Established, in sync	-	40098 / 112770
▼ 10.84.21.15	XMPP	-	Established, in sync	-	35450 / 137599

Details :

- {

name: "b5s29:10.84.21.15",

value: - {

xmppPeerInfoData: - {

state_info: - {

last_state: "Active",

state: "Established",

}

}

}

}

5042501

Configuring the Control Node with BGP from Contrail Command

To configure BGP peering in the control node:

1. From Contrail Command UI select **Infrastructure > Cluster > Advanced** page.

Click the **BGP Routers** tab. A list of control nodes and BGP routers is displayed. See [Figure 88 on page 521](#).

Figure 88: Infrastructure > Cluster > Advanced > BGP Routers

CONTRAIL
COMMAND

INFRASTRUCTURE > Clusters

admin

Admin

Search

FAVORITES >

MONITORING >

INFRASTRUCTURE >

OVERLAY >

WORKLOADS >

IAM >

SERVICES >

SECURITY >

Clusters

STATUS

NAME

●

test-01

Upgrade

Add Cluster

2. (Optional) The global AS number is 64512 by default. You can change the AS number according to your requirement on the **BGP Router** tab, by clicking the **Edit** icon. In the **Edit BGP Router** tab enter AS

number in the range of 1-65,535. You can also enter the AS number in the range of 1-4,294,967,295, when **4 Byte ASN** is enabled in **Global Config**.

3. Click the **Create** button on the **BGP Routers** tab. The **Create BGP Router** window is displayed. See [Figure 89 on page 522](#).

Figure 89: Create BGP Router

The screenshot shows the 'Create BGP Router' window in the Cortex XDR console. The window has a dark header with the Cortex XDR logo and navigation tabs: INFRASTRUCTURE, Cluster, Advanced, and Create BGP Router. On the left, there is a sidebar with a search bar and a list of navigation items: FAVORITES, MONITORING, INFRASTRUCTURE (selected), OVERLAY, WORKLOADS, IAM, and SERVICES. The main content area is divided into three tabs: BGP, Tags, and Permissions. The BGP tab is active, showing various input fields and sections. The fields include Router Type (BGP Router), Hostname (with a red asterisk), Vendor ID (with a red asterisk), IP Address (with a red asterisk and a hint 'Enter valid IPv4'), Router ID (with a red asterisk and a hint 'Enter valid IPv4'), Autonomous System (64512), and BGP Router ASN (ASN). Below these fields are sections for Associate Peers and Advanced Options. The Advanced Options section includes BGP Port, Source Port, Hold Time (seconds) (90), Admin State (checked), Authentication Mode (none), Authentication Key, Control Node Zone, and Physical Router. At the bottom, there are 'Create' and 'Cancel' buttons.

4. In the **Create BGP Router** page, populate the fields with values to create your system. See [Table 31 on page 522](#).

Table 31: Create BGP Router

Fields	Description
Router Type	Select the type of router you want create
Hostname	Enter a name for the node being added.
Vendor ID	Required for external peers. Populate with a text identifier, for example, "MX-0". (BGP peer only)
IP Address	The IP address of the node.
Router ID	Enter the router ID.

Table 31: Create BGP Router (*continued*)

Fields	Description
Autonomous System (AS)	Enter autonomous system (AS) number in the range of 1-65,535. If you enable 4 Byte ASN in Global Config , you can enter 4-byte AS number in the range of 1-4,294,967,295.
BGP Router ASN	Enter the Local-AS number, specific to the associated peers.
Address Families	Select the Internet Address Family from the list, for example, inet-vpn , inet6-vpn , and so on.
Cluster ID	Enter the cluster ID, for example, 0.0.0.100.
Associate Peers	
Peer	Select the configured peers from the list.
Hold Time	Enter the maximum time a BGP session remains active if no Keepalives are received.
Loop Count	Enter the number of times the same ASN can be seen in a route-update. The route is discarded when the loop count is exceeded.
MD5 Auth Key	Enter the MD5 authentication key value.
State	Select the state box when you are associating BGP peers.
Passive	Select the passive box to disable the BGP router from advertising any routes. The BGP router can only receive updates from other peers in this state.
Advanced Options	
BGP Port	Enter BGP Port number. The default is 179; change if needed.
Source Port	Enter source port number for client side connection.
Hold Time (seconds)	BGP session hold time. The default is 90 seconds; change if needed.
Admin State	Select the Admin state box to enable the state as UP and deselect the box to disable the state to DOWN.

Table 31: Create BGP Router (*continued*)

Fields	Description
Authentication Mode	Select MD5 from list if required.
Authentication key	Enter the Authentication Key value.
Control Node Zone	Select the required control node zone from the list.
Physical Router	Select the the physical router from the list.

5. Click **Create** to complete add each node.
6. You can check for peers and details about the control nodes by selecting **Infrastructure > Cluster > Control Nodes**. Click the desired node to check the details on **Summary** and **Detailed Stats** page.

RELATED DOCUMENTATION

Creating a Virtual Network with Juniper Networks Contrail

[Creating a Virtual Network with OpenStack Contrail | 342](#)

Configuring MD5 Authentication for BGP Sessions

Contrail supports MD5 authentication for BGP peering based on RFC 2385.

This option allows BGP to protect itself against the introduction of spoofed TCP segments into the connection stream. Both of the BGP peers must be configured with the same MD5 key. Once configured, each BGP peer adds a 16-byte MD5 digest to the TCP header of every segment that it sends. This digest is produced by applying the MD5 algorithm on various parts of the TCP segment. Upon receiving a signed segment, the receiver validates it by calculating its own digest from the same data (using its own key) and compares the two digests. For valid segments, the comparison is successful since both sides know the key.

The following are ways to enable BGP MD5 authentication and set the keys on the Contrail node.

1. If the **md5** key is not included in the provisioning, and the node is already provisioned, you can run the following script with an argument for md5:


```
contrail-controller/src/config/utils/provision_control.py
```

```
host@<your_node>:/opt/contrail/utils# python provision_control.py --host_name
<host_name> --host_ip <host_ip> --router_asn <asn> --api_server_ip <api_ip>
--api_server_port <api_port> --oper add --md5 "juniper" --admin_user admin
--admin_password <password> --admin_tenant_name admin
```

2. You can also use the web user interface to configure MD5.

- a. Connect to the node's IP address at port 8080 (<node_ip>:8080) and select **Configure->Infrastructure->BGP Routers**. As shown in [Figure 90 on page 525](#), a list of BGP peers is displayed.

Figure 90: Edit BGP Router Window



- b. For a BGP peer, click on the gear icon on the right hand side of the peer entry. Then click **Edit**. This displays the Edit BGP Router dialog box.
- c. Scroll down the window and select **Advanced Options**.
- d. Configure the MD5 authentication by selecting **Authentication Mode>MD5** and entering the **Authentication Key** value.

RELATED DOCUMENTATION

Creating a Virtual Network with Juniper Networks Contrail

Configuring Transport Layer Security-Based XMPP in Contrail

Overview: TLS-Based XMPP

Transport Layer Security (TLS)-based XMPP can be used to secure all Extensible Messaging and Presence Protocol (XMPP)-based communication that occurs in the Contrail environment.

Secure XMPP is based on *RFC 6120, Extensible Messaging and Presence Protocol (XMPP): Core*.

TLS XMPP in Contrail

In the Contrail environment, the Transport Layer Security (TLS) protocol is used for certificate exchange, mutual authentication, and negotiating ciphers to secure the stream from potential tampering and eavesdropping.

The RFC 6120 highlights a basic stream message exchange format for TLS negotiation between an XMPP server and an XMPP client.

NOTE: Simple Authentication and Security Layer (SASL) authentication is not supported in the Contrail environment.

Configuring XMPP Client and Server in Contrail

In the Contrail environment, XMPP based communications are used in client and server exchanges, between the compute node (as the XMPP client), and:

- the control node (as the XMPP server)
- the DNS server (as the XMPP server)

Configuring Control Node for XMPP Server

To enable secure XMPP, the following parameters are configured at the XMPP server.

On the control node, enable the parameters in the configuration file:
`/etc/contrail/contrail-control.conf`.

Parameter	Description	Default
<code>xmpp_server_cert</code>	Path to the node's public certificate	<code>/etc/contrail/ssl/certs/server.pem</code>

Parameter	Description	Default
xmpp_server_key	Path to server's or node's private key	/etc/contrail/ssl/private/server-privkey.pem
xmpp_ca_cert	Path to CA certificate	/etc/contrail/ssl/certs/ca-cert.pem
xmpp_auth_enable=true	Enables SSL based XMPP	Default is set to false, XMPP is disabled. NOTE: The keyword true is case sensitive.

Configuring DNS Server for XMPP Server

To enable secure XMPP, the following parameters are configured at the XMPP DNS server.

On the DNS server control node, enable the parameters in the configuration file:

/etc/contrail/contrail-control.conf

Parameter	Description	Default
xmpp_server_cert	Path to the node's public certificate	/etc/contrail/ssl/certs/server.pem
xmpp_server_key	Path to server's/node's private key	/etc/contrail/ssl/certs/server-privkey.pem
xmpp_ca_cert	Path to CA certificate	/etc/contrail/ssl/certs/ca-cert.pem
xmpp_dns_auth_enable=true	Enables SSL based XMPP	Default is set to false, XMPP is disabled. NOTE: The keyword true is case sensitive.

Configuring Control Node for XMPP Client

To enable secure XMPP, the following parameters are configured at the XMPP client.

On the compute node, enable the parameters in the configuration file:

/etc/contrail/contrail-vrouter-agent.conf

Parameter	Description	Default
xmpp_server_cert	Path to the node's public certificate	/etc/contrail/ssl/certs/server.pem
xmpp_server_key	Path to server's/node's private key	/etc/contrail/ssl/private/server-privkey.pem
xmpp_ca_cert	Path to CA certificate	/etc/contrail/ssl/certs/ca-cert.pem
xmpp_auth_enable=true xmpp_dns_auth_enable=true	Enables SSL based XMPP	Default is set to false, XMPP is disabled. NOTE: The keyword true is case sensitive.

Configuring Graceful Restart and Long-lived Graceful Restart

IN THIS SECTION

- [Application of Graceful Restart and Long-lived Graceful Restart | 528](#)
- [BGP Graceful Restart Helper Mode | 529](#)
- [Feature Highlights | 529](#)
- [XMPP Helper Mode | 529](#)
- [Configuration Parameters | 530](#)
- [Cautions for Graceful Restart | 531](#)
- [Configuring Graceful Restart using Contrail Command | 532](#)

Graceful restart and long-lived graceful restart BGP helper modes are supported for the Contrail control node and XMPP helper mode.

Application of Graceful Restart and Long-lived Graceful Restart

Whenever a BGP peer session is detected as down, all routes learned from the peer are deleted and immediately withdrawn from advertised peers. This causes instantaneous disruption to traffic flowing end-to-end, even when routes kept in the vrouter kernel in the data plane remain intact.

Graceful restart and long-lived graceful restart features can be used to alleviate traffic disruption caused by downs.

When configured, graceful restart features enable existing network traffic to be unaffected if Contrail controller processes go down. The Contrail implementation ensures that if a Contrail control module restarts, it can use graceful restart functionality provided by its BGP peers. Or when the BGP peers restart, Contrail provides a graceful restart helper mode to minimize the impact to the network. The graceful restart features can be used to ensure that traffic is not affected by temporary outage of processes.

Graceful restart is not enabled by default.

With graceful restart features enabled, learned routes are not deleted when sessions go down, and the routes are not withdrawn from the advertised peers. Instead, the routes are kept and marked as 'stale'. Consequently, if sessions come back up and routes are relearned, the overall impact to the network is minimized.

After a certain duration, if a downed session does not come back up, all remaining stale routes are deleted and withdrawn from advertised peers.

The graceful restart and long-lived graceful restart features can be enabled only for BGP peers in Contrail 3.2.

BGP Graceful Restart Helper Mode

The BGP helper mode can be used to minimize routing churn whenever a BGP session flaps. This is especially helpful if the SDN gateway router goes down gracefully, as in an rpd crash or restart on an MX Series Junos device. In that case, the contrail-control can act as a graceful restart helper to the gateway, by retaining the routes learned from the gateway and advertising them to the rest of the network as applicable. In order for this to work, the restarting router (the SDN gateway in this case) must support and be configured with graceful restart for all of the address families used.

The graceful restart helper mode is also supported for BGP-as-a-Service (BGPaaS) clients. When configured, contrail-control can provide a graceful restart or long-lived graceful restart helper mode to a restarting BGPaaS client.

Feature Highlights

The following are highlights of the graceful restart and long-lived graceful restart features.

- Configuring a non-zero restart time enables the ability to advertise graceful restart and long-lived graceful restart capabilities in BGP.
- Configuring helper mode enables the ability for graceful restart and long-lived graceful restart helper modes to retain routes even after sessions go down.
- With graceful restart configured, whenever a session down event is detected and a closing process is triggered, all routes, across all address families, are marked stale. The stale routes are eligible for best-path election for the configured graceful restart time duration.
- When long-lived graceful restart is in effect, stale routes can be retained for a much longer time than that allowed by graceful restart alone. With long-lived graceful restart, route preference is retained and best paths are recomputed. The community marked LLGR_STALE is tagged for stale paths and re-advertised. However, if no long-lived graceful restart community is associated with any received stale route, those routes are not kept, instead, they are deleted.
- After a certain time, if a session comes back up, any remaining stale routes are deleted. If the session does not come back up, all retained stale routes are permanently deleted and withdrawn from the advertised peer.

XMPP Helper Mode

Contrail supports for long-lived graceful restart (LLGR) with XMPP helper mode. Graceful restart and long lived graceful restart can be enabled using the Contrail web UI or by using the provision_control script.

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.

Configuration Parameters

Graceful restart parameters are configured in the **global-system-config** of the schema. They can be configured by means of a provisioning script or by using the Contrail Web UI.

Configure a non-zero restart time to advertise for graceful restart and long-lived graceful restart capabilities from peers.

Configure helper mode for graceful restart and long-lived graceful restart to retain routes even after sessions go down.

Configuration parameters include:

- **enable** or **disable** for all graceful restart parameters:
 - **restart-time**
 - **long-lived-restart-time**
 - **end-of-rib-timeout**
- **bgp-helper-enable** to enable graceful restart helper mode for BGP peers in contrail-control
- **xmpp-helper-enable** to enable graceful restart helper mode for XMPP peers (agents) in contrail-control

The following shows configuration by a provision script.

```
/opt/contrail/utils/provision_control.py
--api_server_ip 10.xx.xx.20
--api_server_port 8082
--router_asn 64512
--admin_user admin
--admin_password <password>
--admin_tenant_name admin
--set_graceful_restart_parameters
--graceful_restart_time 60
--long_lived_graceful_restart_time 300
--end_of_rib_timeout 30
--graceful_restart_enable
--graceful_restart_bgp_helper_enable
```

The following are sample parameters:

```
-set_graceful_restart_parameters
    --graceful_restart_time 300
    --long_lived_graceful_restart_time 60000
    --end_of_rib_timeout 30
    --graceful_restart_enable
    --graceful_restart_bgp_helper_enable
```

When BGP peering with Juniper Networks devices, Junos must also be explicitly configured for graceful restart/long-lived graceful restart, as shown in the following example:

```
set routing-options graceful-restart
set protocols bgp group <a1234> type internal
set protocols bgp group <a1234> local-address 10.xx.xxx.181
set protocols bgp group <a1234> keep all
set protocols bgp group <a1234> family inet-vpn unicast graceful-restart long-lived
  restarter stale-time 20
set protocols bgp group <a1234> family route-target graceful-restart long-lived
  restarter stale-time 20
set protocols bgp group <a1234> graceful-restart restart-time 600
set protocols bgp group <a1234> neighbor 10.xx.xx.20 peer-as 64512
```

The graceful restart helper modes can be enabled in the schema. The helper modes can be disabled selectively in the **contrail-control.conf** for BGP sessions by configuring **gr_helper_disable** in the **/etc/contrail/contrail-control.conf** file.

The following are examples:

```
/usr/bin/openstack-config /etc/contrail/contrail-control.conf DEFAULT gr_helper_bgp_disable 1
```

```
/usr/bin/openstack-config /etc/contrail/contrail-control.conf DEFAULT gr_helper_xmpp_disable 1
```

```
service contrail-control restart
```

For more details about graceful restart configuration, see

<https://github.com/Juniper/contrail-controller/wiki/Graceful-Restart> .

Cautions for Graceful Restart


Be aware of the following caveats when configuring and using graceful restart.

- Using the graceful restart/long-lived graceful restart feature with a peer is effective either to all negotiated address families or to none. If a peer signals support for graceful restart/long-lived graceful restart for only a subset of the negotiated address families, the graceful restart helper mode does not come into effect for any family in the set of negotiated address families.
- Because graceful restart is not yet supported for contrail-vrouter-agent, the parameter should *not* be set for **graceful_restart_xmpp_helper_enable**. If the vrouter agent restarts, the data plane is reset and the routes and flows are reprogrammed anew, which typically results in traffic loss for several seconds for new and /existing flows.
- Graceful restart/long-lived graceful restart is not supported for multicast routes.
- Graceful restart/long-lived graceful restart helper mode may not work correctly for EVPN routes, if the restarting node does not preserve forwarding state for EVPN routes.

Configuring Graceful Restart using Contrail Command


To configure graceful restart in Contrail Command, navigate to **Infrastructure > Cluster > Advanced Options** and select the Edit icon near the top right corner of the screen.

Figure 91: Global Config System Configuration Screen



The **Edit System Configuration** window opens. Click the box for **Graceful Restart** to enable graceful restart, and enter a non-zero number to define the **Restart Time** in seconds. You can also specify the times for the long-lived graceful restart (LLGR) and the end of RIB timers from this window.

Figure 92: Edit System Configuration Window



Scaling Up Contrail Networking Configuration API Server Instances

Certain deployment scenarios may need running of multiple API configuration server instances for improved performance. One of the methods to achieve this is by increasing the number of API configuration server instances on a node *after deploying* Contrail Networking. This is done by modifying the **docker-compose.yaml** file to allow multiple configuration API containers on the same host.

The steps described in this topic are valid for Contrail Networking Releases 5.0 through 2008. For release 2011 and later, refer the topic [“Scaling Up Contrail Networking Configuration API” on page 538](#).



CAUTION: Any change to the Contrail networking Configuration API must be made only with the help of Juniper Networks Professional Services. We strongly recommend that you contact Juniper Networks Professional Services before you make any change to the Configuration API.

1. Edit the `/etc/contrail/config/docker-compose.yaml` file on each Contrail Configuration node. At the end of the API section, at line 56, append the following chunk of code:

```
api-2:
  image: "hub.juniper.net/contrail/contrail-controller-config-api:1912.32-rhel"

  env_file: /etc/contrail/common_config.env
  environment:
    - CONFIG_API_PORT=18082
    - CONFIG_API_INTROSPECT_PORT=17084
  container_name: config_api_2
  command: ["/usr/bin/contrail-api", "--conf_file",
"/etc/contrail/contrail-api.conf", "--conf_file",
"/etc/contrail/contrail-keystone-auth.conf", "--worker_id", "2"]
  network_mode: "host"
  volumes_from:
    - node-init
  depends_on:
    - node-init
  restart: always
  stdin_open: True
  tty: True
  logging:
    driver: "json-file"
    options:
      max-size: "50m"
      max-file: "10"
```

Make sure that the API port, Introspect port, Worker ID and container_name are unique within the node. The default API section should not be changed to allow other Contrail services (for example, Schema-transformer, SVC-Monitor or contrail-status command) to run without configuration changes. The default API runs with port 8082, introspect port 8084, worker_id 0 and container name config_api_1.

2. Run the following commands on each configuration node to apply the changes:

```
cd /etc/contrail/config
docker-compose down
docker-compose up -d
```

3. Run the following commands on each configuration node to verify the configuration:

```
pgrep -al contrail-api
ss -nlpt | grep contrail-api
```

```
# pgrep -al contrail-api
16922 /usr/bin/python /usr/bin/contrail-api --conf_file
/etc/contrail/contrail-api.conf --conf_file
/etc/contrail/contrail-keystone-auth.conf --worker_id 2
17139 /usr/bin/python /usr/bin/contrail-api --conf_file
/etc/contrail/contrail-api.conf --conf_file
/etc/contrail/contrail-keystone-auth.conf --worker_id 0
# ss -nlpt | grep contrail-api
LISTEN      0          128          *:17084      *: *
    users:(("contrail-api",pid=16922,fd=7))
LISTEN      0          128      10.0.0.16:18082  *: *
    users:(("contrail-api",pid=16922,fd=28))
LISTEN      0          128      10.0.0.16:8082   *: *
    users:(("contrail-api",pid=17139,fd=28))
LISTEN      0          128          *:8084      *: *
    users:(("contrail-api",pid=17139,fd=7))
```

4. Add new server instances to load-balancer.

```
listen contrail_config_internal
    bind 192.168.3.90:8082 transparent
    mode http
    balance leastconn
    option httpchk GET /
    option httplog
    option forwardfor
    timeout client 360s
    timeout server 360s
    server 0contrail-ctl0.local 192.168.3.114:8082 check fall 5 inter 2000 rise
2
    server 1contrail-ctl0.local 192.168.3.114:18081 check fall 5 inter 2000 rise
2
    server 2contrail-ctl0.local 192.168.3.114:18082 check fall 5 inter 2000 rise
2
    server 3contrail-ctl0.local 192.168.3.114:18083 check fall 5 inter 2000 rise
2
    server 0contrail-ctl1.local 192.168.3.139:8082 check fall 5 inter 2000 rise
2
    server 1contrail-ctl1.local 192.168.3.139:18081 check fall 5 inter 2000 rise
```



```

2
  server 2contrail-ctl1.local 192.168.3.139:18082 check fall 5 inter 2000 rise
2
  server 3contrail-ctl1.local 192.168.3.139:18083 check fall 5 inter 2000 rise
2
  server 0contrail-ctl2.local 192.168.3.100:8082 check fall 5 inter 2000 rise
2
  server 1contrail-ctl2.local 192.168.3.100:18081 check fall 5 inter 2000 rise
2
  server 2contrail-ctl2.local 192.168.3.100:18082 check fall 5 inter 2000 rise
2
  server 3contrail-ctl2.local 192.168.3.100:18083 check fall 5 inter 2000 rise
2

```

5. Run the following commands to set load-balancer timeouts and balancing method

You cannot configure timeout for the Neutron plugin. Neutron relies on load-balancer to terminate connections. Therefore, it is important that you increase the default timeout value of 30 seconds. This enables Neutron to respond without error even if the Contrail API responds longer than 30 seconds.

Sample timeout options for HAProxy

```

timeout client 360s
timeout server 360s

```

NOTE: It is recommended that you use load balancing methods that distributes the load evenly across the configuration API server instances. The preferred load balancing methods are **leastconn** and **round-robin**.



CAUTION: Increasing the number of configuration API server instances might result in higher load on RabbitMQ and Schema-transformer.

Scaling Up Contrail Networking Configuration API

Starting from Contrail Networking Release 2011, config-api can be scaled vertically, to run up to 10 workers in a single container by using uWSGI. However, the recommended maximum number of workers in each config_api container is 5.

The steps described in this topic are valid for Contrail Networking Releases 2011 and later releases. To see these procedures for earlier releases, see [“Scaling Up Contrail Networking Configuration API Server Instances” on page 534](#).

From Contrail Ansible Deployer

If you are deploying Contrail Networking using Ansible Deployer , you should specify **CONFIG_API_WORKER_COUNT** parameter in the contrail_configuration section of the instances.yml file as shown below.

```
contrail_configuration:
  CONFIG_API_WORKER_COUNT: 5
```

From Contrail Command UI

If you are deploying Contrail Networking from Contrail Command UI, you should specify a new key / value pair (CONFIG_API_WORKER_COUNT / *value*) in the Contrail Configuration section as shown in [Figure 93 on page 539](#).

Figure 93: Scaling API Configuration from Contrail Command

The screenshot shows the Contrail Command web interface. The left sidebar contains a search bar, 'FAVORITES', and a menu with 'INFRASTRUCTURE' selected. The main content area is titled 'INFRASTRUCTURE > Clusters'. A vertical progress bar on the left lists steps: STEP 1 Inventory, STEP 2 Provisioning Options (highlighted), STEP 3 Infrastructure Networks, STEP 4 (optional) Overcloud, STEP 5 (optional) Undercloud Nodes, STEP 6 (optional) Jumphost Nodes, STEP 7 Control Nodes, and STEP 8 (optional) Orchestrator Nodes. The 'Provisioning Options' section contains several input fields: 'Container Registry Password' (with an eye icon), 'Contrail Version' (set to 'latest'), 'Domain Suffix' (set to 'local'), 'NTP Server', 'Default Vrouter Gateway', and 'Encapsulation Priority' (set to 'VXLAN,MPLSoGRE,MPL...'). There is a checkbox for 'Fabric Management'. Below these is a 'Contrail Configuration' section with a table for key-value pairs. The table has two columns: 'Key' and 'Value'. A row is highlighted with a blue border, showing 'CONFIG_API_WORKER_COUNT' as the key and '5' as the value. Below the table is a '+ Add' button. At the bottom of the configuration section are 'Previous' and 'Next' buttons.

From RHOSP Deployer

If you are deploying Contrail Networking using RHOSP deployer, you should specify the desired value for the parameter **CONFIG_API_WORKER_COUNT** in the **ContrailSettings** section of the **contrail-services.yaml** file as shown below.

```
ContrailSettings:
  CONFIG_API_WORKER_COUNT: 3
```

From JUJU Deployer

If you are deploying Contrail Networking using JUJU deployer, you must specify the desired value for **config-api-worker-count** in the **config.yaml** file.

```
config-api-worker-count:
  default: 1
  type: int
  description: |
    Number of workers spawned inside config-api container.
```