



---

Contrail™

## Contrail Getting Started Guide

Release

3.1



---

Modified: 2016-09-08

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

Copyright © 2016, Juniper Networks, Inc. All rights reserved.

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

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

*Contrail™ Contrail Getting Started Guide*

3.1

Copyright © 2016, 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 <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

# Table of Contents

	About the Documentation . . . . .	xxi
	Documentation and Release Notes . . . . .	xxi
	Documentation Conventions . . . . .	xxi
	Documentation Feedback . . . . .	xxiii
	Requesting Technical Support . . . . .	xxiv
	Self-Help Online Tools and Resources . . . . .	xxiv
	Opening a Case with JTAC . . . . .	xxiv
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Understanding Contrail . . . . .</b>	<b>3</b>
	Contrail Overview . . . . .	3
	Contrail Description . . . . .	4
	Contrail Major Components . . . . .	4
	Contrail Control Nodes . . . . .	4
	Contrail Compute Nodes – XMPP Agent and vRouter . . . . .	4
	Contrail Solution . . . . .	4
<b>Part 2</b>	<b>Installing and Upgrading Contrail</b>	
<b>Chapter 2</b>	<b>Supported Platforms and Server Requirements . . . . .</b>	<b>9</b>
	Supported Platforms . . . . .	9
	Server Requirements . . . . .	9
<b>Chapter 3</b>	<b>Installing Contrail and Provisioning Roles . . . . .</b>	<b>11</b>
	Installation Overview . . . . .	11
	Downloading Installation Software . . . . .	12
	Installing the Operating System and Contrail Packages . . . . .	13
	Configuring System Settings . . . . .	14
	Installing the Contrail Packages, Part One (CentOS or Ubuntu) . . . . .	15
	Setting Up the Testbed Definitions File . . . . .	16
	Testbed Definitions File Settings for Deploying Contrail with an Existing OpenStack Node . . . . .	19
	Supporting Multiple Interfaces on Servers and Nodes . . . . .	21
	Support for Multiple Interfaces . . . . .	21
	Number of cfgm Nodes Supported . . . . .	22
	Uneven Number of Database Nodes Required . . . . .	22
	Support for VLAN Interfaces . . . . .	22
	Support for Bonding Options . . . . .	22
	Support for Static Route Options . . . . .	22
	Server Interface Examples . . . . .	23
	Interface Naming and Configuration Management . . . . .	23

	Setting Up Interfaces and Installing . . . . .	23
	Sample testbed.py File With Exclusive Interfaces . . . . .	24
	Installing the Contrail Packages, Part Two (CentOS or Ubuntu) — Installing on the Remaining Machines . . . . .	26
	Configuring the Control Node . . . . .	29
	Adding or Removing a Compute Node in an Existing Contrail Cluster . . . . .	34
<b>Chapter 4</b>	<b>Using Server Manager to Automate Provisioning . . . . .</b>	<b>37</b>
	Installing Server Manager . . . . .	37
	Installation Requirements for Server Manager . . . . .	37
	Platform Support . . . . .	37
	Installation Prerequisites . . . . .	38
	Installing Server Manager . . . . .	38
	Finishing the Provisioning . . . . .	39
	Starting the Server Manager Service . . . . .	39
	Upgrading Server Manager Software . . . . .	40
	Prerequisite to Upgrading . . . . .	40
	Steps for a New Installation . . . . .	40
	Server Manager Installation Completion Checks . . . . .	40
	Server Manager Checks . . . . .	40
	Server Manager Client Checks . . . . .	41
	Server Manager Webui Checks . . . . .	41
	Sample Configurations for Server Manager Templates . . . . .	41
	Sample Settings . . . . .	41
	The dhcp.template File . . . . .	41
	The named.conf.options File . . . . .	41
	The named.template File . . . . .	42
	The sendmail.cf File . . . . .	42
	Using Server Manager to Automate Provisioning . . . . .	42
	Overview of Server Manager . . . . .	42
	Server Manager Requirements and Assumptions . . . . .	43
	Server Manager Component Interactions . . . . .	44
	Configuring Server Manager . . . . .	45
	Configuring the Cobbler DHCP Template . . . . .	46
	User-Defined Tags for Server Manager . . . . .	47
	Server Manager Client Configuration File . . . . .	47
	Restart Services . . . . .	48
	Accessing Server Manager . . . . .	49
	Communicating with the Server Manager Client . . . . .	49
	Server Manager Commands for Configuring Servers . . . . .	50
	Create New Servers or Update Existing Servers . . . . .	50
	Delete Servers . . . . .	52
	Show Server Configuration . . . . .	52
	Server Manager Commands for Managing Clusters . . . . .	53
	Server Manager Commands for Managing Tags . . . . .	56
	Server Manager Commands for Managing Images . . . . .	58
	Server Manager Operational Commands for Managing Servers . . . . .	61
	Reimaging Server(s) . . . . .	61
	Provisioning and Configuring Roles on Servers . . . . .	63

Adding and Deleting Roles . . . . .	64
Restarting Server(s) . . . . .	66
Show Status of Server(s) . . . . .	67
Server Manager REST API Calls . . . . .	68
REST APIs for Server Manager Configuration Database Entries . . . . .	69
API: Add a Server . . . . .	69
API: Delete Servers . . . . .	70
API: Retrieve Server Configuration . . . . .	70
API: Add an Image . . . . .	70
API: Upload an Image . . . . .	71
API: Get Image Information . . . . .	71
API: Delete an Image . . . . .	71
API: Add or Modify a Cluster . . . . .	72
API: Delete a Cluster . . . . .	72
API: Get Cluster Configuration . . . . .	73
API: Get All Server Manager Configurations . . . . .	73
API: Reimage Servers . . . . .	73
API: Provision Servers . . . . .	73
API: Restart Servers . . . . .	74
Example: Reimaging and Provisioning a Server . . . . .	74
Using the Server Manager Web User Interface . . . . .	77
Log In to Server Manager . . . . .	77
Create a Cluster for Server Manager . . . . .	78
Working with Servers in the Server Manager User Interface . . . . .	84
Add a Server . . . . .	84
Edit Tags for Servers . . . . .	86
Using the Edit Config Option for Multiple Servers . . . . .	86
Filter Servers by Tag . . . . .	87
Viewing Server Details . . . . .	87
Configuring Images and Packages . . . . .	87
Add New Image or Package . . . . .	88
Selecting Server Manager Actions for Clusters . . . . .	88
Reimage a Cluster . . . . .	88
Provision a Cluster . . . . .	89
Installing and Using Server Manager Lite . . . . .	89
Server Manager Lite Overview . . . . .	89
Installing Server Manager Lite . . . . .	90
Provisioning Using SM-Lite . . . . .	90
Displaying the Cluster Status . . . . .	91
Displaying the SM-Lite Installation and Provisioning Log Files . . . . .	91
Contrail Provisioning Log Files . . . . .	91
<b>Chapter 5</b>	
<b>    Installing and Using Contrail Storage . . . . .</b>	<b>93</b>
Installing and Using Contrail Storage . . . . .	93
Overview of the Contrail Storage Solution . . . . .	93
Basic Storage Functionality with Contrail . . . . .	94
Ceph Block and Object Storage Functionality . . . . .	94
Using the Contrail Storage User Interface . . . . .	95
Hardware Specifications . . . . .	96

	Software Files for Compute Storage Nodes . . . . .	96
	Contrail OpenStack Nova Modifications . . . . .	96
	Installing the Contrail Storage Solution . . . . .	97
	Using Fabric Commands to Install and Configure Storage . . . . .	97
	Fabric Installation Procedure . . . . .	98
	Using Server Manager to Install and Configure Storage . . . . .	100
	Server Manager Installation Procedure for Storage . . . . .	100
	Example: Configurations for Storage for Reimaging and Provisioning a Server . . . . .	101
	Storage Installation Limits . . . . .	106
<b>Chapter 6</b>	<b>Upgrading Contrail Software . . . . .</b>	<b>109</b>
	Upgrading Contrail Software . . . . .	109
	DKMS for vRouter Kernel Module . . . . .	112
<b>Part 3</b>	<b>Configuring Contrail</b>	
<b>Chapter 7</b>	<b>Configuring Virtual Networks . . . . .</b>	<b>115</b>
	Creating Projects in OpenStack for Configuring Tenants in Contrail . . . . .	116
	Creating Virtual Networks and Policies in Juniper Networks Contrail . . . . .	117
	Creating a Virtual Network—Juniper Networks Contrail . . . . .	118
	Deleting a Virtual Network with Juniper Networks Contrail . . . . .	120
	Creating a Network Policy with Juniper Networks Contrail . . . . .	121
	Associating a Network to a Policy—Juniper Networks Contrail . . . . .	123
	Associating Network Policies Overview . . . . .	123
	Associating a Network Policy to a Network . . . . .	123
	Creating Virtual Networks and Policies in OpenStack Contrail . . . . .	125
	Creating a Virtual Network—OpenStack Contrail . . . . .	125
	Deleting a Virtual Network with OpenStack Contrail . . . . .	127
	Creating a Network Policy with OpenStack Contrail . . . . .	129
	Associating a Network to a Policy—OpenStack Contrail . . . . .	131
	Associating Network Policies Overview . . . . .	132
	Associating a Network Policy to a Network . . . . .	132
	Creating an Image and Launching a Virtual Machine . . . . .	133
	Creating an Image . . . . .	133
	Launching a Virtual Machine (Instance) . . . . .	136
	Creating a Floating IP Address Pool and Allocating it to a Virtual Machine . . . . .	138
	Creating a Floating IP Address Pool . . . . .	139
	Allocating a Floating IP Address to a Virtual Machine . . . . .	140
	Using Security Groups with Virtual Machines (Instances) . . . . .	142
	Security Groups Overview . . . . .	142
	Creating Security Groups and Adding Rules . . . . .	142
	Support for IPv6 Networks in Contrail . . . . .	145
	Overview: IPv6 Networks in Contrail . . . . .	146
	Creating IPv6 Virtual Networks in Contrail . . . . .	146
	Address Assignments . . . . .	146
	Adding IPv6 Peers . . . . .	147
	Configuring EVPN and VXLAN . . . . .	148
	Configuring the VXLAN Identifier Mode . . . . .	150
	Configuring Forwarding . . . . .	151

	Configuring the VXLAN Identifier . . . . .	153
	Configuring Encapsulation Methods . . . . .	154
<b>Chapter 8</b>	<b>Example of Deploying a Multi-Tier Web Application Using Contrail . . . . .</b>	<b>157</b>
	Example: Deploying a Multi-Tier Web Application . . . . .	157
	Multi-Tier Web Application Overview . . . . .	157
	Example: Setting Up Virtual Networks for a Simple Tiered Web Application . . . . .	158
	Verifying the Multi-Tier Web Application . . . . .	160
	Sample Addressing Scheme for Simple Tiered Web Application . . . . .	161
	Sample Physical Topology for Simple Tiered Web Application . . . . .	162
	Sample Physical Topology Addressing . . . . .	162
	Sample Network Configuration for Devices for Simple Tiered Web Application . . . . .	163
<b>Chapter 9</b>	<b>Configuring Services . . . . .</b>	<b>169</b>
	Configuring DNS Servers . . . . .	169
	DNS Overview . . . . .	169
	Defining Multiple Virtual Domain Name Servers . . . . .	170
	IPAM and Virtual DNS . . . . .	170
	DNS Record Types . . . . .	171
	Configuring DNS Using the Interface . . . . .	172
	Configuring DNS Using Scripts . . . . .	177
	Configuring Discovery Service . . . . .	178
	Contrail Discovery Service Introduction . . . . .	178
	Discovery Service Registration and Publishing . . . . .	179
	Discovery Service Subscription . . . . .	179
	Discovery Service REST API . . . . .	180
	Discovery Service Heartbeats . . . . .	182
	Discovery Service Internal Databases . . . . .	182
	Discovery Service Client Library . . . . .	182
	Discovery Service Debugging . . . . .	182
	Support for Multicast . . . . .	182
	Subnet Broadcast . . . . .	183
	All-Broadcast/Limited-Broadcast and Link-Local Multicast . . . . .	183
	Host Broadcast . . . . .	184
	Using Static Routes with Services . . . . .	184
	Static Routes for Service Instances . . . . .	184
	Configuring Static Routes on a Service Instance . . . . .	185
	Configuring Static Routes on Service Instance Interfaces . . . . .	186
	Configuring Static Routes as Host Routes . . . . .	188
	Configuring Metadata Service . . . . .	188

<b>Chapter 10</b>	<b>Configuring Service Chaining</b> . . . . .	<b>191</b>
	Service Chaining . . . . .	191
	Service Chaining Basics . . . . .	191
	Service Chaining Configuration Elements . . . . .	192
	Service Chaining MX Series Configuration . . . . .	195
	ECMP Load Balancing in the Service Chain . . . . .	196
	Customized Hash Field Selection for ECMP Load Balancing . . . . .	197
	Overview: Custom Hash Feature . . . . .	197
	Using ECMP Hash Fields Selection . . . . .	199
	Configuring ECMP Hash Fields Over Service Chains . . . . .	199
	Sample Flows . . . . .	199
	Sample Traffic Flow Path Without Custom ECMP Hash Fields . . . . .	200
	Sample Traffic Flow Path With Custom ECMP Hash Fields . . . . .	200
	Using the Contrail Heat Template . . . . .	201
	Introduction to Heat . . . . .	201
	Heat Architecture . . . . .	202
	Support for Heat Version 2 Resources . . . . .	202
	Deprecation of Heat Version 1 Resources . . . . .	202
	Heat Version 2 with Service Chaining and Port Tuple Sample Workflow . . . . .	203
	Example: Creating a Service Template Using Heat . . . . .	203
	Service Chain Route Reorigination . . . . .	204
	Overview: Service Chaining in Contrail . . . . .	205
	Route Aggregation . . . . .	206
	Schema for Route Aggregation . . . . .	207
	Configuring and Troubleshooting Route Aggregation . . . . .	208
	Routing Policy . . . . .	213
	Applying Routing Policy . . . . .	214
	Routing Policy Configuration . . . . .	216
	Configuring and Troubleshooting Routing Policy . . . . .	217
	Using a VNC Script to Create Routing Policy . . . . .	218
	Verify Routing Policy in API Server . . . . .	219
	Verify Routing Policy in the Control Node . . . . .	220
	Verify Routing Policy Configuration in the Control Node . . . . .	220
	Verify Routing Policy Configuration on the Routing Instance . . . . .	220
	Control for Route Reorigination . . . . .	221
	Configuring and Troubleshooting Reorigination Control . . . . .	222
	Service Instance Health Check . . . . .	223
	Health Check Overview . . . . .	223
	Health Check Object Configuration . . . . .	224
	Health Check Modes . . . . .	224
	Creating a Health Check with the Contrail User Interface . . . . .	225
	Using the Health Check . . . . .	226
	Health Check Process . . . . .	226
<b>Chapter 11</b>	<b>Examples: Configuring Service Chaining</b> . . . . .	<b>227</b>
	Example: Creating an In-Network or In-Network-NAT Service Chain . . . . .	227
	Creating an In-Network or In-Network-NAT Service Chain . . . . .	227
	Example: Creating a Transparent Service Chain . . . . .	235
	Creating a Transparent Mode Service Chain . . . . .	235



	Example: Creating a Service Chain With the CLI . . . . .	239
	CLI for Creating a Service Chain . . . . .	239
	CLI for Creating a Service Template . . . . .	240
	CLI for Creating a Service Instance . . . . .	240
	CLI for Creating a Service Policy . . . . .	240
	Example: Creating a Service Chain with VSRX and In-Network or Routed Mode . . . . .	241
<b>Part 4</b>	<b>Monitoring and Troubleshooting the Network Using Contrail Analytics</b>	
<b>Chapter 12</b>	<b>Understanding Contrail Analytics . . . . .</b>	<b>245</b>
	Contrail Analytics Overview . . . . .	245
	Contrail Alerts . . . . .	246
	Alert API Format . . . . .	246
	Analytics APIs for Alerts . . . . .	247
	Analytics APIs for SSE Streaming . . . . .	248
	Built-in Node Alerts . . . . .	248
	Underlay Overlay Mapping in Contrail . . . . .	249
	Overview: Underlay Overlay Mapping using Contrail Analytics . . . . .	250
	Underlay Overlay Analytics Available in Contrail . . . . .	250
	Architecture and Data Collection . . . . .	251
	New Processes/Services for Underlay Overlay Mapping . . . . .	251
	External Interfaces Configuration for Underlay Overlay Mapping . . . . .	252
	Physical Topology . . . . .	252
	SNMP Configuration . . . . .	252
	Link Layer Discovery Protocol (LLDP) Configuration . . . . .	253
	IPFIX and sFlow Configuration . . . . .	253
	Sending pRouter Information to the SNMP Collector in Contrail . . . . .	254
	pRouter UVEs . . . . .	255
	Contrail User Interface for Underlay Overlay Analytics . . . . .	256
	Viewing Topology to the Virtual Machine Level . . . . .	256
	Viewing the Traffic of any Link . . . . .	257
	Trace Flows . . . . .	257
	Search Flows and Map Flows . . . . .	258
	Overlay to Underlay Flow Map Schemas . . . . .	259
	Module Operations for Overlay Underlay Mapping . . . . .	262
	SNMP Collector Operation . . . . .	262
	Topology Module Operation . . . . .	263
	IPFIX and sFlow Collector Operation . . . . .	263
	Troubleshooting Underlay Overlay Mapping . . . . .	264
	Script to add pRouter Objects . . . . .	265

<b>Chapter 13</b>	<b>Configuring Contrail Analytics . . . . .</b>	<b>267</b>
	Analytics Scalability . . . . .	267
	High Availability for Analytics . . . . .	268
	System Log Receiver in Contrail Analytics . . . . .	269
	Overview . . . . .	269
	Redirecting System Logs to Contrail Collector . . . . .	269
	Exporting Logs from Contrail Analytics . . . . .	269
	Sending Flow Messages to the Contrail System Log . . . . .	270
	Ceilometer Support in a Contrail Cloud . . . . .	271
	Overview . . . . .	271
	Ceilometer Details . . . . .	271
	Verification of Ceilometer Operation . . . . .	272
	Contrail Ceilometer Plugin . . . . .	274
	Ceilometer Installation and Provisioning . . . . .	276
<b>Chapter 14</b>	<b>Using Contrail Analytics to Monitor and Troubleshoot the Network . . . . .</b>	<b>279</b>
	Monitoring the System . . . . .	280
	Debugging Processes Using the Contrail Introspect Feature . . . . .	282
	Monitor > Infrastructure > Dashboard . . . . .	286
	Monitor Dashboard . . . . .	286
	Monitor Individual Details from the Dashboard . . . . .	287
	Using Bubble Charts . . . . .	287
	Color-Coding of Bubble Charts . . . . .	288
	Monitor > Infrastructure > Control Nodes . . . . .	288
	Monitor Control Nodes Summary . . . . .	289
	Monitor Individual Control Node Details . . . . .	289
	Monitor Individual Control Node Console . . . . .	291
	Monitor Individual Control Node Peers . . . . .	293
	Monitor Individual Control Node Routes . . . . .	294
	Monitor > Infrastructure > Virtual Routers . . . . .	295
	Monitor vRouters Summary . . . . .	296
	Monitor Individual vRouters Tabs . . . . .	297
	Monitor Individual vRouter Details Tab . . . . .	297
	Monitor Individual vRouters Interfaces Tab . . . . .	298
	Configuring Interface Monitoring and Mirroring . . . . .	299
	Monitor Individual vRouters Networks Tab . . . . .	300
	Monitor Individual vRouters ACL Tab . . . . .	301
	Monitor Individual vRouters Flows Tab . . . . .	302
	Monitor Individual vRouters Routes Tab . . . . .	303
	Monitor Individual vRouter Console Tab . . . . .	304
	Monitor > Infrastructure > Analytics Nodes . . . . .	306
	Monitor Analytics Nodes . . . . .	306
	Monitor Analytics Individual Node Details Tab . . . . .	307
	Monitor Analytics Individual Node Generators Tab . . . . .	308
	Monitor Analytics Individual Node QE Queries Tab . . . . .	309
	Monitor Analytics Individual Node Console Tab . . . . .	310

Monitor > Infrastructure > Config Nodes . . . . .	311
Monitor Config Nodes . . . . .	311
Monitor Individual Config Node Details . . . . .	312
Monitor Individual Config Node Console . . . . .	313
Monitor > Networking . . . . .	314
Monitor > Networking Menu Options . . . . .	314
Monitor -> Networking -> Dashboard . . . . .	315
Monitor > Networking > Projects . . . . .	316
Monitor Projects Detail . . . . .	317
Monitor > Networking > Networks . . . . .	319
Query > Flows . . . . .	322
Query > Flows > Flow Series . . . . .	322
Example: Query Flow Series . . . . .	325
Query > Flow Records . . . . .	326
Query > Flows > Query Queue . . . . .	328
Query > Logs . . . . .	329
Query > Logs Menu Options . . . . .	329
Query > Logs > System Logs . . . . .	330
Sample Query for System Logs . . . . .	331
Query > Logs > Object Logs . . . . .	332
Example: Debugging Connectivity Using Monitoring for Troubleshooting . . . . .	334
Using Monitoring to Debug Connectivity . . . . .	334

## Part 5

## Index

Index . . . . .	341
-----------------	-----



# List of Figures

<b>Part 2</b>	<b>Installing and Upgrading Contrail</b>	
<b>Chapter 3</b>	<b>Installing Contrail and Provisioning Roles . . . . .</b>	<b>11</b>
	Figure 1: Configure > Infrastructure > BGP Routers . . . . .	30
	Figure 2: BGP Routers Summary . . . . .	30
	Figure 3: Create BGP Router . . . . .	31
	Figure 4: Control Nodes . . . . .	32
	Figure 5: Control Node Details . . . . .	33
	Figure 6: Control Node Peers Tab . . . . .	33
<b>Chapter 4</b>	<b>Using Server Manager to Automate Provisioning . . . . .</b>	<b>37</b>
	Figure 7: Server Manager Component Interactions . . . . .	44
<b>Part 3</b>	<b>Configuring Contrail</b>	
<b>Chapter 7</b>	<b>Configuring Virtual Networks . . . . .</b>	<b>115</b>
	Figure 8: OpenStack Projects . . . . .	116
	Figure 9: Add Project . . . . .	116
	Figure 10: Add IP Address Management . . . . .	118
	Figure 11: Configure Networks . . . . .	119
	Figure 12: Create Network . . . . .	119
	Figure 13: Configure Networks . . . . .	121
	Figure 14: Policies Window . . . . .	121
	Figure 15: Create Policy Window . . . . .	122
	Figure 16: Configure > Networking > Networks . . . . .	124
	Figure 17: Edit Network . . . . .	124
	Figure 18: Networks Window . . . . .	125
	Figure 19: Create Network Window . . . . .	125
	Figure 20: Create Network Window Subnet Tab . . . . .	126
	Figure 21: OpenStack Networks . . . . .	127
	Figure 22: OpenStack Network Detail, Associated Instances Tab . . . . .	127
	Figure 23: Instances . . . . .	128
	Figure 24: Network Policy . . . . .	129
	Figure 25: Create Network Policy . . . . .	129
	Figure 26: Network Policy . . . . .	130
	Figure 27: Edit Policy Rules . . . . .	130
	Figure 28: Networks Screen . . . . .	132
	Figure 29: Edit Network Policy . . . . .	132
	Figure 30: OpenStack Images Window . . . . .	133
	Figure 31: OpenStack Create An Image Window . . . . .	134
	Figure 32: OpenStack Instances . . . . .	136

	Figure 33: Launch Instance, Details Tab . . . . .	137
	Figure 34: Launch Instance, Networking Tab . . . . .	138
	Figure 35: Configure > Networking > Networks . . . . .	139
	Figure 36: Edit Network . . . . .	139
	Figure 37: Manage Floating IPs . . . . .	140
	Figure 38: Allocate Floating IP Window . . . . .	141
	Figure 39: Associate Floating IP . . . . .	142
	Figure 40: Security Groups . . . . .	143
	Figure 41: Edit Security Group Rules . . . . .	143
	Figure 42: Add Rule . . . . .	144
	Figure 43: Create Security Group . . . . .	145
	Figure 44: Associate Security Group at Launch Instance . . . . .	145
	Figure 45: Global Config Window for VXLAN ID . . . . .	150
	Figure 46: Edit Global Config Window for VXLAN Identifier Mode . . . . .	151
	Figure 47: Edit Network Window . . . . .	152
	Figure 48: Edit Network Window for VXLAN Identifier . . . . .	153
	Figure 49: Edit Global Config Window for Encapsulation Priority Order . . . . .	154
<b>Chapter 8</b>	<b>Example of Deploying a Multi-Tier Web Application Using Contrail . . . . .</b>	<b>157</b>
	Figure 50: Simple Tiered Web Use Case . . . . .	158
	Figure 51: Create Floating IP Pool . . . . .	159
	Figure 52: Allocate Floating IP . . . . .	159
	Figure 53: Sample Physical Topology for Simple Tiered Web Application . . . . .	162
	Figure 54: Sample Physical Topology Addressing . . . . .	163
<b>Chapter 9</b>	<b>Configuring Services . . . . .</b>	<b>169</b>
	Figure 55: DNS Servers Examples . . . . .	170
	Figure 56: IPAM and Virtual DNS . . . . .	171
	Figure 57: Example Usage for NS Record Type . . . . .	172
	Figure 58: Configure DNS Records . . . . .	172
	Figure 59: Add DNS . . . . .	173
	Figure 60: Add DNS Record . . . . .	174
	Figure 61: Associate IPAMs to DNS . . . . .	175
	Figure 62: Configure IP Address Management . . . . .	176
	Figure 63: DNS Server . . . . .	176
<b>Chapter 10</b>	<b>Configuring Service Chaining . . . . .</b>	<b>191</b>
	Figure 64: Service Chaining . . . . .	191
	Figure 65: Contrail Service Chain . . . . .	192
	Figure 66: Load Balancing a Service Chain . . . . .	196
	Figure 67: Route Reorigination . . . . .	205
	Figure 68: Route Aggregate Relationships . . . . .	206
	Figure 69: Create Route Aggregate . . . . .	209
	Figure 70: Create Service Instance . . . . .	210
	Figure 71: Create Health Check Screen . . . . .	225
<b>Chapter 11</b>	<b>Examples: Configuring Service Chaining . . . . .</b>	<b>227</b>
	Figure 72: Create Networks . . . . .	227
	Figure 73: Add Service Template . . . . .	228
	Figure 74: Add Service Template Shared IP . . . . .	230

	Figure 75: Service Templates . . . . .	230
	Figure 76: Create Service Instances . . . . .	231
	Figure 77: Create Service Instances . . . . .	232
	Figure 78: Service Instance Details . . . . .	232
	Figure 79: Service Instance Console . . . . .	233
	Figure 80: Create Policy . . . . .	233
	Figure 81: Edit Network . . . . .	234
	Figure 82: Launch Instances . . . . .	234
	Figure 83: Create Networks . . . . .	235
	Figure 84: Add Service Template . . . . .	236
	Figure 85: Create Service Instances . . . . .	237
	Figure 86: Service Instance Details . . . . .	238
	Figure 87: Create Policy . . . . .	238
	Figure 88: Launch Instances . . . . .	239
<b>Part 4</b>	<b>Monitoring and Troubleshooting the Network Using Contrail Analytics</b>	
<b>Chapter 12</b>	<b>Understanding Contrail Analytics . . . . .</b>	<b>245</b>
	Figure 89: Analytics Topology . . . . .	252
	Figure 90: Add Physical Router Window . . . . .	255
	Figure 91: Sample Output From a pRouter REST API . . . . .	255
	Figure 92: Sample Output From a pRouter UVE . . . . .	256
	Figure 93: Physical Topology Related to a vRouter . . . . .	257
	Figure 94: Traffic Statistics Graph . . . . .	257
	Figure 95: List of Active Flows . . . . .	258
	Figure 96: Underlay Path . . . . .	259
<b>Chapter 13</b>	<b>Configuring Contrail Analytics . . . . .</b>	<b>267</b>
	Figure 97: Analytics Scalability . . . . .	268
<b>Chapter 14</b>	<b>Using Contrail Analytics to Monitor and Troubleshoot the Network . . . . .</b>	<b>279</b>
	Figure 98: Monitor Menu . . . . .	280
	Figure 99: Control Nodes Details Tab Window . . . . .	283
	Figure 100: Controller Introspect Window . . . . .	284
	Figure 101: BGP Peer Introspect Page . . . . .	284
	Figure 102: BGP Neighbor Summary Introspect Page . . . . .	285
	Figure 103: Agent Introspect Page . . . . .	286
	Figure 104: Monitor > Infrastructure > Dashboard . . . . .	286
	Figure 105: Dashboard Summary Boxes . . . . .	287
	Figure 106: Bubble Summary Information . . . . .	288
	Figure 107: Control Nodes Summary . . . . .	289
	Figure 108: Individual Control Node—Details Tab . . . . .	290
	Figure 109: Individual Control Node—Console Tab . . . . .	291
	Figure 110: Individual Control Node—Peers Tab . . . . .	293
	Figure 111: Individual Control Node—Routes Tab . . . . .	294
	Figure 112: vRouters Summary . . . . .	296
	Figure 113: Individual vRouters—Details Tab . . . . .	297
	Figure 114: Individual vRouters—Interfaces Tab . . . . .	299
	Figure 115: Individual vRouter . . . . .	300

Figure 116: Interfaces . . . . .	300
Figure 117: Individual vRouters—Networks Tab . . . . .	301
Figure 118: Individual vRouters—ACL Tab . . . . .	302
Figure 119: Individual vRouters—Flows Tab . . . . .	303
Figure 120: Individual vRouters—Routes Tab . . . . .	304
Figure 121: Individual vRouter—Console Tab . . . . .	305
Figure 122: Analytics Nodes Summary . . . . .	307
Figure 123: Monitor Analytics Individual Node Details Tab . . . . .	308
Figure 124: Individual Analytics Node—Generators Tab . . . . .	309
Figure 125: Individual Analytics Node—QE QueriesTab . . . . .	309
Figure 126: Analytics Individual Node—Console Tab . . . . .	310
Figure 127: Config Nodes Summary . . . . .	312
Figure 128: Individual Config Nodes— Details Tab . . . . .	312
Figure 129: Individual Config Node—Console Tab . . . . .	313
Figure 130: Monitor Networking Menu Options . . . . .	315
Figure 131: Traffic Statistics for Domain Window . . . . .	315
Figure 132: Monitor > Networking > Projects . . . . .	316
Figure 133: Monitor Projects Connectivity Details . . . . .	317
Figure 134: Traffic Statistics Between Networks . . . . .	317
Figure 135: Projects Instances Summary . . . . .	318
Figure 136: Instance Traffic Statistics . . . . .	319
Figure 137: Network Summary . . . . .	319
Figure 138: Individual Network Connectivity Details—Summary Tab . . . . .	320
Figure 139: Individual Network— Port Map Tab . . . . .	320
Figure 140: Individual Network— Port Distribution Tab . . . . .	321
Figure 141: Individual Network Instances Tab . . . . .	321
Figure 142: Individual Network Details Tab . . . . .	322
Figure 143: Query Flow Series Window . . . . .	323
Figure 144: Flow Series Select . . . . .	324
Figure 145: Flow Series Filter . . . . .	325
Figure 146: Example: Query Flow Series . . . . .	325
Figure 147: Query Flow Series Tabular Results . . . . .	326
Figure 148: Query Flow Series Graphical Results . . . . .	326
Figure 149: Flow Records . . . . .	327
Figure 150: Flow Records Select Window . . . . .	328
Figure 151: Where Clause Window . . . . .	328
Figure 152: Flows Query Queue . . . . .	329
Figure 153: Query > Logs . . . . .	330
Figure 154: Query > Logs > System Logs . . . . .	330
Figure 155: Edit Where Clause . . . . .	332
Figure 156: Sample Query System Logs . . . . .	332
Figure 157: Query > Logs > Object Logs . . . . .	333
Figure 158: Navigate to Instance . . . . .	334
Figure 159: Traffic Statistics for Instance . . . . .	334
Figure 160: Navigate to Instance . . . . .	335
Figure 161: Traffic Statistics for Instance . . . . .	335
Figure 162: Navigate to a3s18 Interfaces . . . . .	335
Figure 163: Navigate to a3s19 Interfaces . . . . .	335
Figure 164: ACL Connectivity a3s18 . . . . .	336



Figure 165: ACL Connectivity a3s19 .....	336
Figure 166: Routes default-domain:demo:vn0:vn0 .....	336
Figure 167: Routes default-domain:demo:vn16:vn16 .....	337
Figure 168: Verify Route and Next Hop a3s18 .....	337
Figure 169: Verify Route and Next Hop a3s19 .....	338
Figure 170: Flows for a3s18 .....	338
Figure 171: Flows for a3s19 .....	338



# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xxi</b>
	Table 1: Notice Icons . . . . .	xxii
	Table 2: Text and Syntax Conventions . . . . .	xxii
<b>Part 2</b>	<b>Installing and Upgrading Contrail</b>	
<b>Chapter 3</b>	<b>Installing Contrail and Provisioning Roles</b> . . . . .	<b>11</b>
	Table 3: Create BGP Router Fields . . . . .	31
<b>Chapter 4</b>	<b>Using Server Manager to Automate Provisioning</b> . . . . .	<b>37</b>
	Table 4: Server Manager Parameters . . . . .	45
	Table 5: Server Manager Add Server Command Options . . . . .	50
	Table 6: Server Manager Delete Server Command Options . . . . .	52
	Table 7: Server Manager Show Server Command Options . . . . .	53
	Table 8: Server Manager Add Cluster Command Options . . . . .	54
	Table 9: Server Manager Delete Cluster Command Options . . . . .	55
	Table 10: Server Manager Show Cluster Command Options . . . . .	56
<b>Part 3</b>	<b>Configuring Contrail</b>	
<b>Chapter 7</b>	<b>Configuring Virtual Networks</b> . . . . .	<b>115</b>
	Table 11: Add IP Address Management Fields . . . . .	118
	Table 12: Create Network Fields . . . . .	119
	Table 13: Create Policy Fields . . . . .	122
	Table 14: Create Network Fields . . . . .	126
	Table 15: Edit Policy Rules Fields . . . . .	131
	Table 16: Create An Image Fields . . . . .	134
	Table 17: Launch Instance Details Tab Fields . . . . .	137
	Table 18: Add Rule Fields . . . . .	144
<b>Chapter 8</b>	<b>Example of Deploying a Multi-Tier Web Application Using Contrail</b> . . . . .	<b>157</b>
	Table 19: Sample Addressing Scheme for Example . . . . .	161
<b>Chapter 9</b>	<b>Configuring Services</b> . . . . .	<b>169</b>
	Table 20: DNS Record Types Supported . . . . .	171
	Table 21: Add DNS Fields . . . . .	173
	Table 22: Add DNS Record Fields . . . . .	174
	Table 23: Associate IPAMs to DNS Fields . . . . .	175
	Table 24: DNS Modes . . . . .	176
	Table 25: DNS Scripts . . . . .	177
<b>Chapter 10</b>	<b>Configuring Service Chaining</b> . . . . .	<b>191</b>

	Table 26: Health Check Configurable Parameters . . . . .	224
	Table 27: Create Health Check Fields . . . . .	225
<b>Chapter 11</b>	<b>Examples: Configuring Service Chaining . . . . .</b>	<b>227</b>
	Table 28: Add Service Template Fields . . . . .	228
	Table 29: Create Service Instances Fields . . . . .	231
	Table 30: Add Service Template Fields . . . . .	236
	Table 31: Create Service Instances Fields . . . . .	237
<b>Part 4</b>	<b>Monitoring and Troubleshooting the Network Using Contrail Analytics</b>	
<b>Chapter 14</b>	<b>Using Contrail Analytics to Monitor and Troubleshoot the Network . . . . .</b>	<b>279</b>
	Table 32: Monitor Menu Options . . . . .	280
	Table 33: Dashboard Summary Boxes . . . . .	287
	Table 34: Control Nodes Summary Fields . . . . .	289
	Table 35: Individual Control Node—Details Tab Fields . . . . .	290
	Table 36: Control Node: Console Tab Fields . . . . .	291
	Table 37: Control Node: Peers Tab Fields . . . . .	293
	Table 38: Control Node: Routes Tab Fields . . . . .	294
	Table 39: vRouters Summary Fields . . . . .	296
	Table 40: vRouters Details Tab Fields . . . . .	298
	Table 41: vRouters: Interfaces Tab Fields . . . . .	299
	Table 42: vRouters: Networks Tab Fields . . . . .	301
	Table 43: vRouters: ACL Tab Fields . . . . .	302
	Table 44: vRouters: Flows Tab Fields . . . . .	303
	Table 45: vRouters: Routes Tab Fields . . . . .	304
	Table 46: Control Node: Console Tab Fields . . . . .	305
	Table 47: Fields on Analytics Nodes Summary . . . . .	307
	Table 48: Monitor Analytics Individual Node Details Tab Fields . . . . .	308
	Table 49: Monitor Analytics Individual Node Generators Tab Fields . . . . .	309
	Table 50: Analytics Node QE Queries Tab Fields . . . . .	309
	Table 51: Monitor Analytics Individual Node Console Tab Fields . . . . .	310
	Table 52: Config Nodes Summary Fields . . . . .	312
	Table 53: Individual Config Nodes— Details Tab Fields . . . . .	313
	Table 54: Individual Config Node-Console Tab Fields . . . . .	313
	Table 55: Projects Summary Fields . . . . .	316
	Table 56: Projects Summary Fields . . . . .	316
	Table 57: Projects Instances Summary Fields . . . . .	318
	Table 58: Network Summary Fields . . . . .	319
	Table 59: Query Flow Series Fields . . . . .	323
	Table 60: Query Flow Records Fields . . . . .	327
	Table 61: Query Flow Records Fields . . . . .	329
	Table 62: Query System Logs Fields . . . . .	330
	Table 63: Object Logs Query Fields . . . . .	333

# About the Documentation

- Documentation and Release Notes on page xxi
- Documentation Conventions on page xxi
- Documentation Feedback on page xxiii
- Requesting Technical Support on page xxiv

## 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 <http://www.juniper.net/techpubs/>.

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 <http://www.juniper.net/books>.

## Documentation Conventions

---

Table 1 on page xxii defines notice icons used in this guide.

Table 1: 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 2 on page xxii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>
<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@# <b>set system domain-name</b> <i>domain-name</i>

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"><li>To configure a stub area, include the <b>stub</b> statement at the <b>[edit protocols ospf area area-id]</b> hierarchy level.</li><li>The console port is labeled <b>CONSOLE</b>.</li></ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub &lt;default-metric <i>metric</i>&gt;;</b>
(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.	<b>broadcast   multicast</b>  <b>(<i>string1</i>   <i>string2</i>   <i>string3</i>)</b>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [ <i>community-ids</i> ]</b>
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	<pre>[edit] routing-options {   static {     route default {       nexthop <i>address</i>;       retain;     }   } }</pre>
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"><li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li><li>To cancel the configuration, click <b>Cancel</b>.</li></ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

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

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.

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

## Requesting Technical Support

---

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

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

## Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

## Opening a Case with JTAC

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

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



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



## PART 1

# Overview

- [Understanding Contrail on page 3](#)



## CHAPTER 1

# Understanding Contrail

- [Contrail Overview on page 3](#)
- [Contrail Description on page 4](#)

### Contrail Overview

---

Juniper Networks Contrail is an open, standards-based software solution that delivers network virtualization and service automation for federated cloud networks. It provides self-service provisioning, improves network troubleshooting and diagnostics, and enables service chaining for dynamic application environments across enterprise virtual private cloud (VPC), managed Infrastructure as a Service (IaaS), and Networks Functions Virtualization use cases.

Contrail simplifies the creation and management of virtual networks to enable policy-based automation, greatly reducing the need for physical and operational infrastructure typically required to support network management. In addition, it uses mature technologies to address key challenges of large-scale managed environments, including multitenancy, network segmentation, network access control, and IP service enablement. These challenges are particularly difficult in evolving dynamic application environments such as the Web, gaming, big data, cloud, and the like.

Contrail allows a tenant or a cloud service provider to abstract virtual networks at a higher layer to eliminate device-level configuration and easily control and manage policies for tenant virtual networks. A browser-based user interface enables users to define virtual network and network service policies, then configure and interconnect networks simply by attaching policies. Contrail also extends native IP capabilities to the hosts (compute nodes) in the data center to address the scale, resiliency, and service enablement challenges of traditional orchestration platforms.

Using Contrail, a tenant can define, manage, and control the connectivity, services, and security policies of the virtual network. The tenant or other users can use the self-service graphical user interface to easily create virtual network nodes, add and remove IP services (such as firewall, load balancing, DNS, and the like) to their virtual networks, then connect the networks using traffic policies that are simple to create and apply. Once created, policies can be applied across multiple network nodes, changed, added, and deleted, all from a simple browser-based interface.

Contrail can be used with open cloud orchestration systems such as OpenStack or CloudStack. It can also interact with other systems and applications based on Operations

Support System (OSS) and Business Support Systems (BSS), using northbound APIs. Contrail allows customers to build elastic architectures that leverage the benefits of cloud computing — agility, self-service, efficiency, and flexibility — while providing an interoperable, scale-out control plane for network services within and across network domains.

**Related Documentation**

- [Contrail Description on page 4](#)

---

## Contrail Description

- [Contrail Major Components on page 4](#)
- [Contrail Solution on page 4](#)

### Contrail Major Components

The following are the major components of Contrail.

---

#### Contrail Control Nodes

- Responsible for the routing control plane, configuration management, analytics, and the user interface.
- Provide APIs to integrate with an orchestration system or a custom user interface.
- Horizontally scalable, can run on multiple servers.

---

#### Contrail Compute Nodes – XMPP Agent and vRouter

- Responsible for managing the data plane.
- Functionality can reside on a host OS.

### Contrail Solution

Contrail architecture takes advantage of the economics of cloud computing and simplifies the physical network (IP fabric) with a software virtual network overlay that delivers service orchestration, automation, and intercloud federation for public and hybrid clouds.

Similar to the native Layer 3 designs of web-scale players in the market and public cloud providers, the Contrail solution leverages IP as the abstraction between dynamic applications and networks, ensuring smooth migration from existing technologies, as well as support of emerging dynamic applications.

The Contrail solution is software running on x86 Linux servers, focused on enabling multitenancy for enterprise Information Technology as a Service (ITaaS). Multitenancy is enabled by the creation of multiple distinct Layer 3-enabled virtual networks with traffic isolation, routing between tenant groups, and network-based access control for each user group. To extend the IP network edge to the hosts and accommodate virtual machine workload mobility while simplifying and automating network (re)configuration, Contrail maintains a real-time state across dynamic virtual networks, exposes the network-as-a-service to cloud users, and enables deep network diagnostics and analytics down to the host.

In this paradigm, users of cloud-based services can take advantage of services and applications and assume that pooled, elastic resources are orchestrated, automated, and optimized across compute, storage, and network nodes in a converged architecture that is application-aware and independent of underlying hardware and software technologies.

- Related Documentation**
- [Contrail Overview on page 3](#)
  - [Installation Overview on page 11](#)





## PART 2

# Installing and Upgrading Contrail

- [Supported Platforms and Server Requirements on page 9](#)
- [Installing Contrail and Provisioning Roles on page 11](#)
- [Using Server Manager to Automate Provisioning on page 37](#)
- [Installing and Using Contrail Storage on page 93](#)
- [Upgrading Contrail Software on page 109](#)



## CHAPTER 2

# Supported Platforms and Server Requirements

- [Supported Platforms on page 9](#)
- [Server Requirements on page 9](#)

## Supported Platforms

---

Contrail Release 3.1 is supported on the OpenStack Kilo and Liberty releases, on the following operating system versions:

- Ubuntu 14.04.4
- CentOS 7.2
- Red Hat Enterprise Linux (RHEL) 7.2 (RHOSP8)
- VMware vCenter 5.5, 6.0

Contrail Release 3.1 is also supported on OpenStack Mitaka on Ubuntu 14.04.4.

Following is the supported Linux kernel version for each distribution supported on Contrail Release 3.1.

- CentOS 7.2—kernel version 3.10.0-327.10.1
- Ubuntu 14.04.4—kernel version 3.13.0-85-generic
- Red Hat 7.2—kernel version 3.10.0-327.10.1
- vCenter—Ubuntu 14.04.4 kernel version 3.13.0-85-generic

## Server Requirements

---

The minimum requirement for a proof-of-concept (POC) system is 3 servers, either physical or virtual machines. All non-compute roles can be configured in each controller node. For scalability and availability reasons, it is highly recommended to use physical servers.

Each server must have a minimum of:

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

For a production environment, each server must have a minimum of:

- 256 GB memory
- 500 GB hard drive
- 16 CPU cores



**NOTE:** If you are using Contrail Storage, additional hardware requirements can be found in [“Installing and Using Contrail Storage” on page 93](#), Hardware Specifications.

---

**Related  
Documentation**

- [Installation Overview on page 11](#)
- [Downloading Installation Software on page 12](#)

## CHAPTER 3

# Installing Contrail and Provisioning Roles

- [Installation Overview on page 11](#)
- [Downloading Installation Software on page 12](#)
- [Installing the Operating System and Contrail Packages on page 13](#)
- [Configuring System Settings on page 14](#)
- [Installing the Contrail Packages, Part One \(CentOS or Ubuntu\) on page 15](#)
- [Setting Up the Testbed Definitions File on page 16](#)
- [Testbed Definitions File Settings for Deploying Contrail with an Existing OpenStack Node on page 19](#)
- [Supporting Multiple Interfaces on Servers and Nodes on page 21](#)
- [Installing the Contrail Packages, Part Two \(CentOS or Ubuntu\) — Installing on the Remaining Machines on page 26](#)
- [Configuring the Control Node on page 29](#)
- [Adding or Removing a Compute Node in an Existing Contrail Cluster on page 34](#)

## Installation Overview

---

The Contrail Controller is typically installed on multiple servers. The base software image is installed on all servers to be used, then provisioning scripts are run that launch role-based components of the software.

The roles used for the installed system include:

- *cfgm*—Runs Contrail configuration manager (config-node)
- *openstack*—Runs OpenStack services such as Nova, Quantum, and the like
- *collector*—Runs monitoring and analytics services
- *compute*—Runs vRouter service and launches tenant virtual machines (VMs)
- *control*—Runs the control plane service
- *database*—Runs analytics and configuration database services
- *webui*—Runs the administrator web-based user interface service

The roles are run on multiple servers in an operating installation. A single node can have multiple roles. The roles can also run on a single server for testing or demonstration purposes.

[“Installing the Operating System and Contrail Packages” on page 13](#) describes installing the Contrail Controller software onto multiple servers.

Your account team can help you determine the number of servers needed for your specific implementation.

**Related  
Documentation**

- [Server Requirements on page 9](#)
- [Downloading Installation Software on page 12](#)
- [Installing the Operating System and Contrail Packages on page 13](#)
- [Installation Overview on page 11](#)
- [Downloading Installation Software on page 12](#)
- [Installing the Operating System and Contrail Packages on page 13](#)
- [Configuring System Settings on page 14](#)
- [Installing the Contrail Packages, Part One \(CentOS or Ubuntu\) on page 15](#)
- [Setting Up the Testbed Definitions File on page 16](#)
- [Testbed Definitions File Settings for Deploying Contrail with an Existing OpenStack Node on page 19](#)
- [Supporting Multiple Interfaces on Servers and Nodes on page 21](#)
- [Installing the Contrail Packages, Part Two \(CentOS or Ubuntu\) — Installing on the Remaining Machines on page 26](#)
- [Configuring the Control Node on page 29](#)
- [Adding or Removing a Compute Node in an Existing Contrail Cluster on page 34](#)

---

## Downloading Installation Software

All components necessary for installing the Contrail Controller are available as:

- an RPM file (**contrail-install-packages-1.xx-xxx.el6.noarch.rpm**) that can be used to install the Contrail system on an appropriate CentOS operating system.
- a Debian file (**contrail-install-packages-1.xx-xxx~xxxxxx\_all.deb**) that can be used to install the Contrail system on an appropriate Ubuntu operating system.

Versions are available for each Contrail release, for the supported Linux operating systems and versions, and for the supported versions of OpenStack.

All installation images can be downloaded from  
<http://www.juniper.net/support/downloads/?p=contrail#sw>.

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

**Related  
Documentation**

- [Installing the Operating System and Contrail Packages on page 13](#)
- [Configuring System Settings on page 14](#)
- [Setting Up the Testbed Definitions File on page 16](#)
- [Installing the Contrail Packages, Part One \(CentOS or Ubuntu\) on page 15](#)
- [Download Software](#)

---

## Installing the Operating System and Contrail Packages

Install the stock CentOS or Ubuntu operating system image appropriate for your version of Contrail onto the server. See “[Supported Platforms](#)” on [page 9](#). Then install Contrail packages separately.

The following are general guidelines for installing the operating system and preparing to install Contrail.

1. Install a CentOS or Ubuntu minimal distribution as desired on all servers. Follow the published operating system installation procedure for the selected operating system; refer to the website for the operating system.
2. After rebooting all of the servers after installation, verify that you can log in to each of them using the root password defined during installation.
3. After the initial installations on all servers, configure some items specific to your systems, (see “[Configuring System Settings](#)” on [page 14](#)), then begin the first part of the installation (see “[Installing the Contrail Packages, Part One \(CentOS or Ubuntu\)](#)” on [page 15](#)).

**Related  
Documentation**

- [Supported Platforms on page 9](#)

- [Installation Overview on page 11](#)
- [Downloading Installation Software on page 12](#)
- [Installing the Operating System and Contrail Packages on page 13](#)
- [Configuring System Settings on page 14](#)
- [Installing the Contrail Packages, Part One \(CentOS or Ubuntu\) on page 15](#)
- [Setting Up the Testbed Definitions File on page 16](#)
- [Supporting Multiple Interfaces on Servers and Nodes on page 21](#)
- [Installing the Contrail Packages, Part Two \(CentOS or Ubuntu\) — Installing on the Remaining Machines on page 26](#)
- [Configuring the Control Node on page 29](#)
- [Adding or Removing a Compute Node in an Existing Contrail Cluster on page 34](#)
- [Download Software](#)

---

## Configuring System Settings

After installing the base image on all servers being used in the installation, and before running role provisioning scripts, perform the following steps to configure items specific to your environment.

Perform these configuration steps each time you perform an initial installation or an upgrade to a new release.

To configure system settings:

1. Update the `/etc/resolv.conf` file with name server information specific to your system.
2. Update the `/etc/sysconfig/network` file with the hostname and domain information specific to your system.
3. Configure the LAN port with network information specific to your system:
  - a. Use the `show ifconfig -a` command to determine which LAN port you are using, as this might not be obvious on some systems due to the ways interfaces can be named.
  - b. Update the appropriate interface configuration file in `/etc/sysconfig/network-scripts/ifcfg-<int name>` using the following guidelines:
    - `IPADDR = <IP of the host you want to assign>`
    - `NETMASK = <e.g. 255.255.255.0>`
    - `GATEWAY = <gateway router address>`
    - `BOOTPROTO` — delete this, or change `dhcp` to `static`
    - Other settings can remain unchanged.



- Related Documentation**
- [Installing the Contrail Packages, Part One \(CentOS or Ubuntu\) on page 15](#)
  - [Setting Up the Testbed Definitions File on page 16](#)

## Installing the Contrail Packages, Part One (CentOS or Ubuntu)

This procedure includes instructions for installing Contrail for either a CentOS-based system or an Ubuntu-based system. In each step, be sure to follow the instructions for your operating system type.

All installation files are available from  
<http://www.juniper.net/support/downloads/?p=contrail#sw>.

**CentOS Systems** Contrail packages for CentOS are provided either as part of the Contrail ISO installation or separately in an RPM file with the format: **contrail-install-packages-x.xx-xxx~openstack\_version.el6.noarch.rpm**, where **x.xx-xxx~openstack\_version** represents the release number, build number, and OpenStack common version name for the included Contrail install packages.

If you already have a compatible operating system installed, you can choose to copy only the Contrail packages after the base operating system installation is complete. The base operating system can be installed using netboot or a USB, using installation instructions for that operating system.

**Ubuntu Systems** Contrail packages for Ubuntu are provided only as packages in a Debian file of the format: **contrail-install-packages-x.xx-xxx~openstack\_version\_all.deb**, where **x.xx-xxx~openstack\_version** represents the release number, build number, and OpenStack common version name for the included Contrail install packages.

It is expected that you already have a compatible Ubuntu operating system installed before installing the Contrail packages. For more information, see [“Supported Platforms” on page 9](#).



**NOTE:** If the stock kernel version of your Ubuntu system is older than the required version, the following Fabric task can be used to upgrade the kernel version in all nodes:

```
cd /opt/contrail/utils; fab upgrade_kernel_all
```

### *Installing Contrail Packages for CentOS or Ubuntu*

This procedure provides instructions for installing Contrail packages onto either a CentOS-based system or an Ubuntu-based system.

1. Ensure that a compatible base operating system has been installed, using the installation instructions for that system.
2. Download the appropriate Contrail install packages file from  
<http://www.juniper.net/support/downloads/?p=contrail#sw> :

**CentOS:** **contrail-install-packages-x.xx-xxx~openstack\_version.el6.noarch.rpm**

*Ubuntu:* `contrail-install-packages-x.xx-xxx~openstack_version_all.deb`

3. Copy the downloaded Contrail install packages file to `/tmp/` on the first server for your system installation.

4. On one of the config nodes in your cluster, copy the Contrail packages as follows:

*CentOS:* `scp`

```
<id@server>:/path/to/contrail-install-packages-x.xx-xxx~openstack_version.el6.noarch.rpm  
/tmp
```

*Ubuntu:* `scp`

```
<id@server>:/path/to/contrail-install-packages-x.xx-xxx~openstack_version_all.deb  
/tmp
```

5. Install the Contrail packages:

*CentOS:* `yum localinstall`

```
/tmp/contrail-install-packages-x.xx-xxx~openstack_version.el6.noarch.rpm
```

*Ubuntu:* `dpkg -i /tmp/contrail-install-packages-x.xx-xxx~openstack_version_all.deb`

6. Run the `setup.sh` script. This step creates the Contrail packages repository as well as the Fabric utilities (located in `/opt/contrail/utls`) needed for provisioning:

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

7. Populate the `testbed.py` definitions file, see [“Setting Up the Testbed Definitions File” on page 16](#).



**NOTE:** In Contrail, Apache ZooKeeper resides on the database node. Because a ZooKeeper ensemble operates most effectively with an odd number of nodes, it is required to have an odd number (3, 5, 7, and so on) of database nodes in a Contrail system.

---

#### Related Documentation

- [Download Software](#)
- [Supported Platforms on page 9](#)
- [Setting Up the Testbed Definitions File on page 16](#)
- [Supporting Multiple Interfaces on Servers and Nodes on page 21](#)
- [Installing the Contrail Packages, Part Two \(CentOS or Ubuntu\) — Installing on the Remaining Machines on page 26](#)
- [Configuring the Control Node on page 29](#)

---

## Setting Up the Testbed Definitions File

Populate a testbed definitions file, `/opt/contrail/utls/fabfile/testbeds/testbed.py`, with parameters specific to your system, then run the fab commands as provided in [“Installing](#)

the Contrail Packages, Part Two (CentOS or Ubuntu) — Installing on the Remaining Machines” on page 26 to launch the role-based provisioning script tasks.

You can view *example* testbed files on any node in the controller at:

- `/opt/contrail/utils/fabfile/testbeds/testbed_multibox_example.py` for a multiple server system

For a list of all available Fabric commands, refer to the `/opt/contrail/utils/README.fabric` file.

To define the following parameters within the `testbed.py` file:

1. Provide host strings for the nodes in the cluster. Replace the addresses shown in the example with the actual IP addresses of the hosts in your system.

```
host1 = 'root@1.1.1.1'
```

```
host2 = 'root@1.1.1.2'
```

```
host3 = 'root@1.1.1.3'
```

```
host4 = 'root@1.1.1.4'
```

```
host5 = 'root@1.1.1.5'
```

2. Define external routers (MX Series routers and the like) to which the virtual network controller control nodes are peered.

```
ext_routers = [('mx1', '1.1.1.253'), ('mx2', '1.1.1.252')]
```

If there are no external routers, define

```
ext_routers = []
```

3. Provide the BGP autonomous system number.

```
router_asn = 64512
```



**NOTE:** The default ASN 64512 is a private ASN number. A private ASN should be used if an AS is only required to communicate via BGP with a single provider. As the routing policy between the AS and the provider is not visible in the Internet, a private ASN can be used for this purpose. IANA has reserved AS 64512 through to AS 65535 to be used as private ASNs. If these circumstances do not apply, you cannot use the default or any other private ASN number.

4. Define the host on which the Fabric tasks are invoked. Replace the address shown in the example with the actual IP address of the host in your system.

```
host_build = 'user@10.10.10.10'
```

5. Define which hosts operate with which roles.

*For multinode setups:*

```
env.roledefs = {
    'all': [host1, host2, host3, host4, host5],
    'database': [host1, host2, host3],
    'cfgm': [host1, host2],
    'control': [host1, host2],
    'compute': [host4, host5],
    'collector': [host1, host2, host3],
    'webui': [host1],
    'build': [host_build],
}
```

*For single node all-in-one setups:*

```
env.roledefs = {
    'all': [host1],
    'database': [host1],
    'cfgm': [host1],
    'control': [host1],
    'compute': [host1],
    'collector': [host1],
    'webui': [host1],
    'build': [host_build],
}
```

#### 6. Define password credentials for each of the hosts.

```
env.password = 'secret' # Required only for releases prior to 1.10
env.passwords = {
    host1: 'secret',
    host2: 'secret',
    host3: 'secret',
    host4: 'secret',
    host5: 'secret',
}
```



**NOTE:** In releases earlier than Contrail Release 1.10, ensure that *both* the `env.password` and `env.passwords` variables are set.



**NOTE:** Set appropriate permissions for the `testbed.py` file because it contains host credentials.

If your system servers and nodes have multiple interfaces, refer to [“Supporting Multiple Interfaces on Servers and Nodes” on page 21](#) for information about setting up the `testbed.py` file for your system.

To deploy a Contrail High Available cluster, refer to *Juniper OpenStack High Availability* for information about setting up the `testbed.py` file for your system.

To deploy with an existing OpenStack, refer to [“Testbed Definitions File Settings for Deploying Contrail with an Existing OpenStack Node” on page 19](#) for `testbed.py` file definitions.

When you are finished, continue on to [“Installing the Contrail Packages, Part Two \(CentOS or Ubuntu\) — Installing on the Remaining Machines” on page 26](#).

- Related Documentation**
- [Supporting Multiple Interfaces on Servers and Nodes on page 21](#)
  - [Testbed Definitions File Settings for Deploying Contrail with an Existing OpenStack Node on page 19](#)
  - [Installing the Contrail Packages, Part Two \(CentOS or Ubuntu\) — Installing on the Remaining Machines on page 26](#)

---

## Testbed Definitions File Settings for Deploying Contrail with an Existing OpenStack Node

---

It is possible to deploy Contrail when there is already an existing OpenStack node on your system.

The following shows additional **testbed.py** definitions that are required to deploy Contrail when there is already an existing OpenStack node.

1. Update the OpenStack admin password in the **testbed.py**.

For example, use the following to update the OpenStack admin password.

```
env.openstack_admin_password = '<password>'
```

2. Update the Keystone environment section.

In the following example, the entries shown in the **env.keystone** section override the following previously-supported options:

- **service\_token**
- **keystone\_ip**
- **keystone\_admin\_user**
- **keystone\_admin\_password**
- **region\_name**

Options include the following:

- **keystone\_ip**—IP Address of the Keystone server. If using Openstack high availability (HA), provide the OpenStack VIP.
- **auth\_protocol**—The authentication protocol used by Keystone.
- **auth\_port**—The authentication port used by Keystone.
- **admin\_token**—The admin token of Keystone.
- **admin\_user**—The admin user name of Keystone.
- **admin\_password**—The password of the admin user of Keystone.
- **nova\_password**—The password of the Nova service.
- **neutron\_password**—The password of the Neutron networking service.

- **service\_tenant**—The tenant name of services such as Nova, Neutron, Glance, and so on.
- **admin\_tenant**—The name of the tenant admin user.
- **region\_name**—The OpenStack region to use. The default is RegionOne.
- **insecure**—The insecure option set for Keystone. The default is False.



**NOTE:** The option "insecure" is applicable only when the protocol is https.

- **manage\_neutron**—Option to configure a Neutron user or role in the Keystone server. The default = 'yes'.

The following is a sample configuration:

```
env.keystone = {
    'keystone_ip'      : '<ip address>',
    'auth_protocol'    : 'http',
    'auth_port'        : '35357',
    'admin_token'      : '$ABC123',
    'admin_user'       : 'admin',
    'admin_password'   : '<password>',
    'nova_password'    : '$ABC123',
    'neutron_password' : '$ABC123',
    'service_tenant'   : 'services',
    'admin_tenant'     : 'admin',
    'region_name'      : 'RegionOne',
    'insecure'         : 'False',
    'manage_neutron'   : 'no',
}
```

### 3. Update the OpenStack environment section.

Options include the following:

- **service\_token**—the tenant name of services such as Nova, Neutron, Glance, and so on.
- **amqp\_host**—the IP address of the Advanced Message Queuing Protocol (AMQP) server to be used in the OpenStack node. If using OpenStack high availability, provide the OpenStack VIP.
- **manage\_amqp**—Option to manage separate AMQP for OpenStack services in OpenStack nodes. The default = 'no'. If set to 'yes', AMQP is provisioned in OpenStack nodes, and OpenStack services use the AMQP in OpenStack nodes instead of config nodes. The **amqp\_host** is neglected if the **manage\_amqp** value is set.
- **osapi\_compute\_workers**—The default is 40. For a low memory system, reduce the osapi compute workers thread.
- **conductor\_workers**—The default is 40. For a low memory system, reduce the conductor workers thread.

The following is a sample configuration:

```
env.openstack = {
    'service_token'      : '$ABC123',
    'amqp_host'          : '<ip address>',
    'manage_amqp'        : 'no',
    'osapi_compute_workers' : 40,
    'conductor_workers'  : 40,
}
```

4. Update the configuration options for the Contrail config node in the OpenStack node.

Options include the following:

- **amqp\_hosts**—List of customer-deployed AMQP servers to be used by config services.
- **amqp\_port**—Port of the customer-deployed AMQP servers.

The following is a sample configuration:

```
env.cfgm = {
    'amqp_hosts' : ['<ip address>'],
    'amqp_port'  : '5672'
}
```

#### Related Documentation

- [Setting Up the Testbed Definitions File on page 16](#)
- [Supporting Multiple Interfaces on Servers and Nodes on page 21](#)
- [Installing the Contrail Packages, Part Two \(CentOS or Ubuntu\) — Installing on the Remaining Machines on page 26](#)

## Supporting Multiple Interfaces on Servers and Nodes

This section describes how to set up and manage multiple interfaces.

- [Support for Multiple Interfaces on page 21](#)
- [Server Interface Examples on page 23](#)
- [Interface Naming and Configuration Management on page 23](#)
- [Setting Up Interfaces and Installing on page 23](#)
- [Sample testbed.py File With Exclusive Interfaces on page 24](#)

### Support for Multiple Interfaces

Servers and nodes with multiple interfaces should be deployed with exclusive management and control and data networks. In the case of multiple interfaces per server, the expectation is that the management network provides only management connectivity to the cluster, and the control and data network carries the control plane information and the guest traffic data.

Examples of control traffic include the following:

- XMPP traffic between the control nodes and the compute nodes.
- BGP protocol messages across the control nodes.

- Statistics, monitoring, and health check data collected by the analytics engine from different parts of the system.

In Contrail Release 1.10 and later, control and data must share the same interface, configured in the **testbed.py** file in a section named **control\_data**.

---

### Number of cfm Nodes Supported

The Contrail system can have any number of **cfm** nodes.

---

### Uneven Number of Database Nodes Required

In Contrail Release 1.10 and later, Apache ZooKeeper resides on the database node. Because a ZooKeeper ensemble operates most effectively with an odd number of nodes, it is required to have an odd number (3, 5, 7, and so on) of database nodes in a Contrail system.

---

### Support for VLAN Interfaces

A VLAN ID can also be specified in the **testbed.py** file under the **control\_data** section, similar to the following example:

```
control_data= { host1: { 'ip': '<ip address>', 'gw': '<ip address>', 'device':  
    'bond0', 'vlan': '20'},  
                host2: { 'ip': '<ip address>', 'gw': '<ip  
    address>', 'device': 'bond0', 'vlan': '20'} }
```

---

### Support for Bonding Options

Contrail provides support for bond interface options.

The default bond interface options are:

```
miimon=100,mode=802.3ad(lacp),xmit_hash_policy=layer3+4
```

In the **testbed.py** bond section, anything other than name and member are treated as a bond interface option, and provisioned as such. The following is an example:

```
bond= { host1: { 'name': 'bond0', 'member': ['p2p0p2', 'p2p0p3'], 'lacp_rate': 'slow' }
```

---

### Support for Static Route Options

Contrail provides support for adding static routes on target systems. This option is ideal for use cases in which a system has servers with multiple interfaces and has control data or management connections that span multiple networks.

The following shows the use of the **static\_route** stanza in the **testbed.py** file to configure static routes in host2 and host5.

```
static_route = {  
    host2 : [{ 'ip': '<ip address>', 'netmask' : '<ip address>',  
              'gw': '<ip address>', 'intf': 'bond0' },  
            { 'ip': '<ip address>', 'netmask' : '<ip address>',
```



```

        'gw': '<ip address>', 'intf': 'bond0' }],
    host5 : [{ 'ip': '<ip address>', 'netmask' : '<ip address>',
               'gw': '<ip address>', 'intf': 'bond0'
    }],
}

```

## Server Interface Examples

In Contrail Release 1.10 and later, control and data are required to share the same interface. A set of servers can be deployed in any of the following combinations for management, control, and data:

- **mgmt=control=data** -- Single interface use case
- **mgmt, control=data** -- Exclusive management access, with control and data sharing a single network.

In Contrail, the following server interface combinations are not allowed:

- **mgmt=control, data**--Dual interfaces in Layer 3 mode, management and control shared on a single network
- **mgmt, control, data**—Complete exclusivity across management, control, and data traffic.

## Interface Naming and Configuration Management

On a standard Linux installation there is no guarantee that a physical interface will come up with the same name after a system reboot. Linux NetworkManager tries to accommodate this behavior by linking the interface configurations to the hardware addresses of the physical ports. However, Contrail avoids using hardware-based configuration files because this type of solution cannot scale when using remote provisioning and management techniques.

The Contrail alternative is a threefold interface-naming scheme based on **<bus, device, port (or function)>**. As an example, on a server operating system that typically assigns interface names such as **p4p0** and **p4p1** for onboard interfaces, the Contrail system assigns **p4p0p0** and **p4p0p1**, when using the optional **contrail-interface-name** package.

When the **contrail-interface-name** package is installed, it uses the threefold naming scheme to provide consistent interface naming after reboots. The **contrail-interface-name** package is installed by default when a Contrail ISO image is installed. If you are using an RPM-based installation, you should install the **contrail-interface-name** package before doing any network configuration.

If your system already has another mechanism for getting consistent interface names after a reboot, it is not necessary to install the **contrail-interface-name** package.

## Setting Up Interfaces and Installing

As part of the provisioning scheme, there are two additional commands that the administrator can use to set up control and data interfaces.

The **fab setup\_interface** command creates bond interface configurations, if there is a corresponding configuration in the **testbed.py** file (see the sample **testbed.py** file in [“Sample testbed.py File With Exclusive Interfaces”](#) on page 24).

When you use the **fab setup\_interface** command, the interface configurations are generated with the syntax (**ifcfg-\*** files), which is needed for the **network service**.

The **fab add\_static\_route** command creates static routes in a node, if there is a corresponding configuration in the **testbed.py** file (see the sample **testbed.py** file in [“Sample testbed.py File With Exclusive Interfaces”](#) on page 24).

The following is a typical work flow for setting up a cluster with multiple interfaces:

```
Set env.interface_rename = True in the testbed.py file (meaning: install the
contrail-interface-name package on compute nodes)

fab install_contrail (meaning: change the testbed.py file with the renamed interface
name)

fab setup_interface

fab add_static_route

fab setup_all
```



**NOTE:** The **fab setup\_interface** command and **fab add\_static\_route** command can be executed simultaneously by using the **fab setup\_network** command.

In cases where the **fab setup\_interface** command is not used for setting up the interfaces, configurations for the data interface are migrated as part of the **vrouter** installation on the compute nodes.

If the data interface is a bond interface, the bond member interfaces are reconfigured into network service based configurations using appropriate **ifcfg** script files.

## Sample testbed.py File With Exclusive Interfaces

The following is a sample **testbed.py** definitions file that shows the configuration for exclusive interfaces for management and control and for data networks.

```
#testbed file  from fabric.api import env
os_username = 'admin'
os_password = '<password>'
os_tenant_name = 'demo'

host1 = 'host@<ip address>'
host2 = 'host@<ip address>'
host3 = 'host@<ip address>'
host4 = 'host@<ip address>'
host5 = 'host@<ip address>'
host6 = 'host@<ip address>'
host7 = 'host@<ip address>'
host8 = 'host@<ip address>'
```

```

ext_routers = [('mx1', '<ip address>')] router_asn = <asn> public_vn_rtgt =
10003 public_vn_subnet = '<ip address>'

host_build = 'host@<ip address>'

env.roledefs = {
    'all': [host1, host2, host3, host4, host5, host6, host7, host8],
    'cfgm': [host1],
    'openstack': [host6],
    'webui': [host7],
    'control': [host4, host3],
    'compute': [host2, host5],
    'collector': [host2, host3],
    'database': [host8],
    'build': [host_build],
}

env.hostnames = {
    'all': ['nodea10', 'nodea4', 'nodea2', 'nodeb2',
'nodeb12', 'nodea32', 'nodec36', 'nodec31']
}

bond= {
    host2 : { 'name': 'bond0', 'member': ['p2p0p0', 'p2p0p1', 'p2p0p2', 'p2p0p3'],
'mode': 'balance-xor' },
    host5 : { 'name': 'bond0', 'member': ['p4p0p0', 'p4p0p1', 'p4p0p2', 'p4p0p3'],
'mode': 'balance-xor' }, }

control_data = {
    host1 : { 'ip': '<routing prefix address>', 'gw' : '<ip address>',
'device': 'eth0' },
    host2 : { 'ip': '<routing prefix address>', 'gw' : '<ip address>',
'device': 'p0p25p0' },
    host3 : { 'ip': '<routing prefix address>', 'gw' : '<ip address>',
'device': 'eth0' },
    host4 : { 'ip': '<routing prefix address>', 'gw' : '<ip address>',
'device': 'eth3' },
    host5 : { 'ip': '<routing prefix address>', 'gw' : '<ip address>',
'device': 'p6p0p1' },
    host6 : { 'ip': '<routing prefix address>', 'gw' : '<ip address>',
'device': 'eth0' },
    host7 : { 'ip': '<routing prefix address>', 'gw' : '<ip address>',
'device': 'eth1' },
    host8 : { 'ip': '<routing prefix address>', 'gw' : '<ip address>',
'device': 'eth1' }, }

env.password = 'secret' #Required only for releases prior to 1.10

env.passwords = {
    host1: 'secret',
    host2: 'secret',
    host3: 'secret',
    host4: 'secret',
    host5: 'secret',
    host6: 'secret',
    host7: 'secret',
    host8: 'secret',

    host_build: 'secret'
}

```

Related Documentation • [Juniper OpenStack High Availability](#)

## Installing the Contrail Packages, Part Two (CentOS or Ubuntu) — Installing on the Remaining Machines

### Preinstallation Checklist



**NOTE:** This procedure assumes that you have first completed the following procedures:

- [Installing the Contrail Packages, Part One \(CentOS or Ubuntu\) on page 15](#)
- [Setting Up the Testbed Definitions File on page 16](#)

And the following system tasks are accomplished:

- All of the servers are time synced.
- All servers can ping from one to another, both on management and on data and control, if part of the system.
- All servers can **ssh** and **scp** between one another.
- All host names are resolvable.
- If using CentOS or RHEL, SELinux has been disabled (**/etc/sysconfig/selinux**).

Each step in this procedure contains instructions for installing on a CentOS system or an Ubuntu system. Be sure to follow the instructions specific to your operating system.

To copy and install Contrail packages on the remaining machines in your cluster, you can use **scp** and **yum localinstall** as on the first server (for CentOS) or **scp** and **dpkg -i** (for Ubuntu), as in ["Installing the Contrail Packages, Part One \(CentOS or Ubuntu\)" on page 15](#), or you can use a Fabric utility to copy onto all machines at once, as follows:

1. Ensure that the **testbed.py** file has been created and populated with information specific to your cluster at **/opt/contrail/utils/fabfile/testbeds**.

See ["Setting Up the Testbed Definitions File" on page 16](#).

2. Run Fabric commands to install packages as follows:

**CentOS:** **/opt/contrail/utils/fab**

**install\_pkg\_all:/tmp/contrail-install-packages-1.xx-xxx~openstack\_version.el6.noarch.rpm**

**Ubuntu:** **/opt/contrail/utils/fab**

**install\_pkg\_all:/tmp/contrail-install-packages-1.xx-xxx~openstack\_version\_all.deb**



**NOTE:** Fab commands are always run from **/opt/contrail/utils/**.

3. *Ubuntu*: The recommended Kernel version for Ubuntu based system is 3.13.0-40. Nodes can be upgraded to kernel version 3.13.0-40 using below fabric-utils command :

```
fab upgrade_kernel_all
```



**NOTE:** This step upgrades the kernel version to 3.13.0-40 in all nodes and performs reboot. Reconnect to perform remaining tasks.

4. Install the required Contrail packages in each node of the cluster:

```
fab install_contrail
```



**NOTE:** To install Contrail with an existing OpenStack node:

```
fab install_without_openstack # Script will install nova-compute in the compute node
```

or

```
fab install_without_openstack:no # User installs nova-compute in the compute node
```

5. If your installation has multiple interfaces (see [“Supporting Multiple Interfaces on Servers and Nodes” on page 21](#)), run **setup\_interface**:

```
fab setup_interface
```

6. Provision the entire cluster:

```
fab setup_all
```



**NOTE:** To provision Contrail with an existing OpenStack node, use one of the following:

- **fab setup\_without\_openstack** # Script provisions vrouter and nova-compute services in the compute nodes and the compute nodes are rebooted on completion
- **fab setup\_without\_openstack:no** # Only vrouter services are provisioned, the nova-compute service is not provisioned and compute nodes are rebooted on completion
- **fab setup\_without\_openstack:no,False** # Only vrouter services are provisioned, the nova-compute service is not provisioned and the compute nodes are not rebooted on complet

7. *CentOS only*: Alternatively, if the **contrail-install-packages** is already installed (as part of installing the Contrail ISO via **netboot**), follow steps 2, 4, and 5.

When finished, you can proceed to [“Configuring the Control Node” on page 29](#).

- Related Documentation**
- [Configuring the Control Node on page 29](#)

## Configuring the Control Node

---

An important task after a successful installation is to configure the control node. 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.

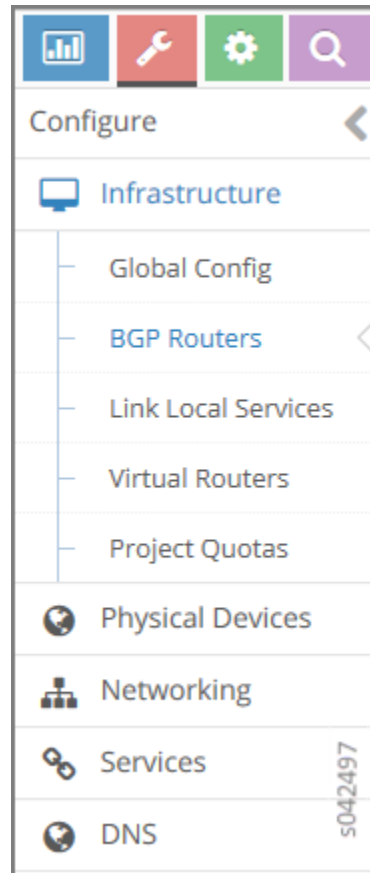
Before you begin, ensure that the following tasks are completed:

- 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 can access the Contrail user interface at **<http://nn.nn.nn.nn:8080>**, where ***nn.nn.nn.nn*** is the IP address of the configuration node server that is running the **contrail-webui** service.

To configure BGP peering in the control node:

1. From the Contrail Controller module control node (<http://nn.nn.nn.nn:8080>), select **Configure > Infrastructure > BGP Routers**; see [Figure 1 on page 30](#).

Figure 1: Configure > Infrastructure > BGP Routers



A summary screen of the control nodes and BGP routers is displayed; see [Figure 2 on page 30](#).

Figure 2: BGP Routers Summary

Configure > Infrastructure > BGP Routers				
BGP Routers				
<input type="checkbox"/> IP Address	Type	Vendor	HostName	
▶ <input type="checkbox"/> 10.84.25.31	Control Node	contrail	b5s31	⚙
▶ <input type="checkbox"/> 10.84.11.252	BGP Router	mx	a3-mx80-1	⚙
▶ <input type="checkbox"/> 10.84.25.30	Control Node	contrail	b5s30	⚙
▶ <input type="checkbox"/> 10.84.25.29	Control Node	contrail	b5s29	⚙
▶ <input type="checkbox"/> 10.84.25.28	Control Node	contrail	b5s28	⚙
▶ <input type="checkbox"/> 10.84.25.27	Control Node	contrail	b5s27	⚙
▶ <input type="checkbox"/> 10.84.11.253	BGP Router	mx	mx1	⚙
Total: 7 records   50 Records ▼				
Page 1 of 1				

2. (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.

3. 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 3 on page 31](#).

**Figure 3: 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 3 on page 31](#).

**Table 3: Create BGP Router Fields**

Field	Description
<b>Hostname</b>	Enter a name for the node being added.
<b>Vendor ID</b>	Required for external peers. Populate with a text identifier, for example, "MX-0". (BGP peer only)
<b>IP Address</b>	The IP address of the node.
<b>Router ID</b>	Enter the router ID.
<b>Autonomous System</b>	Enter the AS number for the node. (BGP peer only)
<b>Address Families</b>	Enter the address family, for example, <b>inet-vpn</b>

Table 3: Create BGP Router Fields (*continued*)

Field	Description
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.

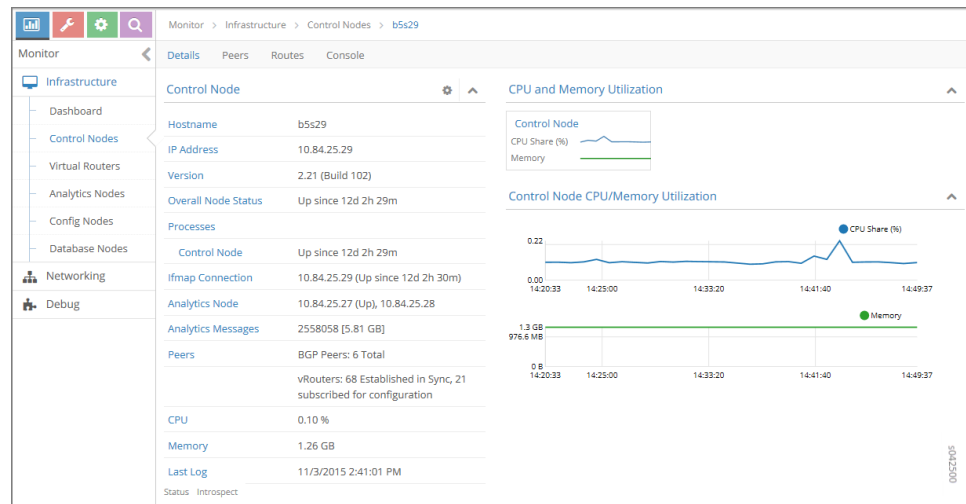
- Click **Save** to add each node that you create.
- To configure an existing node as a peer, select it from the list in the **Available Peers** box, then click >> to move it into the **Configured Peers** box.  
Click << to remove a node from the **Configured Peers** box.
- You can check for peers by selecting **Monitor > Infrastructure > Control Nodes**; see [Figure 4 on page 32](#).

Figure 4: Control Nodes



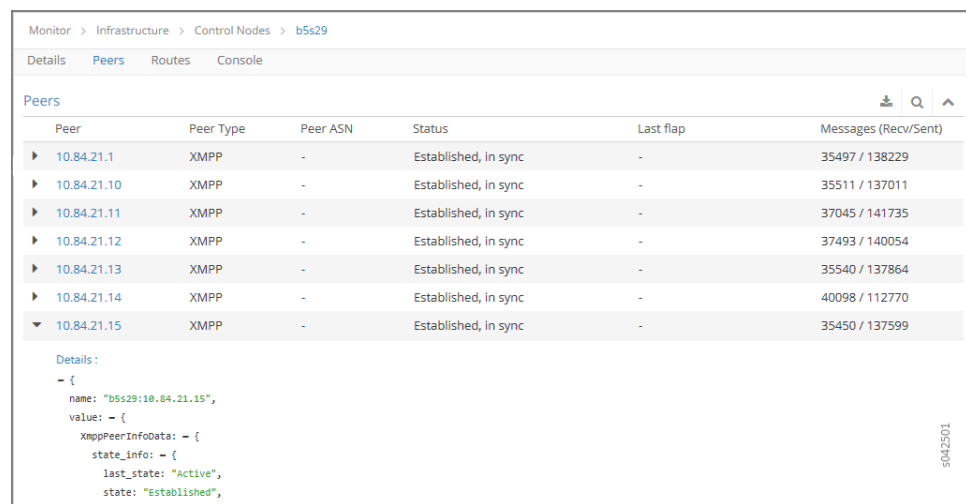
In the **Control Nodes** window, click any hostname in the memory map to view its details; see [Figure 5 on page 33](#).

Figure 5: Control Node Details



8. Click the **Peers** tab to view the peers of a control node; see [Figure 6 on page 33](#).

Figure 6: Control Node Peers Tab



#### Related Documentation

- [Creating Virtual Networks and Policies in Juniper Networks Contrail on page 117](#)
- [Creating Virtual Networks and Policies in OpenStack Contrail on page 125](#)

## Adding or Removing a Compute Node in an Existing Contrail Cluster

---

Use the following procedure to add one or more new compute nodes to an existing Contrail cluster.

1. Add the new information about the new compute node(s) into your existing **testbed.py** file.



**NOTE:** For convenience, this procedure assumes you are adding a node @10.1.1.1, however, replace the 10.1.1.1 with the correct IP for the node or nodes that you are adding.

2. Copy the **contrail-install-packages** file for CentOS or Ubuntu to the **/tmp** directory of the **cfgm** node where the **fab** commands are triggered:

CentOS: **scp <id@server>:/path/to/contrail-install-packages-xxx-xxx.el6.noarch.rpm /tmp**

Ubuntu: **scp <id@server>:/path/to/contrail-install-packages\_xxx-xxx~havana\_all.deb /tmp**

3. Install the **contrail-install-packages** on to the new compute node (or nodes):

*CentOS:* **fab**

**install\_pkg\_node:/tmp/contrail-install-packages\_x.xx-xxx.xxx.noarch.rpm,host@10.1.1.1**

*Ubuntu:* **fab**

**install\_pkg\_node:/tmp/contrail-install-packages\_x.xx-xxx~havana\_all.deb,host@10.1.1.1**

4. For Ubuntu 12.04.4 or 12.04.3 server with a kernel version older than 3.13.0-34, upgrade the kernel by using the following **fab** command:

**cd /opt/contrail/utils; fab upgrade\_kernel\_node:host@10.1.1.1**

where 10.1.1.1 must be replaced with the server's actual IP address.

5. Use **fab** commands to add the new compute node (or nodes):

**fab add\_vrouter\_node:host@10.1.1.1**

6. Manually reboot the compute node after kernel upgrade.

### *Removing a Node*

Use the following procedure to remove one or more compute nodes from an existing Contrail cluster.



**NOTE:** For convenience, this procedure assumes you are adding a node @10.1.1.1, however, replace the 10.1.1.1 with the correct IP address for the node or nodes that you are adding.

1. Use the following **fab** command to remove the new compute node:

```
fab detach_vrouter_node:root@10.1.1.1
```

2. Remove the information about this detached compute node from the existing **testbed.py** file.



## CHAPTER 4

# Using Server Manager to Automate Provisioning

- [Installing Server Manager on page 37](#)
- [Using Server Manager to Automate Provisioning on page 42](#)
- [Using the Server Manager Web User Interface on page 77](#)
- [Installing and Using Server Manager Lite on page 89](#)

## Installing Server Manager

---

- [Installation Requirements for Server Manager on page 37](#)
- [Installing Server Manager on page 38](#)
- [Upgrading Server Manager Software on page 40](#)
- [Server Manager Installation Completion Checks on page 40](#)
- [Sample Configurations for Server Manager Templates on page 41](#)

## Installation Requirements for Server Manager

This document provides details for installing Server Manager.

### Platform Support

---

Server Manager can be installed on, and used to reimage and provision, the following platform operating systems:

- Ubuntu 14.04.4
- Ubuntu 14.04.2
- Ubuntu 14.04.1
- Ubuntu 14.04
- Ubuntu 12.04.3

As of Contrail 3.0, Server Manager installation supports Contrail provisioning for only the following OpenStack versions:

- Juno

- Kilo

Additionally, Server Manager can be used to reimage and provision:

- VMware ESXi 5.5

### Installation Prerequisites

---

Before installing Server Manager ensure the following prerequisites are met.

- The system has Internet access to get dependent packages. Ensure access is available to the Ubuntu archive **mirrors/repos** at **/etc/apt/sources.list**.



**NOTE:** Server Manager is tested only with the following versions of dependent packages: Puppet 3.7.3-1 and Cobbler 2.6.3-1. The tested versions are installed during the Server Manager installation.

---

- Puppet Master requires the fully-qualified domain name (FQDN) of the Server Manager for key generation. The domain name is taken from the **/etc/hosts** file. If the server is part of multiple domains, specify the domain name by using the **–domain** option during the installation.
- On multi-interface systems, specify the interface to which Server Manager needs to listen by using the **–hostip** option. If the listening interface is not specified, the first available interface from the **ifconfig** list is used.

## Installing Server Manager

Server Manager and all of its components (Server Manager, monitoring, Server Manager client, Server Manager Web user interface) are provided together in a wrapper installation package:

Ubuntu: **contrail-server-manager-installer\_<version~sku>.deb**

You can choose to install all components at once or install individual components one at a time.

Use the following steps to install and set up Server Manager and its components.

1. Install the Server Manager packages:

Ubuntu: **dpkg –i contrail-server-manager–installer\_<version-sku>.deb**

---



**NOTE:** Make sure to select the correct version package that corresponds to the platform for which you are installing.

---

2. Set up the Server Manager components. Use the **setup.sh** command to install all of the components, or you can install individual components.



```
cd /opt/contrail/contrail_server_manager; ./setup.sh [--hostip=<ip address>]
[--domain=<domain name>]
```

- To set up all components:

```
./setup.sh --all
```

- To set up only the Server Manager server:

```
./setup.sh --sm=contrail-server-manager_<version-sku>.deb
```

- To set up only the Server Manager client:

```
setup.sh --sm-client=contrail-server-manager_<version-sku>.deb
```

- To set up only the Server Manager user interface:

```
setup.sh --webui=contrail-server-manager_<version-sku>.deb
```

- To set up only Server Manager monitoring:

```
setup.sh --sm-mon=contrail-server-manager_<version-sku>.deb
```

Other options include:

- --sm-cliff-client
- --nowebui
- --nosm-mon

3. Installation logs are located at `/var/log/contrail/install_logs/`.

### Finishing the Provisioning

The Server Manager service does not start automatically upon successful installation. You must finish the provisioning by modifying the following templates. Refer to the sample configuration section included in this topic for details about configuring these files.

```
/etc/cobbler/dhcp.template
/etc/cobbler/named.template
/etc/bind/named.conf.options
/etc/cobbler/setting
/etc/cobbler/modules.conf
/etc/mail/sendmail.cf
```

### Starting the Server Manager Service

When you are finished modifying the templates to match your environment, start the Server Manager service using the following command:

```
service contrail-server-manager start
```

## Upgrading Server Manager Software

If you are upgrading Server Manager software from a previous version to the current version, use the following guidelines to ensure successful installation.

### Prerequisite to Upgrading

---

Before upgrading, you must remove the previous version of the Server Manager installer. Remove any existing Server Manager installer package from the system using the following steps.

1. **dpkg -P contrail-server-manager-installer**
2. **rm -rf /opt/contrail/contrail-server-manager**

### Steps for a New Installation

---

After the existing Server Manager installer package has been removed, use the installation steps for a new installation, see details in previous section:

1. **dpkg -i <contrail-server-manager-installer\*.deb>**
2. **cd /opt/contrail/contrail-server-manager**
3. **./setup.sh -all**
4. After the setup script has completed running, you can restart Server Manager by issuing:

**service contrail-server-manager restart**

It is not necessary to reconfigure the templates of DHCP, bind, and so on. Previous template configurations and configured data are preserved during the upgrade.

## Server Manager Installation Completion Checks

The following are various checks you can use to investigate the status of your Server Manager installation.

### Server Manager Checks

---

Use the following to check that the Server Manager installation is complete.

- Use the following commands to verify that the services are running:

**service contrail-server-manager status**

**service cobblerd status**

**service bind9 status**

**service isc-dhcp-server status**

- Also verify processes using the following command:

**ps auwx | grep Passenger**

### Server Manager Client Checks

---

- Verify the items listed using the following command:  
**which server-manager**
- Check the client configuration at  
**/opt/contrail/contrail-server-manager/client/sm-client-config.ini**
- Make sure **listen\_ip\_addr** is configured with the Server Manager IP address.

### Server Manager Webui Checks

---

- Verify the status of the Server Manager webui using the following command:  
**service supervisor-webui status**
- Check the webui access from the browser:
  - Contrail release 2.2 and lower—**http:<server manager> :8080**
  - Contrail release 3.0 and greater—**http:<server manager> :9080.**

## Sample Configurations for Server Manager Templates

The following are sample parameters for the Server Manager templates. Use settings specific for your environment. Typically you configure to parameters for DHCP, bind, and e-mail services.

### Sample Settings

---

```
bind_master: 10.84.11.6

manage_forward_zones: ['contrail.juniper.net']

manage_reverse_zones: ['10.84.11']

next_server: 10.84.11.6

server: 10.84.11.6
```

### The dhcp.template File

---

Add Server Manager hooks into the DHCP template. When DHCP commit, release, or expire actions occur, the Server Manager is notified. The DHCP servers are detected on the Server Manager and the *Discovered* status is maintained.

Define subnet blocks that the DHCP server needs to support, using the sample format given in the **/etc/cobbler/dhcp.template** file.

### The named.conf.options File

---

You must configure the following:

```
forwarders {
    0.0.0.0;
```

```
};  
  
allow-query { any;};  
  
recursion yes;
```

---

### The `named.template` File

Include the following in the beginning of the `named.template` file:

```
"/etc/bind/named.conf.options";
```

```
"/etc/bind/named.conf.local";
```

---

### The `sendmail.cf` File

The `sendmail.cf` template is present with a `juniper.net` configuration. Populate it with configuration specific to your environment. The Server Manager uses the template to generate e-mails when reimaging or provisioning is completed.

#### Related Documentation

- [Using Server Manager to Automate Provisioning on page 42](#)
- [Using the Server Manager Web User Interface on page 77](#)

---

## Using Server Manager to Automate Provisioning

- [Overview of Server Manager on page 42](#)
- [Server Manager Requirements and Assumptions on page 43](#)
- [Server Manager Component Interactions on page 44](#)
- [Configuring Server Manager on page 45](#)
- [Configuring the Cobbler DHCP Template on page 46](#)
- [User-Defined Tags for Server Manager on page 47](#)
- [Server Manager Client Configuration File on page 47](#)
- [Restart Services on page 48](#)
- [Accessing Server Manager on page 49](#)
- [Communicating with the Server Manager Client on page 49](#)
- [Server Manager Commands for Configuring Servers on page 50](#)
- [Server Manager REST API Calls on page 68](#)
- [Example: Reimaging and Provisioning a Server on page 74](#)

## Overview of Server Manager

The Contrail Server Manager can be used to provision, configure, and reimage a Contrail virtual network system of servers, clusters, and nodes. Server Manager is an alternative to using Fabric commands to provision a Contrail system.

This section describes the functions and usage guidelines for the Contrail Server Manager.

The Server Manager provides a simple, centralized way for users to manage and configure components of a virtual network system running across multiple physical and virtual servers in a cloud infrastructure.

You can use Server Manager to configure, provision, and reimage servers with the correct software version and packages for the nodes that are running on each server in multiple virtual network system clusters.

The Server Manager:

- Provides REST APIs to handle customer requests.
- Manages its own database to store information about the servers.
- Interacts with other open source products such as Cobbler and Puppet to configure servers based on user requests.

## Server Manager Requirements and Assumptions

The following are requirements and assumptions for the Server Manager:

- The Server Manager runs on a Linux server (bare metal or virtual machine) and assumes availability of several software products with which it interacts to provide the functionality of managing servers.
- The Server Manager has network connectivity to the servers it is trying to manage.
- The Server Manager has access to a remote power management tool to power cycle the servers that it manages.
- The Server Manager uses Cobbler software for Linux provisioning to configure and download software to physical servers. Cobbler resides on the same server that is running the Server Manager daemon.
  - Server Manager assumes that DNS and DHCP servers embedded with Cobbler provide IP addresses and names to the servers being managed, although it is possible to use external DNS and DHCP servers.
- The Server Manager uses Puppet software, an open source configuration management tool, to accomplish the configuration management of target servers, including the installation and configuration of different software packages and the launching of various services.
- SQLite3 database management software is used to maintain and manage server configurations and it runs on the same machine where the Server Manager daemon is running.


## Server Manager Component Interactions

The Server Manager runs as a daemon and provides REST APIs for interaction with the client. The Server Manager accepts user input in the form of REST API requests, performs the requested function on the resources, and responds with a REST API response.

Configuration parameters required by the Server Manager are provided in the Server Manager configuration file, however, the parameters can be overridden by Server Manager command line parameters.

Figure 7 on page 44 illustrates several high-level components with which the Server Manager interacts.

**Figure 7: Server Manager Component Interactions**



Internally, the Server Manager uses a SQLite3 database to hold server configuration information. The Server Manager coordinates the database configuration information and user requests to manage the servers defined in the database.

While managing the servers, the Server Manager also communicates with other software components. It uses Cobbler for reimaging target servers and it uses Puppet for provisioning, thereby ensuring necessary software packages are installed and configured, required services are running, and so on.

A Server Manager agent runs on each of the servers and communicates with the Server Manager, providing the information needed to monitor the operation of the servers. The Server Manager agent also uses REST APIs to communicate with the Server Manager, and it can use other software tools to fetch other information, such as Nagios infrastructure monitoring, Razor spam filtering, Intelligent Platform Interface (IPMI), and so on. Monitoring functionality is not part of Server Manager.

## Configuring Server Manager

When the installation of all Server Manager components and dependent packages is finished, configure the Server Manager with parameters that identify your environment and make it available for clients to serve REST API requests.

Upon installation, a sample Server Manager configuration file is created at:

`/opt/contrail/server_manager/sm-config.ini`

Modify the **sm-config.ini** configuration file to include parameter values specific to your environment.

The environment-specific configuration section of the **sm-config.ini** file is named **SERVER-MANAGER**.

The following example shows the format and parameters of the **SERVER-MANAGER** section. Typically, only the **listen\_ip\_addr**, **cobbler\_username**, and **cobbler\_passwd** values need to be modified.

```
[SERVER-MANAGER]

listen_ip_addr = <IP-Address-of-SM>

listen_port    = <port-number>

database_name  = <database-file-name>

server_manager_base_dir = <base-dir-where-SM-files-are-created>

html_root_dir  = <html-root-dir>

cobbler_ip_address = <cobbler-ip-address>

cobbler_port    = <cobbler-port-number>

cobbler_username = <cobbler-username>

cobbler_password = <cobbler-password>

puppet_dir     = <puppet-directory>

ipmi_username   = <IPMI username>

ipmi_password   = <IPMI password>

ipmi_type       = <IPMI type>
```

[Table 4 on page 45](#) provides details for each of the parameters in the **SERVER-MANAGER** section.

**Table 4: Server Manager Parameters**

Parameter	Configuration
<b>listen_ip_addr</b>	Specify the IP address of the server on which the Server Manager is listening for REST API requests.

Table 4: Server Manager Parameters (*continued*)

Parameter	Configuration
<code>listen_port</code>	The port number on which the Server Manager is listening for REST API requests. The default is 9001.
<code>database_name</code>	The name of the database file where the Server Manager stores configuration information. This file is created under <code>server_manager_base_dir</code> .
<code>server_manager_base_dir</code>	The base directory where all of the Server Manager configuration files are created. The default is <code>/etc/contrail_smgr</code> .
<code>html_root_dir</code>	The HTML root directory ( <code>/var/www/html</code> ).
<code>cobbler_ip_address</code>	The IP address used to access Cobbler. This address <b>MUST</b> be the same address as the <code>listen_ip_address</code> . The Server Manager assumes that the Cobbler service is running on the same server as the Server Manager service.
<code>cobbler_port</code>	The port on which Cobbler listens for user requests. Leave this field blank.
<code>cobbler_username</code>	Specify the user name to access the Cobbler service. Specify <b>testing</b> unless your Cobbler settings have been modified to use a different user name.
<code>cobbler_password</code>	Specify the password to access the Cobbler service. Specify <b>testing</b> unless your Cobbler settings have been modified to use a different password.
<code>puppet_dir</code>	The directory where the Puppet manifests and templates are created. This is <code>/etc/puppet</code> , unless your Puppet configuration has been modified to use another directory.
<code>ipmi_username</code>	The IPMI username for power management.
<code>ipmi_password</code>	The IPMI password for power management.
<code>ipmi_type</code>	The IPMI type (ipmilan, or other cobbler supported types).

## Configuring the Cobbler DHCP Template

In addition to configuring the `server_config.ini` file, you must manually change the settings in the `/etc/cobbler/dhcp.template` file to use the correct subnet address, mask, and DNS domain name for your environment. Optionally, you can also restrict the use of the current instance of Server Manager and Cobbler to a subset of servers in the network.

Below is a snippet from the `dhcp.template` file showing the fields to be modified.



**NOTE:** The IP addresses and other values in the following are shown for example purposes only. Be sure to use values that are correct for your environment.

```
subnet <subnet address> netmask <netmask address> {
```



```

option routers                <ip-address>;
option subnet-mask            <subnet-address>;
option domain-name-servers    $next_server, 8.8.8.8 ;
option domain-search          "<domain name1>t", "<domain name 2>";
option domain-name            "<domain name>" ;
option ntp-servers            $next_server;
default-lease-time            21600;
max-lease-time                43200;
next-server                   $next_server;
filename                      "/pxelinux.0";
}

```

## User-Defined Tags for Server Manager

Server Manager allows you to define tags that can be used to group servers for performing a particular operation such as show information, reimage, provision, and so on. The Server Manager supports up to seven different tags that can be configured and used for grouping servers.

The names of user-defined tags are kept in the **tags.ini** file, at **/etc/contrail\_smgr/tags.ini**.

It is possible to modify tag names, and add or remove tags dynamically using the Server Manager REST API interface. However, if a tag is already being used to group servers, the tag must be removed from the servers before tag modification is allowed.

The following is a sample **tags.ini** file that is copied on installation. In the sample file, five tags are defined – **datacenter**, **floor**, **hall**, **rack**, and **user\_tag**. Use the tags to group servers together.

```

[TAGS]
tag1 = datacenter
tag2 = floor
tag3 = hall
tag4 = rack
tag5 = user_tag

```

## Server Manager Client Configuration File

The Server Manager client application installation copies the **/opt/contrail/server\_manager/client/sm-client-config.ini**, sample configuration file. The sample file contains parameter values such as the IP address to reach the Server Manager, the port used by Server Manager, default values for cluster table entries, default values for server table entries, and so on. You must modify the values in the **sm-client-config.ini** file to match your environment.

Use values from the **CLUSTER** and **SERVER** subsections in the **ini** file during creation of new clusters and servers, unless you explicitly override the values at the time of creation.

The **CLUSTER** and **SERVER** subsections have fields that represent the password for a host or a service. If a value for the password field is not explicitly provided, the Server Manager selects a default password.

Starting with Contrail Release 3.0.2, if you don't explicitly specify a password, a password is automatically generated by the system. This makes the clusters provisioned by Server Manager more secure.

The following fields now autogenerate a password whenever an explicit password is not provided.

- Previous Parameters:
  - mysql\_root\_password
  - keystone\_password
  - heat\_encryption\_key
- New Parameters:
  - mysql:root\_password
  - mysql:service\_password
  - keystone:admin\_password
  - heat\_encryption\_key

The following is a sample client configuration file.

```
[SERVER-MANAGER]
; ip address of the server manager
; replace the following with proper server manager address
listen_ip_addr = 10.xx.xx.xx
; server manager listening port
listen_port    = 9001
[CLUSTER]
subnet_mask = <subnet address>
domain = <domain name>
database_dir = /home/cassandra
encapsulation_priority = MPLSoUDP,MPLSoGRE,VXLAN
router_asn = 64512
keystone_username = admin
keystone_password = <password>
password = <password>
analytics_data_ttl = 168
haproxy = disable
use_certificates = False
multi_tenancy = False
database_token =
service_token = <password>
analytics_data_ttl = 168
[SERVER]
```

## Restart Services

When all user changes have been made to the configuration files, restart the Server Manager by issuing the following command so that it runs with the modifications :

## service contrail-server-manager restart

### Accessing Server Manager

When the Server Manager configuration has been customized to your environment, and the required daemon services are running, clients can request and use services of the Server Manager by using REST APIs. Any standard REST API client can be used to construct and send REST API requests and process Server Manager responses.

The following steps are typically required to fully implement a new cluster of servers being managed by the Server Manager.

1. Configure elements such as servers, clusters, and images in the Server Manager.

Before managing servers, the Server Manager needs to have configuration details of the servers that the Server Manager is managing.

2. Specify the name and location of boot images, packages, and repositories used to bring up the servers with needed software.

Currently, the servers can be imaged with CentOS, Ubuntu Linux, or VMWare ESXi distributions.

3. Provision or configure the servers by installing necessary packages, creating configuration files, and bringing up the correct services so that each server can perform the functions or role(s) configured for that server.

A Contrail system of servers has several components or roles that work together to provide the functionality of the virtual network system, including: control, config, analytics, compute, web-ui, OpenStack, and database. Each of the roles has different requirements for the software and services needed. The provisioning REST API enables the client to configure the roles on servers using the Server Manager.

4. Set up API calls for monitoring servers.

Once the servers in the Contrail system are correctly reimaged and provisioned to run configured roles, the server monitoring REST API calls allow clients to monitor performance of the servers as they provide one or more role functions. Monitoring functionality is not available in Server Manager for Contrail Release 1.10.

### Communicating with the Server Manager Client

Server Manager provides a REST API interface for clients to talk to the Server Manager software. Any client that can send and receive REST API requests and responses can be used to communicate with Server Manager, for example, Curl or Postman. Additionally, the Server Manager software provides a client with a simplified CLI interface, in a separate package. The Server Manager client can be installed and run on the Server Manager machine itself or on another server with an IP connection to the Server Manager machine.

Prior to using the Server Manager client CLI commands, you need to modify the **sm-client-config.ini** file to specify the IP address and the port for the Server Manager, along with other parameters.

Each of the commands described in this section takes a set of parameters you specify, constructs a REST API request to the Server Manager, and provides the server's response.

The following describes each Server Manager client CLI command in detail.

## Server Manager Commands for Configuring Servers

This section describes commands that are used to configure servers and server parameters in the Server Manager database. These commands allow you to add, modify, delete, or view servers.

- [Create New Servers or Update Existing Servers on page 50](#)
- [Delete Servers on page 52](#)
- [Show Server Configuration on page 52](#)
- [Server Manager Commands for Managing Clusters on page 53](#)
- [Server Manager Commands for Managing Tags on page 56](#)
- [Server Manager Commands for Managing Images on page 58](#)
- [Server Manager Operational Commands for Managing Servers on page 61](#)
- [Reimaging Server\(s\) on page 61](#)
- [Provisioning and Configuring Roles on Servers on page 63](#)
- [Adding and Deleting Roles on page 64](#)
- [Restarting Server\(s\) on page 66](#)
- [Show Status of Server\(s\) on page 67](#)

### Create New Servers or Update Existing Servers

Use the **server-manager add** command to create a new server or update a server in the Server Manager database.

```
server-manager [-h] [--config_file CONFIG_FILE] server [-f FILE_NAME]
```

[Table 5 on page 50](#) lists the optional arguments.

**Table 5: Server Manager Add Server Command Options**

Option	Description
<b>-h, --help</b>	Show the options available for the current command and exit.
<b>--config_file CONFIG_FILE, -c CONFIG_FILE</b>	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code>
<b>--file_name FILE_NAME, -f FILE_NAME</b>	The JSON file that contains the server parameter values.

If no JSON file is specified, the client program accepts all the needed server parameter values interactively, then builds a JSON file and makes a REST API call to the Server Manager. The JSON file contains a number of server entries, in the format shown in the following example:

```
{
  "server": [
    {
      "id": "demo2-server",
      "mac_address": "<mac address>",
      "ip_address": "<ip address>",
      "parameters" : {
        "interface_name": "eth1",
        "partition": ""
      },
      "roles" :
["config","openstack","control","compute","collector","webui","database"],
      "cluster_id": "demo-cluster",
      "subnet_mask": "<subnet mask address>",
      "gateway": "<ip address>",
      "password": "<password>",
      "domain": "<domain name>",
      "ipmi_address": "<ipmi address>",
      "tag" : {
        "datacenter" : "demo-dc",
        "floor" : "demo-floor",
        "hall" : "demo-hall",
        "rack" : "demo-rack",
        "user_tag" : "demo-user"
      }
    }
  ]
}
```

Most of the parameters in the JSON sample file are self-explanatory. **Cluster\_id** defines the cluster to which the server belongs. The **interface\_name** is the Ethernet interface name on the server used to configure the server and **roles** define the roles that can be configured for the server. The sample **roles** array in the example lists all valid role values. **Tag** defines the list of tag values for grouping and classifying the server.

The **server-manager add** command adds a new entry if the server with the given ID or `mac_address` does not exist in the Server Manager database. If an entry already exists, the add command modifies the fields in the existing entry with any new parameters specified.

### Delete Servers

Use the **server-manager delete** command to delete one or more servers from the Server Manager database.

Table 6 on page 52 lists the optional arguments.

**Table 6: Server Manager Delete Server Command Options**

Option	Description
<code>-h, --help</code>	Show the options available for the current command and exit.
<code>--config_file CONFIG_FILE, -c CONFIG_FILE</code>	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code>
<code>--server_id SERVER_ID</code>	The server ID for the server or servers to be deleted.
<code>--mac MAC</code>	The MAC address for the server or servers to be deleted.
<code>--ip IP</code>	The IP address for the server or servers to be deleted.
<code>--cluster_id CLUSTER_ID</code>	The cluster ID for the server or servers to be deleted.
<code>--tag TagName=TagValue</code>	The TagName that is to be matched with the Tagvalue. Up to seven TagName and Tagvalue pairs separated by commas can be provided.

The criteria for identifying servers to be deleted can be specified by providing the **server\_id** or the server: **mac address**, **ip**, **cluster\_id**, or the **TagName = TagValue**.

Provide one of the server matching criteria to display a list of servers available to be deleted.

### Show Server Configuration

Use the **server-manager show** command to display the configuration of servers from the Server Manager database.

```
server-manager show [--config_file CONFIG_FILE]
                    server (--server_id SERVER_ID | --mac MAC | --ip IP |
--cluster_id CLUSTER_ID | --tag <tag_name=tag_value>.. ) [--detail]
```

Table 7 on page 53 lists the optional arguments.

**Table 7: Server Manager Show Server Command Options**

Option	Description
-h, --help	Show the options available for the current command and exit.
--config_file CONFIG_FILE, -c CONFIG_FILE	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code>
--server_id SERVER_ID	The server ID for the server or servers to be deleted.
--mac MAC	The MAC address for the server or servers to be displayed.
--ip IP	The IP address for the server or servers to be displayed.
--cluster_id CLUSTER_ID	The cluster ID for the server or servers to be displayed.
--tag TagName=TagValue	The TagName that is to be matched with the Tagvalue. Up to seven TagName and Tagvalue pairs separated by commas can be provided.
--detail, -d	Flag to indicate if details are requested.

The criteria for identifying servers to be displayed can be specified by providing the **server\_id** or the server: **mac address**, **ip**, **VNS\_id**, **cluster\_id**, or **TagName=TagValue**.

Provide one or more of the server matching criteria to display a list of servers.

### Server Manager Commands for Managing Clusters

A cluster is used to store parameter values that are common to all servers belonging to that cluster. The commands in this section facilitate managing clusters in the Server Manager database, enabling you to add, modify, delete, and view clusters.



**NOTE:** Whenever a server is created with a specific **cluster\_id**, Server Manager checks to see if a cluster with that ID has already been created. If there is no matching **cluster\_id** already in the database, an error is returned.

- [Create a New Cluster or Update an Existing Cluster on page 54](#)
- [Delete a Cluster on page 55](#)
- [Show Cluster Configuration on page 55](#)

**Create a New Cluster or Update an Existing Cluster**

Use the **server-manager add** command to create a new cluster or update an existing cluster in the Server Manager database.

```
server-manager [-h] [--config_file CONFIG_FILE]
               cluster [--file_name FILE_NAME]
```

Table 8 on page 54 lists the optional arguments.

**Table 8: Server Manager Add Cluster Command Options**

Option	Description
<b>-h, --help</b>	Show the options available for the current command and exit.
<b>--config_file CONFIG_FILE, -c CONFIG_FILE</b>	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code>
<b>--file_name FILE_NAME, -f FILE_NAME</b>	The JSON file that contains the cluster parameter values.

If no JSON file is specified, the client program accepts all the needed cluster parameter values interactively, then builds a JSON file and makes a REST API call to the Server Manager. The JSON file contains a number of cluster entries, in the format shown in the following example:

```
{
  "cluster" : [
    {
      "id" : "demo-cluster",
      "parameters" : {
        "router_asn": "<asn number>",
        "database_dir": "/home/cassandra",
        "database_token": "",
        "use_certificates": "False",
        "multi_tenancy": "False",
        "encapsulation_priority": "MPLSoUDP,MPLSoGRE,VXLAN",
        "service_token": "<password>",
        "keystone_username": "admin",
        "keystone_password": "<password>",
        "keystone_tenant": "admin",
        "analytics_data_ttl": "168",
```



```

        "haproxy": "disable",
        "subnet_mask": "<subnet mask address>",
        "gateway": "<ip address>",
        "password": "<password>",
        "external_bgp": "",
        "domain": "<domain name>"
    }
}
]
}

```

Server membership to a cluster is determined by specifying the ID corresponding to the cluster when defining the server. All of the cluster parameters are available to the server when provisioning roles on the server.

#### **Delete a Cluster**

Use the **server-manager delete** command to delete a cluster from the Server Manager database that are no longer needed. Use this command after all servers in the cluster have been deleted.

```

server-manager delete [-h] [--config_file CONFIG_FILE]

cluster [--cluster_id CLUSTER_ID]

```

[Table 9 on page 55](#) lists the optional arguments.

**Table 9: Server Manager Delete Cluster Command Options**

Option	Description
<b>-h, --help</b>	Show the options available for the current command and exit.
<b>--config_file CONFIG_FILE, -c CONFIG_FILE</b>	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code>

#### **Show Cluster Configuration**

Use the **server-manager show** command to list the configuration of a cluster.

```

server-manager show [-h] [ --config_file CONFIG_FILE]

cluster [--cluster_id CLUSTER_ID] [--detail]

```

[Table 10 on page 56](#) lists the optional arguments.

Table 10: Server Manager Show Cluster Command Options

Option	Description
<code>-h, --help</code>	Show the options available for the current command and exit.
<code>--config_file CONFIG_FILE, -c CONFIG_FILE</code>	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code>
<code>--detail, -d</code>	Flag to indicate if details are requested.
<code>--cluster_id CLUSTER_ID</code>	The cluster ID for the cluster or clusters.

You can optionally specify a cluster ID to get information about a particular cluster. If the optional parameter is not specified, information about all clusters in the system is returned.

### Server Manager Commands for Managing Tags

Tags are used for grouping servers together so that an operation such as get, reimage, provision, status, and so on can be easily performed on servers that have matching tags. The Server Manager provides a flexible way for you to define your own tags, then use those tags to assign values to servers. Servers with matching tag values can be easily grouped together. The Server Manager can store a maximum of seven tag values. At initialization, the Server Manager reads the tag names from the configuration file. The tag names can be retrieved or modified using CLI commands. When modifying tag names, the Server Manager ensures that the tag name being modified is not used by any of the server entries.

- [Create a New Tag or Update an Existing Tag on page 56](#)
- [Show Tag Configuration on page 57](#)

#### Create a New Tag or Update an Existing Tag

Use the **server-manager add** command to create a new tag or update an existing tag in the Server Manager database.

```
server-manager add [-h] [--config_file CONFIG_FILE]
                  tag [--file_name FILE_NAME]
```

Optional arguments include the following:

Option	Description
<code>-h, --help</code>	Show the options available for the current command and exit.
<code>--config_file CONFIG_FILE, -c CONFIG_FILE</code>	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code>
<code>--file_name FILE_NAME, -f FILE_NAME</code>	The JSON file that contains the tag names.

If no JSON file is specified, the client program prompts you for tag names, then builds a JSON file and makes a REST API call to the Server Manager. The JSON file contains a number of tag entries, in the format shown in the following example:

```
{
  "tag1" : "data-center",
  "tag2" : "floor",
  "tag3" : "",
  "tag4" : "pod",
  "tag5" : "rack",
}
```

In the example, you specify a JSON file to add or modify the tags, tag1 thru tag5. For tag3, the "" value specifies that if the tag is defined prior to the CLI command, it is removed on execution of the command. The tag name for tag1 is set to data-center. This is allowed if, and only if, none of the server entries are using tag1.

#### **Show Tag Configuration**

Use the **server-manager show** command to list the configuration of a tag.

```
server-manager show [-h] [ --config_file CONFIG_FILE] tag
```

Optional arguments include the following:

Option	Description
<b>-h, --help</b>	Show the options available for the current command and exit.
<b>--config_file CONFIG_FILE, -c CONFIG_FILE</b>	The name of the Server Manager client configuration file. The default file is: <b>/opt/contrail/server_manager/client/sm-client-config.ini</b>

The following is sample output for the **show tag** command.

```
{
  "tag1": "datacenter",
  "tag2": "floor",
  "tag3": "hall",
  "tag4": "rack",
  "tag5": "user_tag"
}
```

## Server Manager Commands for Managing Images

In addition to servers and clusters, the Server Manager also manages information about images and packages that can be used to reimage and configure servers. Images and packages are both stored in the database as images. When new images are added to the database, or existing images are modified or deleted, the Server Manager interfaces with Cobbler to make corresponding modifications in the Cobbler distribution profile for the specified image.

The image types supported are summarized in the following table:

Image Type	Description
<b>centos</b>	Manages the CentOS stock ISO, and does not include the Contrail packages repository packaged with the ISO.
<b>contrail-centos-package</b>	Maintains a repository of the package to be installed on the CentOS system image.
<b>ubuntu</b>	Manages the base Ubuntu ISO.
<b>contrail-ubuntu-package</b>	Maintains a repository of packages that contain Contrail and dependent packages to be installed on an Ubuntu base system.
<b>ESXi5.1/ESXi5.5</b>	Manages VMware ESXi 5.1 or 5.5 ISO.

- [Creating New Images or Updating Existing Images on page 58](#)
- [Add an Image on page 58](#)
- [Upload an Image on page 60](#)
- [Delete an Image on page 60](#)
- [Show Image Configuration on page 61](#)

### Creating New Images or Updating Existing Images

The Server Manager maintains five types of images – CISO, Ubuntu ISO, ESXi hypervisor ISO, Contrail CentOS package, and Contrail Ubuntu package.

Use the **server-manager add** command or the **server-manager upload** command to add new images to the Server Manager database.

- Use **add** when the new image is present locally on the Server Manager machine. The path provided is the image path on the Server Manager machine.
- Use **upload\_image** when the new image is present on the machine where the client program is being invoked. The path provided is the image path on the client machine.

### Add an Image

```
server-manager add [-h] [ --config_file CONFIG_FILE]
                    image [--file_name FILE_NAME]
```

Optional arguments include the following:

Option	Description
<b>-h, --help</b>	Show the options available for the current command and exit.
<b>--config_file CONFIG_FILE, -c CONFIG_FILE</b>	The name of the Server Manager client configuration file. The default file is: <b>/opt/contrail/server_manager/client/sm-client-config.ini</b>
<b>--file_name FILE_NAME, -f FILE_NAME</b>	The name of the JSON file that contains the image parameter values.

If no JSON file is specified, the client program accepts parameter values interactively, then builds a JSON file and makes a REST API call to the Server Manager.

The JSON file contains an array of possible entries, in the following sample format. The sample shows three images: one CentOS ISO containing Contrail packages, one Ubuntu base ISO, and one Contrail Ubuntu package. When the images are added, corresponding distribution, profile, and repository entries are created in Cobbler by the Server Manager.

```
{
  "image": [
    {
      "id": "ubuntu-12.04.3",
      "type": "ubuntu",
      "version": "ubuntu-12.04.3",
      "path": "/iso/ubuntu-12.04.3-server-amd64.iso"
    },
    {
      "id": "centos-6.4",
      "type": "centos",
      "version": "centos-6.4",
      "path": "/iso/CentOS-6.4-x86_64-minimal.iso"
    },
    {
      "id": "contrail-ubuntu-r11-b33",
      "type": "contrail-ubuntu-package",
      "version": "contrail-ubuntu-r11-b33",
```

```

        "path": "/iso/contrail-install-packages_x.xx-xx_all.deb"
    }
]
}

```

### Upload an Image

The **server-manager upload\_image** command is similar to the **server-manager add** command, except that the path provided for the image being added is the local path on the client machine. This command is useful if the client is being run remotely, not on the Server Manager machine, and the image being added is not physically present on the Server Manager machine.

```

server-manager upload_image [-h]

--config_file CONFIG_FILE]

image_id image_version image_type file_name

```

Positional arguments include the following:

Option	Description
<b>image_id</b>	Name of the new image.
<b>image_version</b>	Version number of the new image.
<b>image_type</b>	Type of image: <b>fedora</b> , <b>centos</b> , <b>ubuntu</b> , <b>contrail-ubuntu-package</b> , <b>contrail-centos-package</b>
<b>file_name</b>	Complete path for the file.

Optional arguments include the following:

Option	Description
<b>-h, --help</b>	Show the options available for the current command and exit.
<b>--config_file CONFIG_FILE, -c CONFIG_FILE</b>	The name of the Server Manager client configuration file. The default file is: <b>/opt/contrail/server_manager/client/sm-client-config.ini</b>

### Delete an Image

Use the **server-manager delete** command to delete an image from the Server Manager database. When an image is deleted from the Server Manager database, the corresponding distribution, profile, or repository for the image is also deleted from the Cobbler database.

```

server-manager delete [-h] [ --config_file CONFIG_FILE]
image image_id

```

Positional arguments include the following:

Option	Description
<code>image_id</code>	The image ID for the image to be deleted.

Optional arguments include the following:

Option	Description
<code>-h, --help</code>	Show the options available for the current command and exit.
<code>--config_file CONFIG_FILE, -c CONFIG_FILE</code>	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code>

### Show Image Configuration

Use the **server-manager show** command to list the configuration of images from the Server Manager database. If the detail flag is specified, detailed information about the image is returned. If the optional **image\_id** is not specified, information about all the images is returned.

```
server-manager [-h] [--config_file CONFIG_FILE] image [--image_id IMAGE_ID]
               [--detail]
```

Optional arguments include the following:

Option	Description
<code>-h, --help</code>	Show the options available for the current command and exit.
<code>--config_file CONFIG_FILE, -c CONFIG_FILE</code>	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code>
<code>image_id</code>	The image ID for the image or images.
<code>--detail, -d</code>	Flag to indicate if details are requested.

## Server Manager Operational Commands for Managing Servers

The Server Manager commands in the following sections are operational commands for performing a specific operation on a server or a group of servers. These commands assume that the base configuration of entities required to execute the operational commands is already completed using configuration CLI commands.

### Reimaging Server(s)

Use the **server-manager reimage** command to reimage a server or servers with a provided base ISO and package. Servers are specified by providing match conditions to select them from the database.

Before issuing the **reimage** command, the images must be added to the Server Manager using the **create image** command, which also adds the images to Cobbler. The set of servers to be reimaged can be specified by providing match criteria for servers already added to the Server Manager database. Use the **server\_id** or **server: vns\_id, cluster\_id, pod\_id, or rack\_id**.

You must identify the base image ID to be used to reimage, plus any optional Contrail package to be used. When a Contrail package is provided, a local repository is created that can be used for subsequent provisioning of reimaged servers.

The command prompts for a confirmation before making the REST API call to the Server Manager to start reimaging the servers. This confirmation message can be bypassed by specifying the optional **--no\_confirm** or **-F** parameter on the command line.

```
server-manager reimage [-h]
                        [ --config_file CONFIG_FILE]
                        [--package_image_id PACKAGE_IMAGE_ID]
                        [--no_reboot]
                        (--server_id SERVER_ID | --cluster_id CLUSTER_ID | --tag
                        <tag_name=tag_value>)
                        [--no_confirm]
                        base_image_id
```

Positional arguments include the following:

Option	Description
<b>base_image_id</b>	The image ID of the base image to be used.

Optional arguments include the following:

Option	Description
<b>-h, --help</b>	Show the options available for the current command and exit.
<b>--config_file CONFIG_FILE, -c CONFIG_FILE</b>	The name of the Server Manager client configuration file. The default file is: <b>/opt/contrail/server_manager/client/sm-client-config.ini</b>
<b>--package_image_id PACKAGE_IMAGE_ID, -p PACKAGE_IMAGE_ID</b>	The optional Contrail package to be used to reimage the server or servers.
<b>--no_reboot, -n</b>	Optional parameter to indicate that the server should not be rebooted following the reimage setup.
<b>--server_id SERVER_ID</b>	The server ID for the server or servers to be reimaged.
<b>--cluster_id CLUSTER_ID</b>	The cluster ID for the server or servers to be reimaged.



Option	Description
<code>--tag TagName=TagValue</code>	TagName which is to be matched with Tagvalue
<code>--no_confirm, -F</code>	Flag to bypass confirmation message, default = do NOT bypass.

### Provisioning and Configuring Roles on Servers

Use the **server-manager provision** command to provision identified server(s) with configured roles for the virtual network system. The servers can be selected from the database configuration (using standard server match criteria), identified in a JSON file, or provided interactively.

From the configuration of servers in the database, the Server Manager determines which roles to configure on which servers and uses this information along with other server and VNS parameters from the database to achieve the task of configuring the servers with specific roles.

When the **server-manager provision** command is used, the Server Manager builds the manifest files corresponding to each of the servers and pushes them to the Puppet agent for execution upon the servers.

```
server-manager provision [-h]
                        [--config_file CONFIG_FILE]
                        (--server_id SERVER_ID | --cluster_id CLUSTER_ID | --tag
<tag_name=tag_value> | --provision_params_file PROVISION_PARAMS_FILE |
--interactive)
                        [--no_confirm]
                        package_image_id
```

Positional arguments include the following:

Option	Description
<code>package_image_id</code>	The Contrail package image ID to be used for provisioning.

Optional arguments include the following:

Option	Description
<code>-h, --help</code>	Show the options available for the current command and exit.
<code>--config_file CONFIG_FILE, -c CONFIG_FILE</code>	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code>
<code>--server_id SERVER_ID</code>	The server ID for the server or servers to be provisioned.
<code>--cluster_id CLUSTER_ID</code>	The cluster ID for the server or servers to be provisioned.
<code>--tag TagName=TagValue</code>	TagName to be matched with Tagvalue.

Option	Description
<code>--provision_params_file PROVISION_PARAMS_FILE, -f PROVISION_PARAMS_FILE</code>	Optional JSON file containing the parameters for provisioning the server(s).
<code>--interactive, -I</code>	Flag indicating that you are manually entering the server parameters for provisioning.
<code>--no_confirm, -F</code>	Flag to bypass confirmation message, default = do NOT bypass.

You can specify roles different from what is configured in the database by using the JSON file option parameter. When using the file option, the rest of the server parameters, the cluster parameters, and the list of servers must be configured before using the provision command. The following is a sample format for the file option:

```
{
  "roles" : {
    "database" : ["demo2-server"],
    "openstack" : ["demo2-server"],
    "config" : ["demo2-server"],
    "control" : ["demo2-server"],
    "collector" : ["demo2-server"],
    "webui" : ["demo2-server"],
    "compute" : ["demo2-server"]
  }
}
```

The final option for specifying roles for provisioning servers is to specify the **--interactive** option flag. When the provision command is used, you are prompted to enter role definitions interactively.

---

### Adding and Deleting Roles

You can add or delete roles through server manager to expand or shrink your cluster. For example, you might want to expand the number of controllers or compute nodes when there is more traffic. Conversely, you could shrink the number of controllers or compute nodes when there is less traffic.

Adding and deleting roles is done by modifying the roles section in the server's JSON and issuing the provision command.

The following example shows the JSON to remove all of the controller roles from a server cluster.

```
{
  "cluster_id": "test-cluster",
```

```

        "id": "<ID>",
        "ip_address": "<ip address>",
        "roles": [
            "config",
            "control",
            "collector",
            "webui",
            "database",
            "openstack"
        ]
    },

```

The resulting role changes are shown in the following.

```

{
    "cluster_id": "test-cluster",
    "id": "<ID>",
    "ip_address": "<ip address>",
    "roles": [
    ]
},

```

Use **server-manager add** and **server-manager provision** for the whole cluster. Provisioning the whole cluster enables the references to the server to be added or removed in the other controllers of the cluster.

The following example shows how to add roles to a server cluster, assuming the server initially has no roles, as shown.

```

{
    "cluster_id": "test-cluster",
    "id": "<ID>",
    "ip_address": "<ip address>",
    "roles": [
    ]
},

```

The following shows the addition of a controller and its roles.

```

{
    "cluster_id": "test-cluster",
    "id": "<ID>",
    "ip_address": "<ip address>",
    "roles": [
        "config",
        "control",
        "collector",
        "webui",
        "database",
        "openstack"
    ]
},

```

Use **server-manager add** and **server-manager provision** for the whole cluster. Provisioning the whole cluster enables the references to the server to be added or removed in the other controllers of the cluster.

A similar procedure can be followed to add a compute node.



**NOTE:** The following are caveats for adding and deleting roles:

- Adding and deleting roles is supported only on a high availability (HA) cluster.
- You cannot add and delete in a single provision command. You can perform all adds in a single provision or all deletes in a single provision.
- A node that is down or unreachable cannot be deleted, because the Puppet agent cannot run to perform the add or delete.

### Restarting Server(s)

Use the **server-manager restart** command to reboot identified server(s). Servers can be specified from the database by providing standard match conditions. The **restart** command provides a way to reboot or power-cycle the servers, using the Server Manager REST API interface. If reimaging is intended, use the **restart** command with the **net-boot** flag enabled. When netbooted, the Puppet agent is also installed and configured on the servers. If there are Puppet manifest files created for the server prior to rebooting, the agent pulls those from the Server Manager and executes the configured Puppet manifests. The **restart** command uses an IPMI mechanism to power cycle the servers, if available and configured. Otherwise, the **restart** command uses SSH to the server and the existing reboot command mechanism is used.

```
server-manager restart [-h]
                        [ --config_file CONFIG_FILE]
                        (--server_id SERVER_ID | --cluster_id CLUSTER_ID | --tag
                        <tag_name=tag_value>)
                        [--net_boot]
                        [--no_confirm]
```

Optional arguments include the following:

Option	Description
-h, --help	Show the options available for the current command and exit.
--config_file CONFIG_FILE, -c CONFIG_FILE	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code> .
--server_id SERVER_ID	The server ID for the server or servers to be restarted.
--cluster_id CLUSTER_ID	The cluster ID for the server or servers to be restarted.
--tag TagName=TagValue	TagName to be matched with Tagvalue.
--net_boot, -n	Optional parameter to indicate if the server should be netbooted.

Option	Description
<code>--no_confirm, -F</code>	Flag to bypass confirmation message, default = do NOT bypass.

### Show Status of Server(s)

Use the **server-manager status** command to view the reimaging or provisioning status of server(s).

```
server-manager status server [-h]
                             [--config_file CONFIG_FILE]
                             (--server_id SERVER_ID | --cluster_id CLUSTER_ID
                             | --tag <tag_name=tag_value>)
```

Optional arguments include the following:

Option	Description
<code>-h, --help</code>	Show the options available for the current command and exit.
<code>--config_file CONFIG_FILE, -c CONFIG_FILE</code>	The name of the Server Manager client configuration file. The default file is: <code>/opt/contrail/server_manager/client/sm-client-config.ini</code>
<code>--server_id SERVER_ID</code>	The server ID for the server whose status is to be fetched.
<code>--cluster_id CLUSTER_ID</code>	The cluster ID for the server or servers to be restarted.
<code>--tag TagName=TagValue</code>	TagName to be matched with Tagvalue.

The status command provides a way to fetch the current status of a server.

Status outputs include the following:

restart\_issued  
reimage\_started  
provision\_started  
provision\_completed  
database\_started  
database\_completed  
openstack\_started  
openstack\_completed  
config\_started  
config\_completed  
control\_started  
control\_completed  
collector\_started  
collector\_completed  
webui\_started  
webui\_completed  
compute\_started  
compute\_completed

## Server Manager REST API Calls

This section describes all of the REST API calls to the Server Manager. Each description includes an example configuration.

- [REST APIs for Server Manager Configuration Database Entries on page 69](#)
- [API: Add a Server on page 69](#)
- [API: Delete Servers on page 70](#)
- [API: Retrieve Server Configuration on page 70](#)
- [API: Add an Image on page 70](#)
- [API: Upload an Image on page 71](#)
- [API: Get Image Information on page 71](#)
- [API: Delete an Image on page 71](#)
- [API: Add or Modify a Cluster on page 72](#)
- [API: Delete a Cluster on page 72](#)
- [API: Get Cluster Configuration on page 73](#)
- [API: Get All Server Manager Configurations on page 73](#)
- [API: Reimage Servers on page 73](#)

- [API: Provision Servers on page 73](#)
- [API: Restart Servers on page 74](#)

### REST APIs for Server Manager Configuration Database Entries

The REST API calls in this section help in configuring different elements in the Server Manager database.



**NOTE:** The IP addresses and other values in the following are shown for example purposes only. Be sure to use values that are correct for your environment.

#### API: Add a Server

To add a new server to the service manager configuration database:

URL: `http://<SM-IP-Address>:<SM-Port>/server`

Method: **PUT**

Payload: JSON payload containing an array of servers to be added. For each server in the array, all the parameters are specified as JSON fields. The mask, gateway, password, and domain fields are optional, and if not specified, the values of these fields are taken from the cluster to which the server belongs.

The following is a sample JSON file for adding a server.

```
{
  "server": [
    {
      "id": "demo2-server",
      "mac_address": "<mac address>",
      "ip_address": "<ip address>",
      "parameters": {
        "interface_name": "eth1"
      },
      "roles": [
        "config",
        "openstack",
        "control",
        "compute",
        "collector",
        "webui",
        "database"
      ],
      "cluster_id": "<cluster name>",
      "mask": "<mask address>",
      "gateway": "<ip address>",
      "password": "<password>",
      "domain": "<domain name>",
      "email": "id@company.net",
      "tag": {
        "datacenter": "demo-dc",
        "rack": "demo-rack"
      },
    },
  ],
}
```

```
    }
  ]
}
```

### API: Delete Servers

---

Use one of the following formats to delete a server.

URL: `http://<SM-IP-Address>:<SM-Port>/server?server_id=SERVER_ID`

`http://<SM-IP-Address>:<SM-Port>/server?cluster_id=CLUSTER_ID`

`http://<SM-IP-Address>:<SM-Port>/server?mac=MAC`

`http://<SM-IP-Address>:<SM-Port>/server?ip=IP`

`http://<SM-IP-Address>:<SM-Port>/server[?tag=<tag_name>=<tag_value>,,]`

Method : DELETE

Payload : None

### API: Retrieve Server Configuration

---

Use one of the following methods to retrieve a server configuration. The detail argument is optional, and specified as part of the URL if details of the server entry are requested.

URL: `http://<SM-IP-Address>:<SM-Port>/server[?server_id=SERVER_ID&detail]`

`http://<SM-IP-Address>:<SM-Port>/server[?cluster_id=CLUSTER_ID&detail]`

`http://<SM-IP-Address>:<SM-Port>/server[?tag=<tag_name>=<tag_value>,,]`

`http://<SM-IP-Address>:<SM-Port>/server[?mac=MAC&detail]`

`http://<SM-IP-Address>:<SM-Port>/server[?ip=IP&detail]`

`http://<SM-IP-Address>:<SM-Port>/server[?tag=<tag_name>=<tag_value>,,]`

Method : GET

Payload : None

### API: Add an Image

---

Use the following to add a new image to the Server Manager configuration database from the Server Manager machine.

An image is either an ISO for a CentOS or Ubuntu distribution or an Ubuntu Contrail package repository. When adding an image, the image file is assumed to be available on the Server Manager machine.

URL : `http://<SM-IP-Address>:<SM-Port>/image`

Method: PUT



Payload: Specifies all the parameters that define the image being added.

```
{
  "image": [
    {
      "id": "Image-id",
      "type": "image_type", <ubuntu or centos or esxi5.1 or esxi5.5 or
      contrail-ubuntu-package or contrail-centos-package>
      "version": "image_version",
      "path": "path-to-image-on-server-manager-machine"
    }
  ]
}
```

### API: Upload an Image

Use the following to upload a new image from a client to the Server Manager configuration database.

An image is an ISO for a CentOS or Ubuntu distribution or an Ubuntu Contrail package repository. Add image assumes the file is available on the Server Manager, whereas upload image transfers the image file from the client machine to the Server Manager machine.

URL : **http://<SM-IP-Address>:<SM-Port>/image/upload**

Method: **PUT**

Payload: Specifies all the parameters that define the image being added.

```
{
  "image": [
    {
      "id": "Image-id",
      "type": "image_type", <ubuntu or centos or esxi5.1 or esxi5.5 or
      contrail-ubuntu-package or contrail-centos-package>
      "version": "image_version",
      "path": "path-to-image-on-client-machine"
    }
  ]
}
```

### API: Get Image Information

Use the following to get image information.

URL : **http://<SM-IP-Address>:<SM-Port>/image[?image\_id=IMAGE\_ID&detail]**

Method: **GET**

Payload: Specifies criteria for the image being sought. If no match criteria is specified, information about all the images is provided. The details field specifies if details of the image entry in the database are requested.

### API: Delete an Image

Use the following to delete an image.

URL: `http://<SM-IP-Address>:<SM-Port>/image?image_id=IMAGE_ID`

Method: **DELETE**

Payload: Specifies criteria for the image being deleted.

---

### API: Add or Modify a Cluster

Use the following to add a cluster to the Server Manager configuration database. A cluster maintains parameters for a set of servers that work together in different roles to provide complete functions for a Contrail cluster.

URL: `http://<SM-IP-Address>:<SM-Port>/cluster`

Method: **PUT**

Payload: Contains the definition of the cluster, including all the global parameters needed by all the servers in the cluster. The `subnet_mask`, `gateway`, `password`, and `domain` fields define parameters that apply to all servers in the VNS. These parameter values can be individually overridden for a server by specifying different values in the server entry.

```
{
  "cluster" : [
    {
      "id" : "demo-cluster",
      "parameters" : {
        "router_asn": "<asn number>",
        "database_dir": "/home/cassandra",
        "database_token": "",
        "use_certificates": "False",
        "multi_tenancy": "False",
        "encapsulation_priority": "MPLSoUDP,MPLSoGRE,VXLAN",
        "service_token": "<password>",
        "keystone_user": "admin",
        "keystone_password": "<password>",
        "keystone_tenant": "admin",
        "analytics_data_ttl": "168",
        "subnet_mask": "<subnet mask address>",
        "gateway": "<ip address>",
        "password": "<password>",
        "haproxy": "disable",
        "external_bgp": "",
        "domain": <domain name>
      }
    }
  ]
}
```

---

### API: Delete a Cluster

Use this API to delete a cluster from the Server Manager database.

URL: `http://<SM-IP-Address>:<SM-Port>/cluster?cluster_id=CLUSTER_ID`

Method: **DELETE**

Payload: None

### API: Get Cluster Configuration

---

Use this API to get a cluster configuration.

URL: `http://<SM-IP-Address>:<SM-Port>/cluster[?cluster_id=CLUSTER_ID&detail]`

Method: **GET**

Payload: None

The optional detail argument is specified as part of the URL if details of the VNS entry are requested.

### API: Get All Server Manager Configurations

---

Use this API to get all configurations of Server Manager objects, including servers, clusters, images, and tags.

URL: `http://<SM-IP-Address>:<SM-Port>/all[?detail]`

Method: **GET**

Payload: None

The optional detail argument is specified as part of the URL if details of the Server Manager configuration are requested.

### API: Reimage Servers

---

Use one of the following API formats to reimage one or more servers.

URL: `http://<SM-IP-Address>:<SM-Port>/server/reimage?server_id=SERVER_ID`

`http://<SM-IP-Address>:<SM-Port>/server/reimage?cluster_id=CLUSTER_ID`

`http://<SM-IP-Address>:<SM-Port>/server/reimage?mac=MAC`

`http://<SM-IP-Address>:<SM-Port>/server/reimage?ip=IP`

`http://<SM-IP-Address>:<SM-Port>/server/reimage [?tag=<tag_name>=<tag_value>,.]`

Method: **POST**

Payload: None

### API: Provision Servers

---

Use this API to provision or configure one or more servers for roles configured on them.

URL: `http://<SM-IP-Address>:<SM-Port>/server/provision`

Method: **POST**

Payload: Specifies the criteria to be used to identify servers which are being provisioned. The servers can be identified by server\_id, mac, cluster\_id or tags. See the following example.

```
{
  server_id : <server_id> OR
```

```
mac : <server_mac_address> OR
cluster_id : <cluster_id> OR
tag : {"data-center" : "dc1"} OR
provision_parameters = {
  "roles" : {
    "database" : ["demo2-server"],
    "openstack" : ["demo2-server"],
    "config" : ["demo2-server"],
    "control" : ["demo2-server"],
    "collector" : ["demo2-server"],
    "webui" : ["demo2-server"],
    "compute" : ["demo2-server"]
  }
}
```

---

### API: Restart Servers

This REST API is used to power cycle the servers and reboot either with net-booting enabled or disabled.

If the servers are to be reimaged and reprovisioned, the **net-boot** flag should be set.

If servers are only being reprovisioned, the **net-boot** flag is not needed, however, the Puppet agent must be running on the target systems with the correct puppet configuration to communicate to the puppet master running on the Server Manager.

URL: `http://<SM-IP-Address>:<SM-Port>/server/restart?server_id=SERVER_ID`  
`http://<SM-IP-Address>:<SM-Port>/server/restart?[netboot&]cluster_id=CLUSTER_ID`  
`http://<SM-IP-Address>:<SM-Port>/server/restart? [netboot&]mac=MAC`  
`http://<SM-IP-Address>:<SM-Port>/server/restart? [netboot&]ip=IP`  
`http://<SM-IP-Address>:<SM-Port>/server/restart ?`  
`[netboot&]tag=<tag_name>=<tag_value>`

Method: **POST**

Payload: Specifies the criteria to be used to identify servers which are being restarted. The servers can be identified by their **server\_id**, **mac**, **cluster\_id**, or **tag**. The netboot parameter specifies if the servers being power-cycled are to be booted from Cobbler or locally.

### Example: Reimaging and Provisioning a Server

This example shows the steps used in Server Manager software to configure, reimage, and provision a server running all roles of the Contrail system in a single-node configuration.



**NOTE:** Component names and IP addresses in the following are used for example only. To use this example in your own environment, be sure to use addresses and names specific to your environment.

---

The Server Manager client configuration file used for the following CLI commands, is `/opt/contrail/server_manager/client/sm-client-config.ini`. It contains the values for the server IP address and port number as follows:

**[SERVER-MANAGER]**

**listen\_ip\_addr = 192.168.1.10 (Server Manager IP address)**

**listen\_port = 9001**

The steps to be followed include:

1. Configure cluster.
2. Configure servers.
3. Configure images.
4. Reimage servers (either using servers configured above or using explicitly specified reimage parameters with the request).
5. Provision servers (either using servers configured above or using explicitly specified provision parameters with the request).

1. Configure a cluster.

**server-manager add cluster -f cluster.json**

Where cluster.json contains :

```
{
  "cluster" : [
    {
      "id" : "demo-cluster",
      "parameters" : {
        "router_asn": "<asn number>",
        "database_dir": "/home/cassandra",
        "database_token": "",
        "use_certificates": "False",
        "multi_tenancy": "False",
        "encapsulation_priority": "MPLSoUDP,MPLSoGRE,VXLAN",
        "service_token": "<password>",
        "keystone_user": "admin",
        "keystone_password": "<password>",
        "keystone_tenant": "admin",
        "analytics_data_ttl": "168",
        "subnet_mask": "<subnet mask address>",
        "gateway": "<ip address>",
        "password": "<password>",
        "haproxy": "disable",
        "external_bgp": "",
        "domain": "<domain name>"
      }
    }
  ]
}
```

2. Configure the server.

```
server-manager add server -f server.json
```

Where server.json contains :

```
{
  "server": [
    {
      "id": "demo2-server",
      "mac_address": "<mac address>",
      "ip_address": "<ip-address>",
      "parameters": {
        "interface_name": "eth1",
      },
      "roles": [
        "config", "openstack", "control", "compute", "collector", "webui", "database"
      ],
      "cluster_id": "demo-cluster",
      "subnet_mask": "<subnet mask address>",
      "gateway": "<ip-address>",
      "password": "<password>",
      "domain": "<domain name>t",
      "ipmi_address": "<ip-address>"
    }
  ]
}
```

### 3. Configure images.

In the example, the image files for **ubuntu-12.04.3** and **contrail-ubuntu-164** are located at the corresponding image path specified on the Server Manager.

```
server-manager add -c smgr_client_config.ini image -f image.json
```

Where image.json contains:

```
{
  "image": [
    {
      "id": "ubuntu-12.04.3",
      "type": "ubuntu",
      "version": "ubuntu-12.04.3",
      "path": "/iso/ubuntu-12.04.3-server-amd64.iso"
    },
    {
      "id": "contrail-ubuntu-xxx",
      "type": "contrail-ubuntu-package",
      "version": "contrail-ubuntu-xxx",
      "path": "/iso/contrail-install-packages_x.xx-xxx-xxxxxxxxx_all.deb"
    }
  ]
}
```

### 4. Reimage servers.

This step can be performed after the configuration from the previous steps is in the Server Manager database.

```
server-manager reimage --server_id demo-server -r contrail-ubuntu-164 ubuntu-12.04.3
```

### 5. Provision servers.

```
server-manager provision --server_id demo-server contrail-ubuntu-164
```



**NOTE:** Optionally, the Contrail package to be used can be specified with the `reimage` command, in which case the repository with the Contrail packages is created and made available to the target nodes as part of the `reimage` process. The repository is a mandatory parameter of the `provision` command.

## Using the Server Manager Web User Interface

When the Server Manager is installed on your Contrail system, you can also install a Server Manager Web user interface that you can use to access the features of Server Manager.

- [Log In to Server Manager on page 77](#)
- [Create a Cluster for Server Manager on page 78](#)
- [Working with Servers in the Server Manager User Interface on page 84](#)
- [Add a Server on page 84](#)
- [Edit Tags for Servers on page 86](#)
- [Using the Edit Config Option for Multiple Servers on page 86](#)
- [Filter Servers by Tag on page 87](#)
- [Viewing Server Details on page 87](#)
- [Configuring Images and Packages on page 87](#)
- [Add New Image or Package on page 88](#)
- [Selecting Server Manager Actions for Clusters on page 88](#)
- [Reimage a Cluster on page 88](#)
- [Provision a Cluster on page 89](#)

### Log In to Server Manager

The Server Manager user interface can be accessed using:

**`http://<server-manager-user-interface-ip>:9080`**

Where **`<server-manager-user-interface-ip>`** is the IP address of the server on which the Server Manager web user interface is installed.

From the Contrail user interface, select **Setting > Server Manager** to access the Server Manager home page. From this page you can manage Server Manager settings for clusters, servers, images, and packages.

## Create a Cluster for Server Manager

Select **Add Cluster** to identify a cluster to be managed by the Server Manager. Select **Setting > Server Manager > Clusters**, to access the **Clusters** page, as shown:

Name	Email	New Servers	Configured Servers	In-Reimage Servers	Reimage Servers	In-Provision Servers	Provisioned Servers	Total Servers
setup1	cluster-admin@domain.com	0	0	0	0	0	1	1
setup2	cluster-admin@domain.com	0	0	0	0	0	1	1

To create a new cluster, click the plus icon in the upper right of the **Clusters** page. The **Add Cluster** window is displayed. In the **Add Cluster** window, you can add a new cluster ID and the domain e-mail address of the cluster.

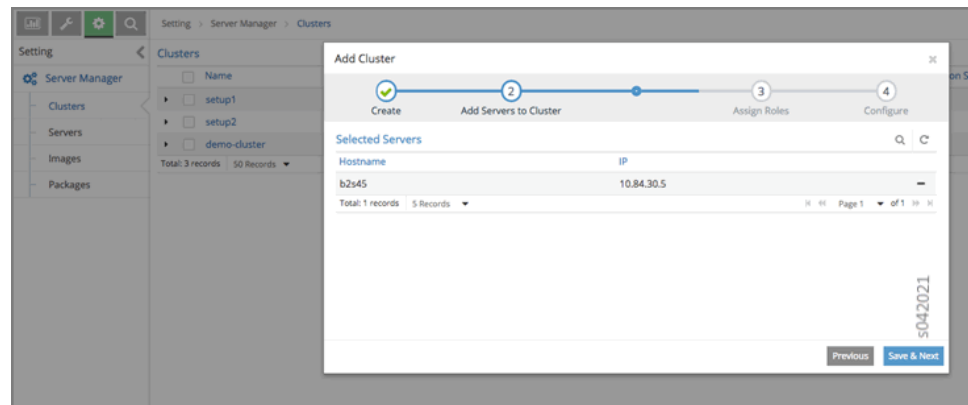
When you are finished adding information about the new cluster in the **Add Clusters** window, click **Save & Next**. Now you can add servers to the cluster, as shown in the following.

Hostname	IP	
<input checked="" type="checkbox"/> b2s45	10.84.30.5	
<input type="checkbox"/> b2s40	10.84.30.1	

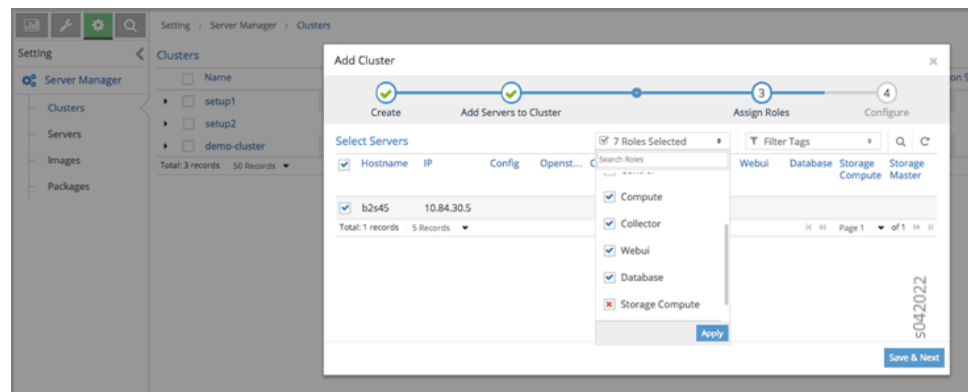
Click the check box of each server to be added to the cluster.



When you are finished, click **Next**. The selected servers are added to the cluster.

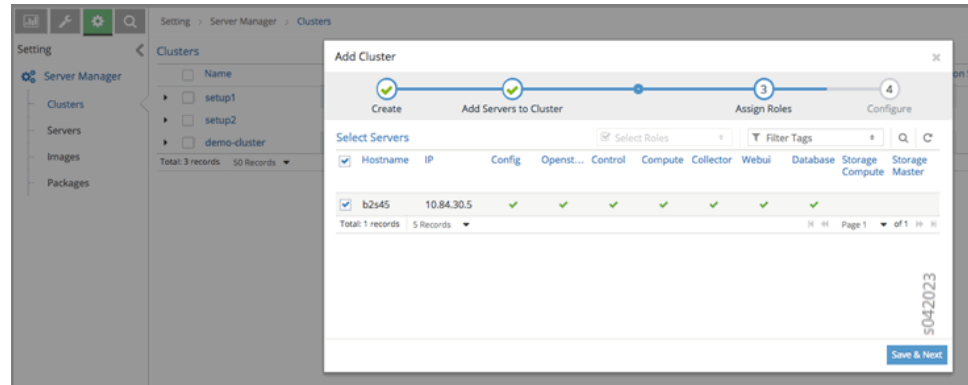


When you are finished adding servers, click **Save & Next**. Now you can assign Contrail roles to servers that you select in the cluster. Roles available are Config, OpenStack, Control, Compute, and Collector. Select each role assignment for the selected server. You can also unselect any assigned role. The assigned roles correspond to the role functions in operation on the server.

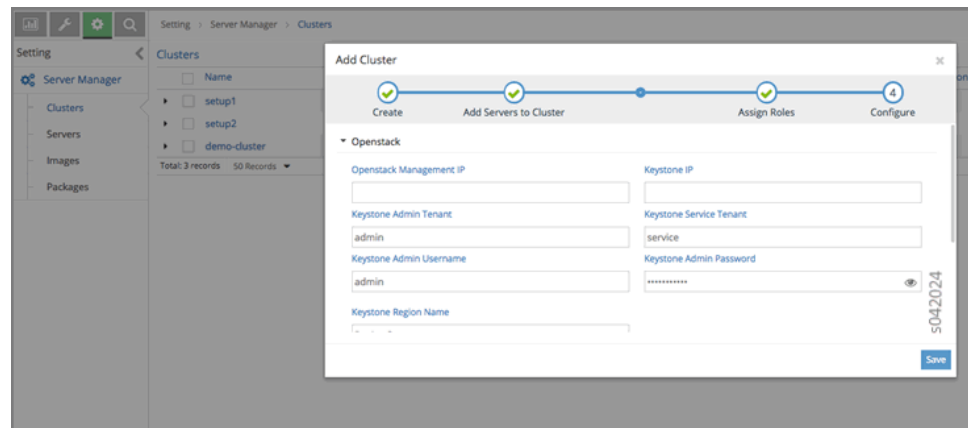


When you are finished selecting roles for the selected server in the **Roles** window, click **Apply** to save your choices.

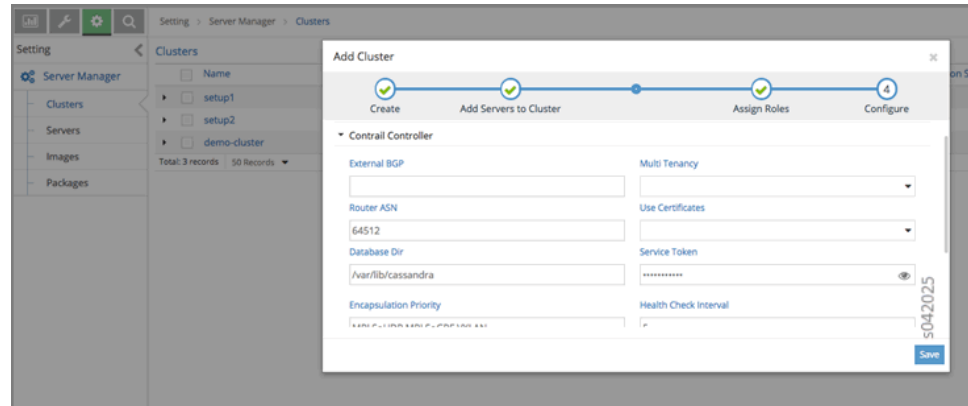
Click **Save & Next** to view your selections. Check marks are displayed in the columns of the **Add Cluster** window, see the following image.




The next step after roles are assigned is to enter the cluster configuration information for OpenStack. After viewing the assigned roles, click **Save & Next**. The **Add Cluster** window is displayed. Click an icon that opens a set of fields where you can enter OpenStack or Contrail configuration information for the cluster. In the following image, the **Openstack** icon is selected. You can enter configuration information, such as OpenStack or Keystone passwords and usernames in the fields.



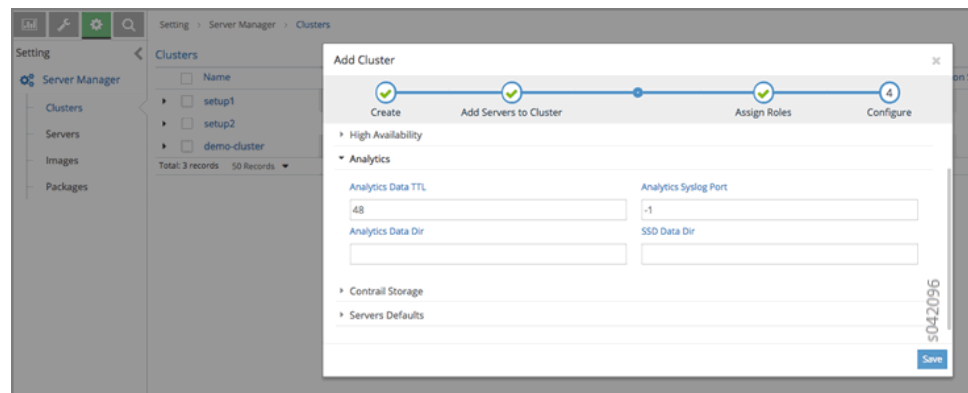
In the following image, the Contrail controller icon is selected. You can enter configuration information for Contrail, such as **External BGP**, **Multi Tenancy**, **Router ASN**, **HA Proxy**, and so on.




In the following image, the High Availability (HA) icon is selected. You can configure high availability parameters such as **HA Proxy**, **Internal** and **External VIP**, and so on.



In the following image, the **Analytics** icon is selected. Here the user can configure parameters for Contrail Analytics, including **TTL**, **Syslog Port**, **Data Dir**, and so on.




In following image, the **Contrail Storage** icon is selected. You can configure parameters for Contrail Storage, including **Monitoring Key**, **OSD Bootstrap Key**, **Admin Key**, and so on.



When you are finished entering all of the cluster configuration information, click **Save** to submit the configurations. You can view all configured clusters on the **Clusters** window by selecting **Setting > Server Manager > Clusters**.

Name	Email	New Servers	Configured Servers	In-Reimage Servers	Reimage Servers	In-Provision Servers	Provisioned Servers
setup1	cluster-admin@domain.com	0	0	0	0	0	0
setup2	cluster-admin@domain.com	0	0	0	0	0	1
demo-cluster	demo-admin@domain.com	0	0	0	0	0	1

To perform an action on one of the configured clusters, click the gear wheel icon at the right to select from a menu of actions available for that cluster, including **Add Servers**, **Remove Servers**, **Assign Roles**, **Edit Config**, **Reimage**, **Provision**, and **Delete**, as shown.



You can also click the expansion icon on the left side of the cluster name to display the details of that cluster in an area below the name line, as shown.

The screenshot shows the Juniper Server Manager interface. On the left, a sidebar contains 'Setting' and 'Server Manager'. Under 'Server Manager', there are links for 'Clusters', 'Servers', 'Images', and 'Packages'. The main area is titled 'Clusters' and contains a table with columns: Name, Email, New Servers, Configured Servers, In-Reimage Servers, Reimage Servers, In-Provision Servers, and Provisioned Servers. Three clusters are listed: 'setup1', 'setup2', and 'demo-cluster'. The 'demo-cluster' is selected, and its details are shown below the table. The details are organized into sections: 'Details' (ID, Email), 'Openstack' (Management IP, Admin Tenant, Service Tenant, Username, Region Name), 'Contrail Controller' (Encapsulation Priority, Router ASN, Database Dir, Health Check Interval), 'High Availability' (HA Proxy), 'Status' (Total Servers, New Servers, Configured Servers, In-Reimage Servers, Reimage Servers, In-Provision Servers, Provisioned Servers), 'Analytics' (Data TTL, Syslog Port), 'Contrail Storage' (Storage Virsh UUID, Storage FSID), and 'Servers Defaults' (Domain, Subnet Mask). A date stamp 's042028' is visible on the right side of the details panel.

Name	Email	New Servers	Configured Servers	In-Reimage Servers	Reimage Servers	In-Provision Servers	Provisioned Servers
setup1	cluster-admin@domain.com	0	0	0	0	0	0
setup2	cluster-admin@domain.com	0	0	0	0	0	1
demo-cluster	demo-admin@domain.com	0	0	0	0	0	1

**demo-cluster Details**

- ID:** demo-cluster
- Email:** demo-admin@domain.com
- Openstack:**
  - Openstack Management IP: -
  - Keystone Admin Tenant: admin
  - Keystone Service Tenant: service
  - Keystone Admin Username: admin
  - Keystone Region Name: RegionOne
- Contrail Controller:**
  - Encapsulation Priority: MPLSoUDP,MPLSoGRE,VXLAN
  - Router ASN: 64512
  - Database Dir: /var/lib/cassandra
  - Health Check Interval: 5
- High Availability:**
  - HA Proxy: disable
- Status:**
  - Total Servers: 1
  - New Servers: 0
  - Configured Servers: 0
  - In-Reimage Servers: 0
  - Reimage Servers: 0
  - In-Provision Servers: 0
  - Provisioned Servers: 1
- Analytics:**
  - Analytics Data TTL: 48
  - Analytics Syslog Port: -1
- Contrail Storage:**
  - Storage Virsh UUID: c542fac0-0c28-4ed8-82a8-abe476c7ba2d
  - Storage FSID: 266b08f2-7746-42b4-8ba8-b56b3977a54e
- Servers Defaults:**
  - Domain: englab.juniper.net
  - Subnet Mask: 255.255.255.0

Total: 3 records | 50 Records

Click the upper right icon to switch to the JSON view to see the contents of the JSON file for the cluster.

The screenshot shows the Juniper Server Manager interface with the 'demo-cluster' selected. The details panel is now in JSON view, displaying the configuration for the cluster. A date stamp 's042029' is visible on the right side of the details panel.

```

- {
  parameters: - {
    domain: englab.juniper.net
    keystone_ip:
    analytics_data_dir:
    keystone_region_name: RegionOne
    encapsulation_priority: MPLSoUDP,MPLSoGRE,VXLAN
    keystone_username: admin
    analytics_data_ttl: 48
    subnet_mask: 255.255.255.0
    nfs_glance_path: null
    admin_key: null
    keystone_password: contrail123
    router_asn: 64512
    database_dir: /var/lib/cassandra
    analytics_syslog_port: -1
    keystone_tenant: admin
    gateway: null
    database_token:
    haproxy: disable
    storage_virsh_uuid: c542fac0-0c28-4ed8-82a8-abe476c7ba2d
    password: c0ntrail123
    uuid: 4fa141b2-1eed-4b69-b2a3-e8293492da2f
    service_token: contrail123
    external_bgp:
    storage_fs_id: 266b08f2-7746-42b4-8ba8-b56b3977a54e
    internal_vip:
    ssd_data_dir:
    hc_interval: 5
    osd_bootstrap_key: null
    openstack_passwd: c0ntrail123
    multi_tenancy:
    storage_mon_secret: null
    use_certificates:
    openstack_mgmt_ip: null
    keystone_service_tenant: service
    nfs_server:
    external_vip:
  }
  package_image_id:

```

The cluster name is a link, click the cluster name to display the cluster **Details** page, as shown.

The screenshot shows the 'demo-cluster' details page in the Contrail Server Manager. The left sidebar has 'Server Manager' selected, with 'Clusters' highlighted. The main content area is divided into sections: Details, Openstack, Contrail Controller, High Availability, Status, Analytics, Contrail Storage, and Servers Defaults. The 'Servers' section at the bottom shows a table with one server entry.

ID	Tags	IP	IPMI	Config	Openstack	Control	Compute	Collector	Webui	Database	Storage	Compute
b2s45	control-lab rack-b2 floor-6	10.84.30.5	10.84.60.148	✓	✓	✓	✓	✓	✓	✓	✓	✓

Total: 1 records 50 Records

## Working with Servers in the Server Manager User Interface

Select **Setting > Server Manager** and click the **Servers** link in the left sidebar at to view a list of all servers.

The screenshot shows the 'Servers' page in the Contrail Server Manager. The left sidebar has 'Server Manager' selected, with 'Servers' highlighted. The main content area shows a table with two server entries.

ID	Cluster	Tags	IP	IPMI	Status
b2s45	demo-cluster	control-lab rack-b2 floor-6	10.84.30.5	10.84.60.148	provision_completed
b2s40	setup2	control-lab rack-b2 floor-6	10.84.30.1	10.84.60.146	provision_completed

Total: 2 records 50 Records

## Add a Server

To add a new server, select **Setting > Server Manager > Servers** and click the plus (+) icon at the upper right side in the header line. The **Add Server** window is displayed, as in the following.

The screenshot shows the 'Servers' page with the 'Add Server' dialog box open. The dialog box has two tabs: 'System Management' and 'Interfaces'. The 'System Management' tab is active, showing fields for ID, Password, Host Name, Domain, Static IP, IPMI Address, IPMI Username, IPMI Password, and Partition.

**Add Server**

**System Management**

ID: demo-server Password: [password icon]

Host Name: demo-server Domain: englab.juniper.net

Static IP: IPMI Address: 10.84.60.148

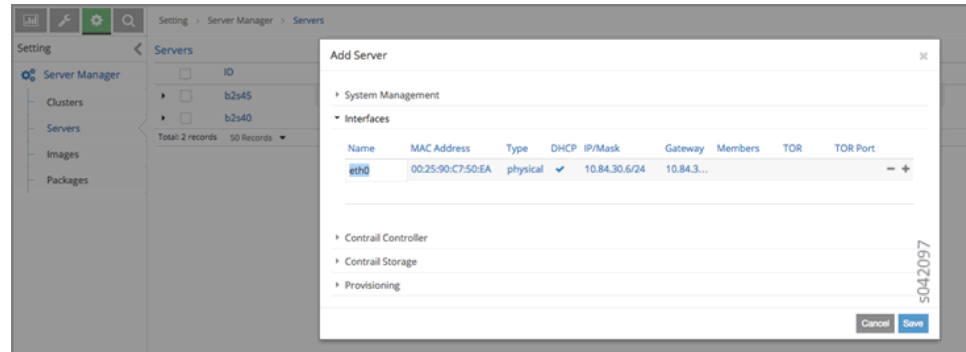
IPMI Username: ADMIN IPMI Password: [password icon]

Partition: [text input]

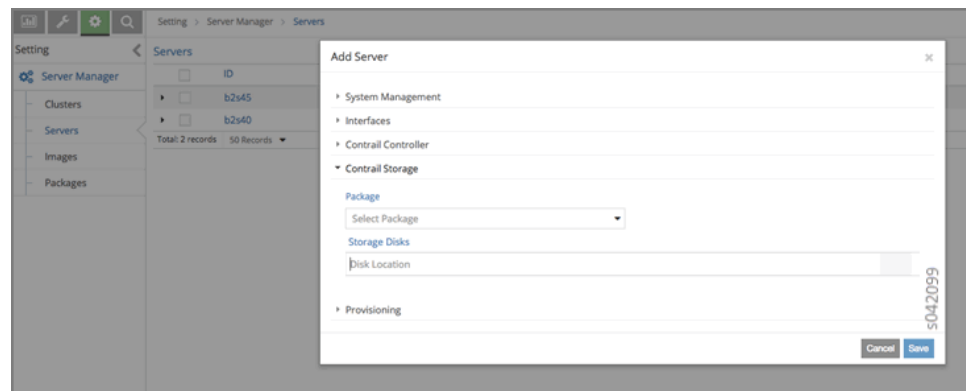
**Interfaces**

Cancel Save

In the following image, the **Interfaces** icon is selected. You can add new interfaces or edit existing interfaces. To enable editing for any field, hover the cursor on any selected field to open it.




In the following image, the **Contrail Storage** icon is selected. You can configure parameters for Contrail Storage, including selecting a package and adding storage disks locations.




When you are finished entering new server details in the **Add Server** window, click **Save** to add the new server configuration to the list of servers.

You can change details of the new server by clicking the gear wheel icon to the right side to get a list of actions available, including **Edit Config**, **Edit Tags**, **Reimage**, **Provision**, and **Delete**, as shown in the following.



## Edit Tags for Servers

Select **Edit Tags** from the gear wheel icon menu. The **Edit Tags** window is displayed. Enter any user-defined tags to be associated with the selected server, then click **Save** to add the tags to the server configuration.




## Using the Edit Config Option for Multiple Servers

You can also edit the configuration of multiple servers at one time. From the **Servers** window at **Setting > Server Manager > Servers**, select the servers you want to edit, then click a gear wheel icon at the right to open the action menu, and select **Edit Config**.

The **Edit Config** window is displayed, as shown.


Click a pencil icon to open configuration fields that can be edited. Fields include **System Management**, **Contrail Controller**, **Contrail Storage**, and so on.





## Filter Servers by Tag

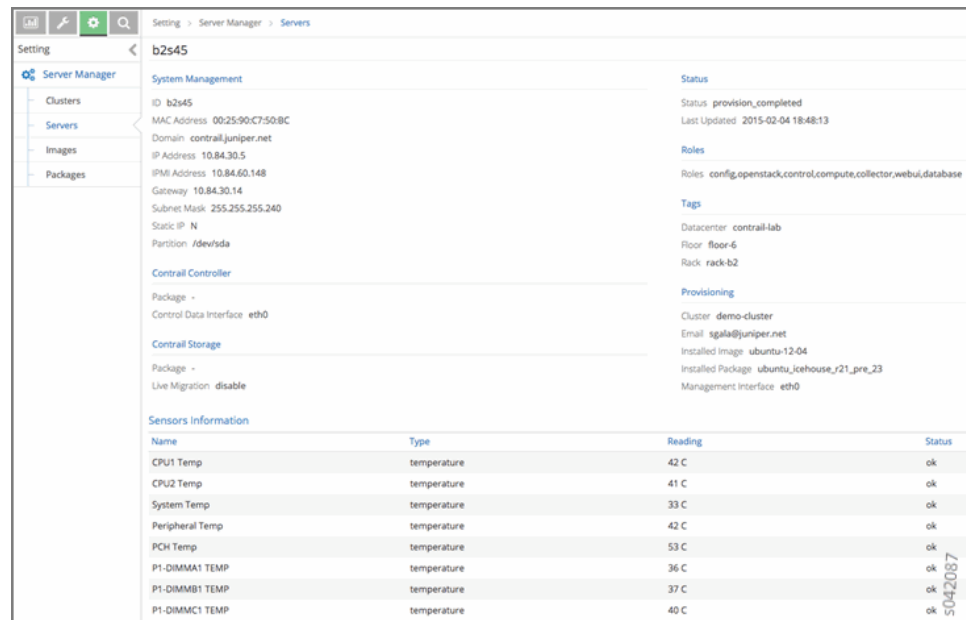
You can filter servers according to the tags defined for them. In the **Servers** window, click the **Filter Tags** field in the upper right heading. A list of configured tags is displayed. Select a tag by which to filter the list of servers.



## Viewing Server Details

Each server name on the **Servers** page is a link to the details page for that server. Click any server name to open the details for that server, as shown in the following image.

For each server, the **Sensors Information** area shows the **Name**, **Type**, **Reading**, and **Status** for the **temperature** (degrees Celsius), **fan** (rpm), and **power** (watts) sensor types.



## Configuring Images and Packages

Use the sidebar **Images** and **Packages** options to configure the software images and packages to be used by the Server Manager. Images are typically used to reimage clusters with an operating system version. Packages are used to provision clusters with a Contrail setup.

Both areas of the Server Manager user interface operate in a similar fashion. The figure shows the **Images** section. The **Packages** section has similar options.

Select **Images**. The Images page is displayed, as shown.



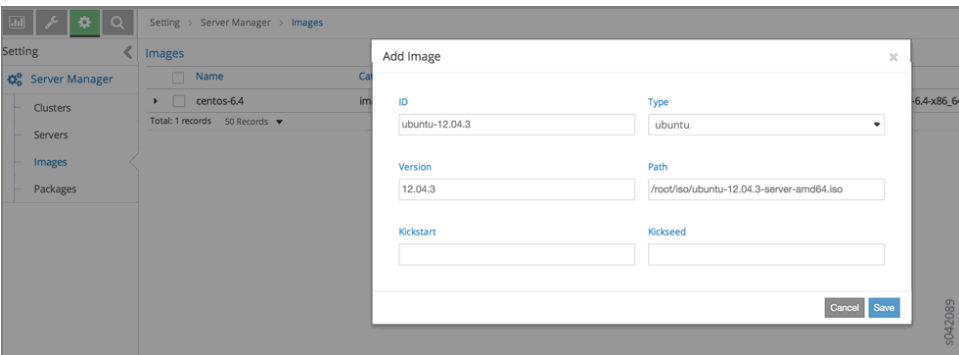
Name	Category	Type	Version	Path
ubuntu-12.04.3	image	ubuntu	12.04.3	/root/ubuntu-12.04.3-server-amd64.iso
centos-6.4	image	centos	6.4	/root/CentOS-6.4-x86_64-minimal.iso

## Add New Image or Package

To add a new image or package, on the respective **Images** or **Packages** page, click the plus (+) icon in the upper right header. The **Add Image** window is displayed. Enter the information for the new image (or package) and click **Save** to add the new item to the list of configured items.



**NOTE:** The path field requires the path of the image where it is located on the server upon which the server-manager process is running.




## Selecting Server Manager Actions for Clusters

After all aspects of a cluster are configured, you can select actions for the Server Manager to perform on the cluster, such as **Reimage** or **Provision**.

## Reimage a Cluster

Select **Setting > Servers > Clusters**. The **Clusters** window is displayed. Click the right side gear wheel icon of the cluster to be reimaged, then select **Reimage** from the action menu.


The **Reimage** dialog box is displayed, as shown. Verify that the correct image is selected in the **Default Image** field, then click **Save** to initiate the reimage action.



## Provision a Cluster

The process to provision a cluster is similar to the process to reimage a cluster. Select **Setting > Servers > Clusters**. The **Clusters** window is displayed. Click the right side gear wheel icon of the cluster to be provisioned, then select **Provision** from the action menu.

The **Provision Cluster** dialog box is displayed, as shown. Verify that the correct package for provisioning is selected in the **Default Package** field, then click **Save** to initiate the provisioning action.



## Installing and Using Server Manager Lite

This topic describes how to install and troubleshoot Server Manager Lite.

### Server Manager Lite Overview

Server Manager Lite (SM-Lite), is a streamlined version of the Server Manager software that does not include the reimage function.

SM-Lite supports the Server Manager functions of provisioning, monitoring, inventory, and webui. SM-Lite is intended to replace fab command provisioning. It allows easy deployment of Contrail provisioning and enables developers to work in isolated environments for Contrail provisioning.

SM-Lite eliminates installation and configuration of DHCP, DNS, and Cobbler services. Additionally, SM-Lite installation setup scripts are enhanced to reduce installation time.

SM-Lite provides a single command to install SM-Lite and provision a Contrail cluster.

SM-Lite introduces additional capabilities into Server Manager. The SM-Lite package is part of the Contrail Server Manager installer Debian package (`contrail-server-manager-installer_<version string>.deb`).

SM-Lite works with or without having a separate node for the SM-Lite installation, it can be installed on any Contrail node, but it is recommended to install it on the config node.

SM-Lite preserves the existing Server Manager webui functionality and it can be run on the same node as the Contrail webui. Because of that, the default port for the Server Manager webui has been changed to port **9080**.

It is important to note that the code base used for SM-Lite and Server Manager is common. Therefore, any changes or enhancements made to Server Manager provisioning functionality are automatically available in the SM-Lite software.

## Installing Server Manager Lite

The SM-Lite package is included as part of the Server Manager installer package.

The installer package also has other packages such as Server Manager, Server Manager client, Server Manager webui, and Server Manager inventory. Before provisioning commands can be executed using SM-Lite, you need to install the Server Manager installer package.

Use the following command to install the Server Manager installer package.

```
dpkg -i <contrail-server-manager-installer-deb>
```

After the Server Manager installer package is installed, all necessary Server Manager packages, scripts, and so on are made available on the server where it is installed. You can then start using Server Manager Lite commands.

## Provisioning Using SM-Lite

To provisioning the target systems, use the **provision.sh** script.

The full syntax and available options of the **provision.sh** script is:

```
/opt/contrail/contrail_server_manager/provision.sh --testbed <testbed>  
--contrail-package <contrail-package> --cluster-id <cluster-id>
```

When you include the **help** option the system displays the following:

```
/opt/contrail/contrail_server_manager/provision.sh --testbed <testbed>  
--contrail-package <contrail-package> --cluster-id <cluster-id>  
-h --help  
-c | --contrail-package <pkg>  
-t | --testbed <testbed.py>  
-cid | --cluster-id <cluster-id>
```

The **provision.sh** script performs the following functions:

1. Installs SM-Lite.

Uses the **setup.sh** installation script with the **-smlite** option to install the SM-Lite package. (**contrail-server-manager-lite\_<version-sku>.all.deb**) and all other needed packages on the system.

2. Prepares the cluster for Contrail provisioning.

Translates the parameters in the **testbed.py** file into Server Manager objects and stores them in the Server Manager database. This specifies the servers in the cluster and the configuration parameters. The cluster-id value is used, if it is specified.

3. Performs a pre-check on the target systems to ensure that they are ready for running provisioning via Puppet. Uses the Puppet modules and manifests from the Contrail package to provision the Contrail cluster.

This step issues provisioning commands for the cluster with the given Contrail package.

Server Manager Lite can be installed on any node. We recommend that you install it on the config node. Server Manager Lite can be installed on a separate node other than the Contrail cluster nodes.

With the introduction of Server Manager Lite in Contrail Release 3.0, the Server Manager webui default port is changed to **9080**. You can change the port by editing the `/etc/contrail/config.global.sm.js` file, and then restarting the `supervisor-webui-sm` process.



**NOTE:** It is also still possible to use the `provision.sh` script to install the full Server Manager (not SM-Lite) software to provision a target cluster.

## Displaying the Cluster Status

The `server-manager show cluster -detail` command displays the provisioning status of a cluster by role and by role progress.

Use the `server-manager status server` command to display the current status of the servers.

## Displaying the SM-Lite Installation and Provisioning Log Files

Log files that provide information during installation and use of SM-Lite software are available at:

- `/var/log/contrail/install_logs/install_<timestamp>.log` (SM-Lite install)
- `/var/log/contrail/install_logs/provision_<timestamp>.log` (provisioning command logs)
- `testbed_parser.log` and `preconfig.log`

## Contrail Provisioning Log Files

For each Puppet run, log files are automatically uploaded to the Server Manager at the following locations:

- `http:<sm-lite-ip-address>/logs`
- `/var/log/contrail_server_manager/<target>/<timestamp>.log`
- `/var/log/contrail/*`

You can also display the status of the processes and services using the `contrail-status` command.

### Related Documentation

- [Using Server Manager to Automate Provisioning on page 42](#)

- [Using the Server Manager Web User Interface on page 77](#)
- [Installing Server Manager on page 37](#)

## CHAPTER 5

# Installing and Using Contrail Storage

- [Installing and Using Contrail Storage on page 93](#)

## Installing and Using Contrail Storage

---

- [Overview of the Contrail Storage Solution on page 93](#)
- [Basic Storage Functionality with Contrail on page 94](#)
- [Ceph Block and Object Storage Functionality on page 94](#)
- [Using the Contrail Storage User Interface on page 95](#)
- [Hardware Specifications on page 96](#)
- [Software Files for Compute Storage Nodes on page 96](#)
- [Contrail OpenStack Nova Modifications on page 96](#)
- [Installing the Contrail Storage Solution on page 97](#)
- [Using Fabric Commands to Install and Configure Storage on page 97](#)
- [Fabric Installation Procedure on page 98](#)
- [Using Server Manager to Install and Configure Storage on page 100](#)
- [Server Manager Installation Procedure for Storage on page 100](#)
- [Example: Configurations for Storage for Reimaging and Provisioning a Server on page 101](#)
- [Storage Installation Limits on page 106](#)

## Overview of the Contrail Storage Solution

Contrail provides a storage support solution using OpenStack Cinder configured to work with Ceph. Ceph is a unified, distributed storage system whose infrastructure provides storage services to Contrail. This Contrail storage solution provides a validated Network File System (NFS) storage service, however, it is not the Ceph FS distributed file system.

The Contrail storage solution has the following features:

- Provides storage class features to Contrail clusters, including replication, reliability, and robustness.
- Uses open source components.
- Uses Ceph block and object storage functionality.

- Integrates with OpenStack Cinder functionality.
- Does not require virtual machines (VMs) to configure mirrors for replication.
- Allows nodes to provide both compute and storage services.
- Provides easy installation of basic storage functionality based on Contrail roles.
- Provides services necessary to perform virtual machine migrations between compute nodes, and supports both migratable and non-migratable virtual machines.
- Provides a Contrail-integrated user interface from which the user can monitor Ceph components and drill down for more information about components.

## Basic Storage Functionality with Contrail

The following are basic interaction points between Contrail and the storage solution.

- Cinder volumes must be manually configured prior to installing the Contrail storage solution. The Cinder volumes can be attached to virtual machines (VMs) to provide additional storage.
- The storage solution stores virtual machine boot images and snapshots in Glance, using Ceph object storage functionality.
- All storage nodes can be monitored through a graphical user interface (GUI).
- It is possible to migrate virtual machines that have ephemeral storage in Ceph.

## Ceph Block and Object Storage Functionality

Installing the Contrail storage solution creates the following Ceph configurations.

- Each disk is configured as a standalone storage device, enhancing optimal performance and creating proper failure boundaries. Ceph allocates and assigns a process called object storage daemon (OSD) to each disk.
- A replication factor of 2 is configured, consisting of one original instance plus one replica copy. Ceph ensures that each replica is on a different storage node.
- A Ceph monitor process (mon) is configured on each storage node.
- The correct number of placement groups are automatically configured, based on the number of disk drives in the cluster.
- Properly identified SSD drives are set up for use as Ceph OSD journals to reduce write latencies.
- An NFS server is created in a virtual machine within the cluster to support virtual machine migration. The NFS file system is mounted on all storage nodes, and every storage node has a shared Nova directory under the `/var/lib/nova/instances` directory. By default, this NFS file system is configured to utilize 30% of the total initial Contrail storage capacity.




## Using the Contrail Storage User Interface

The Contrail storage solution provides a user interface integrated into the Contrail user interface. The storage solution user interface displays the following:

- Customer usable space, which is different from Ceph total space. The displayed usable space does not display the space used by replication and other Ceph functions.
- Monitor OSDs (disks), monitoring processes (MON), and state changes, enabling quick identification of resource failures within storage components.
- Total cluster I/O statistics and individual drive statistics.
- Ceph-specific information about each OSD (disk).
- Ceph logs, Ceph nodes, and Ceph alerts.


Select **Monitor > Infrastructure > Dashboard** to display an at-a-glance view of the system infrastructure components, including the numbers of virtual routers, control nodes, analytics nodes, config nodes, and storage nodes currently operational, and a bubble chart of storage nodes showing the Available (%) and Total Storage (GB). See the following figure.



Bubble charts use the following color-coding scheme for storage nodes:

- Blue—working as configured.
- Red—error, node is down.
- Yellow—one of the node disks is down.

Select **Monitor > Storage > Dashboard** to see a summary of cluster health, usage, pools, and disk status, and to gain insight into activity statistics for all nodes. See the following figure.



## Hardware Specifications

The following are additional hardware specifications needed for the Contrail storage solution.

Additional minimum specifications:

- Two 500 GB, 7200 RPM drives in the server 4 and server 5 cluster positions (those with the compute storage role) in the Contrail installation. This configuration provides 1 TB of clustered, replicated storage.

Recommended compute storage configuration:

- For every 4-5 HDD devices on one compute storage node, use one SSD device to provide the OSD journals for that set of HDD devices.

## Software Files for Compute Storage Nodes

The Contrail storage solution is only supported with the Ubuntu operating system.

For each compute storage node, ensure the following software is downloaded:

- The storage Debian package: **contrail-storage-packages\_x.xx-xx~xxxxxx\_all.deb**.
- NFS VM **qcow2** image from Juniper Networks.

## Contrail OpenStack Nova Modifications

Contrail's OpenStack Nova function has been modified to spawn both migratable and non-migratable virtual machines.

- Nova's typical virtual machine storage directory, `/var/lib/nova/instances`, is used for non-migratable virtual machine ephemeral storage.
- Contrail storage creates a new directory, `/var/lib/nova/instances/global`, used for the ephemeral storage for migratable virtual machines. The `/var/lib/nova/instances/global` must be mounted on a shared storage device (NFS with Contrail Storage), accessible from all the compute nodes.
- To start a non-migratable virtual machine with the Nova CLI command `nova boot`, the additional argument "`--meta storage_scope=local`" must be provided.
- To start a migratable virtual machine with `nova boot`, the additional argument "`--meta storage_scope=global`" must be provided. To force Nova and the Horizon UI to spawn migratable virtual machines by default, the storage scope must be set to global. This task is described in the next section.

## Installing the Contrail Storage Solution

The Contrail storage solution can be installed using the same tools used to install Contrail, either by using Fabric (fab) commands or by using the Contrail Server Manager.

Both installation methods are described in the following sections.

### *Installation Notes*

- When installing a base operating system on any compute storage node, the operating system must be installed only on a single drive. The other drives must be configured as individual devices, and should not be concatenated together in a logical volume manager (LVM) device.
- For best performance, it is recommended to use solid state devices (SSD) for the Ceph OSD journals. Each SSD device can provide OSD journal support to 3-6 HDD OSD devices, depending on the model of SSD device. Most SSD devices can support up to 4 HDDs, assuming the HDDs are running at capacity.

## Using Fabric Commands to Install and Configure Storage

Use the information in this section to install storage using Fabric (fab) commands.

When installing the operating system on a compute storage node, install the operating system on a single drive and leave all other drives as unbundled.

Installing the Contrail storage solution with Fabric commands gives the following:

- Base Ceph block device and object support.
- Easy configuration of SSD devices for OSD journals.
- Virtual machine migration support.
- Limited Cinder multi-backend support.

### *Cautions*

Before installing, ensure the following:

- Manually ensure that the UID or GID of the Nova user is identical on all compute nodes before provisioning any software.
- Manually ensure that the time is identical on all nodes by configuring NTP.

## Fabric Installation Procedure

This section provides guidelines and steps for using Fabric (fab) commands to install the Contrail storage solution. The installation is similar to a regular Contrail fab installation, however, you define additional storage information in the **testbed.py** file, including:

- Define new roles: **storage-master** and **compute-storage**.
  - Define how each additional non-root drive is used in the cluster.
  - Define potential additional virtual machine migration variables.
  - Copy and install the additional storage package to systems.
1. Install the storage Debian package on all nodes:  
**fab install\_storage\_pkg\_all:/YYYY/contrail-storage-package-XXX.deb**
  2. After using the **fab install\_contrail** command, use the **fab install\_storage** command.
  3. After using the **fab setup\_all** command, use the **fab setup\_storage** command.
  4. If you need to enable Contrail-based live virtual machine migration, use the **fab setup\_nfs\_livem** command or the **fab setup\_nfs\_livem\_global** command, as described in the following.



**NOTE:** If virtual machine migration is not needed, do not use either command.

- Use the **fab setup\_nfs\_livem** command to store the virtual machine's ephemeral storage on local drives.
  - Use the **fab setup\_nfs\_livem\_global** command to store the virtual machine's ephemeral storage within Contrail's storage (using Ceph). This command sets the cluster storage scope to global.
5. Add two new Contrail storage roles: **compute-storage** and **storage-master**.
    - Define the **storage-master** role on all nodes running OpenStack. Although Ceph has no notion of a master, define this role because Ceph must be run on the node that runs the OpenStack software. OpenStack nodes typically do not have any cluster storage defined, only local storage.
    - The **storage-compute** role is an add-on role. It means that compute nodes have the option of providing storage functionality. Standalone storage nodes are not supported.
  6. Change the **testbed.py** file details as needed for your environment.

In the base configuration, define the **storage\_node\_config** values, which gives device details. See the following example.

```
storage_node_config = {

    host2 : { 'disks' : ['/dev/sdb', '/dev/sdc'], 'ssd-disks': ['/dev/sdd',
'/dev/sde'], 'local-disks' : ['/dev/sdf', '/dev/sdh'], 'nfs' :
['10.87.140.156:/test', '10.87.140.156:/test1']},

    host3 : { 'disks' : ['/dev/sdb:/dev/sde', '/dev/sdc:/dev/sde',
'/dev/sdd:/dev/sde', '/dev/sdf:/dev/sdj', '/dev/sdg:/dev/sdj',
'/dev/sdh:/dev/sdj', '/dev/sdi:/dev/sdj', 'local-ssd-disks' : ['/dev/sdk',
'/dev/sdl']},}
```

Available device details parameters include:

- **disks** and **ssd-disks** are Ceph disks.
  - **local-disk** and **local-ssd-disks** are LVM disks.
  - **host2** in the example shows all the storage types that can be configured using Cinder multi-backend.
  - **disks** is a list of HDD disks used for a Ceph HDD pool.
  - **ssd-disks** is a list of SSD disks used for a Ceph SSD pool.
  - **local-disks** is a list of disks used for local LVM storage.
  - **nfs** is an NFS device.
  - In the example, **host3** is a more typical configuration.
  - **/dev/sde** and **/dev/sdj** are SSD disks that are used as OSD journals for other HDD drives.
  - **local-ssd-disks** is a list of disks used for local SSD LVM storage.
7. Add virtual machine migration as needed, using the following parameters.

```
live_migration = True
```

```
ceph_nfs_livevm = True
```

```
ceph_nfs_livem_subnet = '192.168.10.0/24' # Private subnet to be provided for live
migration VM
```

```
ceph_nfs_livem_image = '/Ubuntu/libmnfs.qcow2' # path of live migration qcow2 image.
This image is provided by Juniper Networks.
```

```
ceph_nfs_livem_host = host3 # host in which the NFS VM will run
```

For external NFS server-based live migration, use the following configuration.

```
live_migration = True
```

```
ext_nfs_livevm = True
```

```
ext_nfs_livem_mount = '10.10.10.10:/nfsmount' # External NFS server mount path
```



**NOTE:** When using an external NFS server, make sure the NFS server maps the uids and gids correctly, and provides read and write access for all the uids. If there is any issue related to the permission, either the VM launch errors out or the live migration fails with permission-related errors.

## Using Server Manager to Install and Configure Storage

This section provides notes and guidelines to install the storage solution using the Contrail Server Manager. Installing the Contrail Storage solution using Server Manager provides:

- Base Ceph block device and object support.
- Easy configuration of SSD journals.
- Support for live migration configuration.

Before installing the base operating system with Server Manager, ensure that the compute storage nodes have been configured with single operating system device installs.

### *Cautions*

- Virtual machine migration support uses a fixed IP address (192.168.101.3) for the **livernfs** virtual machine.
- There is no Cinder multi-backend support.
- There is no support for single server provisioning, the entire cluster must be provisioned.

## Server Manager Installation Procedure for Storage

This section provides notes and guidelines if you choose to install the storage solution using the Contrail Server Manager.

1. Upload the storage package: **server-manager add image -f <filename.json>**

where <filename.json> has content similar to the following example:

```
{
  "image": [
    {
      "id": "contrail-storage-packages_1.10-xx-xxxxxx_all",
      "parameters": "{}",
      "path":
        "/store/contrail-storage-packages_1.10-xx-xxxxxx_all.deb",
      "type": "contrail-storage-ubuntu-package",
      "version": "1.10-xx"
    },
  ],
}
```

```
]
}
```

2. Use the **ceph-authtool** command if you need to generate unique keys for administration, monitor, and OSD.

- a. To install **ceph-authtool** on CentOS, use the following command:

```
yum install
http://ceph.com/rpm/el6/x86_64/ceph-common-0.80.5-0.el6.x86_64.rpm
```

- b. To install **ceph-authtool** on Ubuntu:

```
apt-get install ceph-common
```

- c. Use the following command once for each key:

```
ceph-authtool --gen-print-key
```

- d. Add the generated keys to the **cluster.json** file:

```
cluster.json:

"storage_mon_secret":

"AQBDCdpTsB5FChAA0zI2++uosfmtj7tjmhPu0g==",

"osd_bootstrap_key":

"AQBKCdpTmN+HGRAA16rmStq5iYoPnANzSXLcXA==",

"admin_key":

"AQBLCdpu0S6FhAAFDW0SsdzyDAUeuwOr/h61A=="
```

- e. In Contrail Release 2.10 and later, add live-migration configuration to the **cluster.json** file, if you are using live migration.

```
"live_migration": "enable",

"live_migration_nfs_vm_host": "compute-node-01",

"live_migration_storage_scope": "global",
```

### Example: Configurations for Storage for Reimaging and Provisioning a Server

Use the following example configurations as guidelines for reimaging and provisioning a server for storage. Examples are given for configurations for releases prior to Release 2.10 and for configurations for Release 2.10 and later.

1. Define storage in the cluster. The following example configurations show new key-value pairs added to the configuration. The **cluster** section should appear similar to the following when storage is defined in a cluster.

Example: Storage and key-value pairs defined in releases prior to 2.10:

```
{
  "cluster" : [
    {
      "id" : "demo-cluster",
      "parameters" : {
        "router_asn": "<asn>",
        "database_dir": "/home/cassandra",
        "database_token": "",
        "use_certificates": "False",
        "multi_tenancy": "False",
        "encapsulation_priority": "MPLSoUDP,MPLSoGRE,VXLAN",
        "service_token": "<password>",
        "keystone_user": "admin",
        "keystone_password": "<password>",
        "keystone_tenant": "admin",
        "analytics_data_ttl": "168",
        "subnet_mask": "<ip address>",
        "gateway": "<ip address>",
        "password": "<password>",
        "haproxy": "disable",
        "external_bgp": "",
        "domain": "demo.company.net",
        "storage_mon_secret": "$ABC123",
        "osd_bootstrap_key": "$ABC123",
        "admin_key": "$ABC123"
      }
    }
  ]
}
```



Example: Storage and key-value pairs defined in releases 2.10 and later:

```
{
  "cluster" : [
    {
      "id" : "demo-cluster",
      "parameters" : {
        "router_asn": "<asn>",
        "database_dir": "/home/cassandra",
        "database_token": "",
        "use_certificates": "False",
        "multi_tenancy": "False",
        "encapsulation_priority": "MPLSoUDP,MPLSoGRE,VXLAN",
        "service_token": "<password>",
        "keystone_user": "admin",
        "keystone_password": "<password>",
        "keystone_tenant": "admin",
        "analytics_data_ttl": "168",
        "subnet_mask": "<ip address>",
        "gateway": "<ip address>",
        "password": "<password>",
        "haproxy": "disable",
        "external_bgp": "",
        "domain": "demo.company.net",
        "storage_mon_secret": "$ABC123",
        "osd_bootstrap_key": "$ABC123",
        "admin_key": "$ABC123",
        "live_migration" : "enable",
        "live_migration_nfs_vm_host": "compute-host-01",
        "live_migration_storage_scope": "global",
      }
    }
  ]
}
```

```
    ]  
  }  
}
```

2. Add the **disks** key, the **storage-compute** role value, and the **storage\_repo\_id** key.

- The **storage\_repo\_id** key must be added to servers with the **storage-master** or **storage-compute** roles.
- The **disks** key-value pair must be added to servers with the **storage-compute** roles.
- The **storage-master** value must be added to the **roles** key for the server that has the **storage-master** role.
- The **storage-compute** value must be added to the **roles** key for the servers that have the **storage-compute** role.

The following server section is an example, showing the **storage\_repo\_id** and **disks** keys, and the **storage-compute** and **storage-master** values.

In the example, one server contains the **storage-compute** role and has 3 HDD drives (**/dev/sdb**, **/dev/sdc**, **/dev/sdd**), supporting 3 OSDs.

Each OSD uses one partition of an SSD drive (**/dev/sde**) as its OSD journal.

The server manager software correctly partitions **/dev/sdd** and assign one partition to each OSD. The **storage\_repo\_id** contains the base name of the Contrail storage package which has been added as an image to Server Manager.

**Example: Server.json updates defined in releases earlier than 2.10:**

```
{ "server": [  
  {  
    "id": "demo2-server",  
    "mac_address": "<mac address>",  
    "ip_address": "<password>",  
    "parameters" : {  
      "interface_name": "eth1",  
      "compute_non_mgmt_ip": "",  
      "compute_non_mgmt_gway": "",  
      "storage_repo_id": "contrail-storage-packages",  
      "disks": ["/dev/sdb:/dev/sde", "/dev/sdc:/dev/sde",  
        "/dev/sdd:/dev/sde"]  
    },  
    "roles" :  
    ["config", "openstack", "control", "compute", "collector", "webui", "database", "storage-compute", "storage-master"],  
  }  
]
```

```

        "cluster_id": "demo-cluster",
        "subnet_mask": "<ip address>",
        "gateway": "1<ip address>",
        "password": "<password>",
        "domain": "demo.company.net",
        "email": "id@company.net"
    } ]
}

```

**Example: Server.json updates defined in releases 2.10 and later:**

```

Server.json :
{
  "server": [
    {
      "id": "demo2-server",
      "mac_address": "<mac address>",
      "ip_address": "<ip address>",
      "parameters" : {
        "interface_name": "eth1",
        "compute_non_mgmt_ip": "",
        "compute_non_mgmt_gway": "",
        "storage_repo_id": "contrail-storage-packages",
        "disks": ["/dev/sdb:/dev/sde", "/dev/sdc:/dev/sde",
"/dev/sdd:/dev/sde"]
      },
      "roles" :
["config", "openstack", "control", "compute", "collector", "webui", "database", "storage-compute", "storage-master"],
      "contrail": {
        "control_data_interface": "p3p2"
      },
      "network": {
        "interfaces": [
          {
            "default_gateway": "<ip address>",
            "dhcp": true,

```

```

        "ip_address": "<ip address>",
        "mac_address": "<mac address>",
        "member_interfaces": "",
        "name": "eth1",
        "tor": "",
        "tor_port": "",
        "type": "physical"
    },
    {
        "default_gateway": "<ip address>",
        "dhcp": "",
        "ip_address": "<ip address>",
        "mac_address": "<mac address>",
        "member_interfaces": "",
        "name": "p3p2",
        "tor": "",
        "tor_port": "",
        "type": "physical"
    }
],
"management_interface": "eth1"
},
"cluster_id": "demo-cluster",
"subnet_mask": "<ip address>",
"gateway": "<ip address>",
"password": "<password>",
"domain": "demo.company.net",
"email": "id@company.net"
} ]
}

```

3. Use the following commands to provision the entire cluster:

```

# /opt/contrail/server_manager/client/server-manager -c

# /opt/contrail/server_manager/smgr_config.ini provision --cluster_id test-cluster
contrail_test_pkg

```

## Storage Installation Limits

### General Limitations

- Minimum number of storage nodes to configure: 2
- The number of storage nodes should always be an even number (2, 4, 12, 22, etc.).

### Fab Storage Install Limitations

There are no additional limitations to installation when using fab commands.

### Server Manager Storage Install Limitations

- There is no integrated way to add OSDs or drives to a storage node.
- There is no integrated way to add new storage nodes to a cluster.
- Provisioning a single server is not supported. You can add a server to Server Manager and then provision the entire cluster.
- The live migration overlay network is preset to use 192.168.101.0/24.
- The user must copy the image **livemnfs.qcow2.gz** to the folder **/var/www/html/contrail/images** before provisioning live migration using Server Manager.



## CHAPTER 6

# Upgrading Contrail Software

- Upgrading Contrail Software on page 109
- DKMS for vRouter Kernel Module on page 112

## Upgrading Contrail Software

---

Use the following procedure to upgrade an installation of Contrail software from one release to a more recent release. This procedure is valid for upgrading Contrail Release 2.21 and later to Contrail Release 3.1.



**NOTE:** If you are installing Contrail for the first time, refer to the full documentation and installation instructions in “Installing the Operating System and Contrail Packages” on page 13.

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

To upgrade Contrail software from Contrail Releases 2.21 and later to Release 3.1:

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

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

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



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

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

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

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

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

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

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

See [“Setting Up the Testbed Definitions File” on page 16](#).

5. In release packages prior to Contrail Release 3.0, the packaged Cassandra version is 1.2.11. In the 3.0 release, the packaged Cassandra version is 2.1.9. Upgrading Cassandra from 1.2.11 to 2.1.9 directly is not supported by Cassandra. For more information, refer to [DataStax Upgrade Guide, Cassandra 2.1.x restrictions](#).

The `fab upgrade_contrail` command sequence enables upgrading Cassandra from 1.2.11 to 2.1.9 by performing necessary intermediate upgrades. Consequently, during the Contrail upgrade procedure (`fab upgrade_contrail`), the Cassandra SSTables are upgraded, which takes a long time if the Cassandra data is huge (usually because the Contrail Analytics keyspace is huge).

There is an option to minimize upgrade down time by dropping the Contrail Analytics keyspace before the upgrade, by issuing the following `fab` command:

`fab drop_analytics_keyspace`

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

Change directory to the `utils` folder:

`cd /opt/contrail/utils; \`

Select the correct upgrade procedure from the following to match your operating system and vRouter. In the following, `<from>` refers to the currently installed release number, such as 2.21.2, 2.22.2, 3.0.2.0, and so on:

*CentOS Upgrade Procedure:*

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

*Ubuntu 14.04 Upgrade, Two Procedures:*

There are two different upgrade procedures for the upgrade to Contrail Release 3.1, depending on which vRouter (**3.13.0-X-generic** or **contrail-vrouter-dkms**) is installed in your current setup.

In Contrail Release 3.0.2.0 and later, the recommended kernel version for an Ubuntu 14.04-based system is 3.13.0-85-. Both procedures can use the command `fab upgrade_kernel_all` to upgrade the kernel.



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

Use the following upgrade procedure for Contrail systems based on Ubuntu 14.04 without the **contrail-vrouter-3.13.0-85-generic** installed. The command sequence upgrades the kernel version and also reboots the compute nodes when finished.

```
fab
install_pkg_all:/tmp/contrail-install-packages-x.x.x.x-xx~<openstack_version>_all.deb;

fab migrate_compute_kernel;

fab
upgrade_contrail:<from>,/tmp/contrail-install-packages-x.x.x.x-xx~<openstack_version>_all.deb;

fab upgrade_kernel_all;

fab restart_openstack_compute;
```

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

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

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

All nodes in the cluster can be upgraded to kernel version 3.13.0-85 by using the following **fab** command:

```
fab upgrade_kernel_all
```

#### 7. (For Contrail Storage option, only.)

Contrail Storage has its own packages.

To upgrade Contrail Storage, download the file:

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

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

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

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

Use the following statement to upgrade the software:

```
cd /opt/contrail/utlis; \
```

```
Ubuntu: fab
```

```
upgrade_storage:<from>,/tmp/contrail-storage-packages_x.x.x.x-xx~<openstack_version>_all.deb;
```

## DKMS for vRouter Kernel Module

---

Dynamic Kernel Module Support (DKMS) is a framework provided by Linux to automatically build out-of-tree driver modules for Linux kernels whenever the Linux distribution upgrades the existing kernel to a newer version.

In Contrail, the vRouter kernel module is an out-of-tree, high performance packet forwarding module that provides advanced packet forwarding functionality in a reliable and stable manner. Contrail provides a DKMS-compatible source package for Ubuntu so that if you deploy an Ubuntu-based Contrail system you do not need to manually compile the kernel module each time the Linux deployment gets upgraded.

The **contrail-vrouter-dkms** package provides the DKMS compatibility for Contrail. Prior to installing the **contrail-vrouter-dkms** package, you must install both the DKMS package and the **contrail-vrouter-utils** package, because the **contrail-vrouter-dkms** package is dependent on both. Installing the **contrail-vrouter-dkms** package adds the vRouter sources to the DKMS database, builds the vRouter module, and installs it in the existing kernel modules tree. When a kernel upgrade occurs, DKMS ensures that the module is compiled for the newer kernel and installed in the proper location so that upon reboot, the newer module can be used with the upgraded kernel.

For more information about DKMS, refer to:

- DKMS Ubuntu documentation at <https://help.ubuntu.com/community/DKMS>
- DKMS Ubuntu manual pages at <http://manpages.ubuntu.com/manpages/lucid/man8/dkms.8.html>
- Linux Journal article on DKMS at <http://www.linuxjournal.com/article/6896>

## PART 3

# Configuring Contrail

- [Configuring Virtual Networks on page 115](#)
- [Example of Deploying a Multi-Tier Web Application Using Contrail on page 157](#)
- [Configuring Services on page 169](#)
- [Configuring Service Chaining on page 191](#)
- [Examples: Configuring Service Chaining on page 227](#)



## CHAPTER 7

# Configuring Virtual Networks

- [Creating Projects in OpenStack for Configuring Tenants in Contrail on page 116](#)
- [Creating Virtual Networks and Policies in Juniper Networks Contrail on page 117](#)
- [Creating Virtual Networks and Policies in OpenStack Contrail on page 125](#)
- [Creating an Image and Launching a Virtual Machine on page 133](#)
- [Creating a Floating IP Address Pool and Allocating it to a Virtual Machine on page 138](#)
- [Using Security Groups with Virtual Machines \(Instances\) on page 142](#)
- [Support for IPv6 Networks in Contrail on page 145](#)
- [Configuring EVPN and VXLAN on page 148](#)

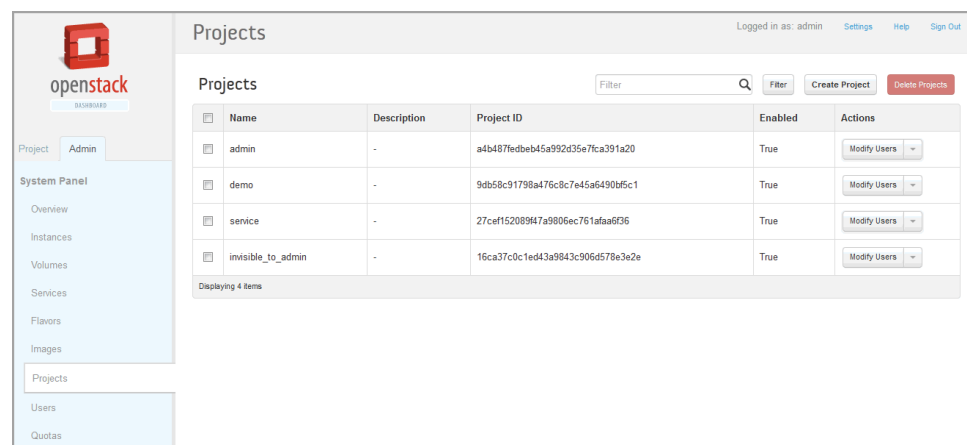
## Creating Projects in OpenStack for Configuring Tenants in Contrail

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 8 on page 116](#).

**Figure 8: OpenStack Projects**



2. In the upper right, click the **Create Project** button to access the **Add Project** window; see [Figure 9 on page 116](#).

**Figure 9: Add Project**

**Add Project**

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.

**Related  
Documentation**

- [Creating Virtual Networks and Policies in Juniper Networks Contrail on page 117](#)
- [Creating Virtual Networks and Policies in OpenStack Contrail on page 125](#)
- [OpenStack documentation](#)

---

## Creating Virtual Networks and Policies in Juniper Networks Contrail

---

- [Creating a Virtual Network—Juniper Networks Contrail on page 118](#)
- [Deleting a Virtual Network with Juniper Networks Contrail on page 120](#)
- [Creating a Network Policy with Juniper Networks Contrail on page 121](#)
- [Associating a Network to a Policy—Juniper Networks Contrail on page 123](#)

Creating a Virtual Network—Juniper Networks Contrail

Contrail makes creating a virtual network very easy for a self-service user. You create networks and network policies at the user dashboard, then associate policies with each network. The following procedure shows how to create a virtual network when using Juniper Networks Contrail.

- 1. Before creating a virtual network, create an IP address management (IPAM) for your project. Select **Configure > Networking > IP Address Management**, then click the **Create** button.

The **Add IP Address Management** window appears, see [Figure 10 on page 118](#).

Figure 10: Add IP Address Management

The screenshot shows a web-based form titled "Add IP Address Management". It includes the following fields:

- Name:** A text input field with the placeholder text "IPAM Name".
- DNS Method:** A dropdown menu currently set to "Default".
- NTP Server IP:** An empty text input field.
- Domain Name:** An empty text input field.

At the bottom right of the form are two buttons: "Cancel" and "Save". A vertical identifier "s041838" is visible on the right side of the window frame.

- 2. Complete the fields in **Add IP Address Management**: The fields are described in [Table 11 on page 118](#).


Table 11: Add IP Address Management Fields

Field	Description
Name	Enter a name for the IPAM you are creating.
DNS Method	Select from a drop-down list the domain name server method for this IPAM: <b>Default</b> , <b>Virtual DNS</b> , <b>Tenant</b> , or <b>None</b> .
NTP Server IP	Enter the IP address of an NTP server to be used for this IPAM.
Domain Name	Enter a domain name to be used for this IPAM.

- 3. Select **Configure > Networking > Networks** to access the **Configure Networks** screen; see [Figure 11 on page 119](#).



Figure 11: Configure Networks




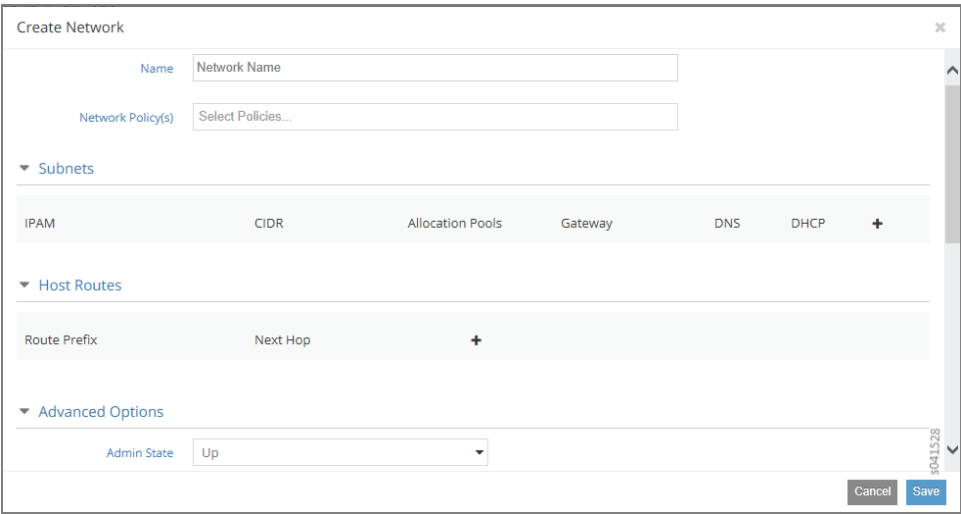
4. Verify that your project is displayed as active in the upper right field, then click the  icon. The **Create Network** window is displayed. See [Figure 12 on page 119](#). Use the scroll bar to access all sections of this window.

Figure 12: Create Network



5. Complete the fields in the **Create Network** window with values that identify the network name, network policy, and IP options as needed. See field descriptions in [Table 12 on page 119](#).

Table 12: Create Network Fields

Field	Description
Name	Enter a name for the virtual network you are creating.
Network Policy(s)	Select the policy to be applied to this network from the drop-down list of available policies. You can select more than one policy by clicking each one needed.

Table 12: Create Network Fields (*continued*)

Field	Description
<b>Subnets</b>	Use this area to identify and manage subnets for this virtual network. Click the + icon to open fields for IPAM, CIDR, Allocation Pools, Gateway, DNS, and DHCP. Select the subnet to be added from a drop down list in the IPAM field. Complete the remaining fields as necessary. You can add multiple subnets to a network. When finished, click the + icon to add the selections into the columns below the fields. Or click the - icon to remove the selections.
<b>Host Routes</b>	Use this area to add or remove host routes for this network. Click the + icon to open fields where you can enter the Route Prefix and the Next Hop. Click the + icon to add the information, or click the - icon to remove the information.
<b>Advanced Options</b>	Use this area to add or remove advanced options, including identifying the Admin State as Up or Down, to identify the network as Shared or External, to add DNS servers, or to define a VxLAN Identifier.
<b>Floating IP Pools</b>	Use this area to identify and manage the floating IP address pools for this virtual network. Click the + icon to open fields where you can enter the Pool Name and Projects. Click the + icon to add the information, or click the - icon to remove the information.
<b>Route Target(s)</b>	Move the scroll bar down to access this area, then specify one or more route targets for this virtual network. Click the + icon to open fields where you can enter route target identifiers. Click the + icon to add the information, or click the - icon to remove the information.

- To save your network, click the **Save** button, or click **Cancel** to discard your work and start over.


Now you can create a network policy, see [“Creating a Network Policy with Juniper Networks Contrail” on page 121](#).

## Deleting a Virtual Network with Juniper Networks Contrail

You can delete any of the virtual networks in your system. However, you must first disassociate any virtual machines (instances) that are associated with that network. Use OpenStack to view and delete the virtual machines associated with a virtual network, see [“Deleting a Virtual Network with OpenStack Contrail” on page 127](#). When you have deleted the virtual machines associated with a virtual network, you can delete the network in OpenStack, or you can delete the network in Juniper Networks Contrail, using the following procedure.

- To view the virtual networks in the current project, select **Configure > Networks**. The **Configure Networks** window is displayed. See [Figure 13 on page 121](#).

Figure 13: Configure Networks



2. Select the network you want to delete, then click the Delete (trashcan) icon at the top right. A confirm window is displayed.
3. Click **Confirm** to delete the network, or click **Cancel** to quit the delete activity.

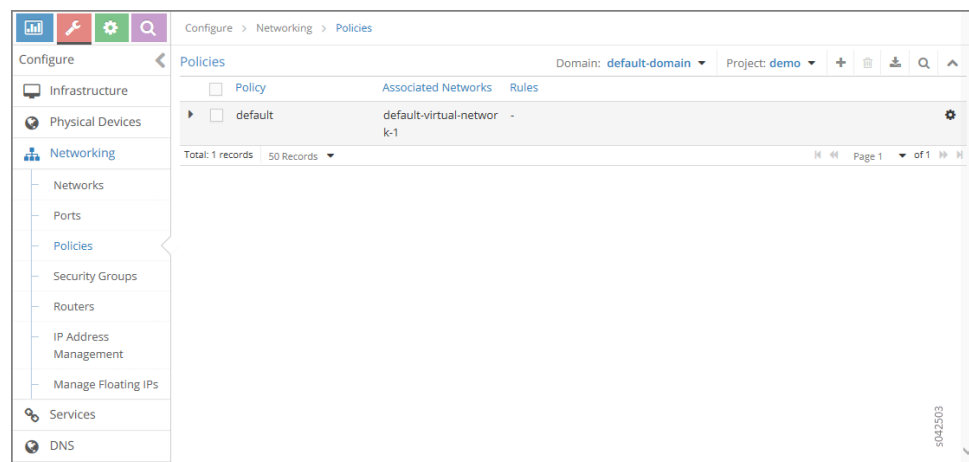
## Creating a Network Policy with Juniper Networks Contrail

The Contrail Controller makes creating network traffic policies very simple. You work from the self-service user interface to define a policy, then define a rule or rules to be applied in that policy. You can define such parameters as the type and direction of traffic for the rule, the source and destination of that traffic, traffic originating from or destined for specific ports, the sequence in which to apply a rule, and so on.

To create a network policy when using Juniper Networks Contrail:

1. In the Contrail Web user interface, select **Configure > Networking > Policies**. The **Policies** window is displayed. See [Figure 14 on page 121](#).

Figure 14: Policies Window



2. Click the + icon.

The **Create Policy** window is displayed. See [Figure 15 on page 122](#). Click the + icon in the Create Policy window.

**Figure 15: Create Policy Window**

3. Enter the policy name and select the values from the menus in the **Create Policy** window. [Table 13 on page 122](#) describes the selections.

**Table 13: Create Policy Fields**

Field	Description
<b>Name</b>	Enter a name for the policy you are creating.
<b>Policy Rules</b>	Use this area to define the rules for the policy you are creating. Click the + (plus sign) to open up the fields for defining the rules. Click the – (minus sign) to delete any rule. Multiple rules can be added to a policy. Each policy rule field is described in the following table rows.
<b>Action</b>	Define the action to take with traffic that matches the current rule. Select from a list: <b>Pass, Deny</b> .
<b>Protocol</b>	Define the protocol associated with traffic for this policy rule. Select from a list of available protocols (or <b>ANY</b> ): <b>ANY, TCP, UDP, ICMP</b> .
<b>Source</b>	Select the source network for traffic associated with this policy rule. Choose <b>ANY</b> or select from the menu list of all available sources. Sources are displayed in the form: <i>domain-name:project-name:network-name</i> .
<b>Ports</b>	Use this field to specify that traffic from particular source ports are associated with this policy rule. Identify traffic from <b>any</b> port or enter a specific port, a list of ports separated with commas, or a range of ports in the form <i>nnnn-nnnnn</i> .
<b>Direction</b>	Define the direction of traffic to match the rule. For traffic moving in and out, select <> (bidirectional). For traffic moving in one direction, select > (unidirectional).
<b>Destination</b>	Select the destination network for traffic to match this rule. Choose <b>ANY</b> or select from the menu of all available destinations. Destinations are displayed in the form: <i>domain-name:project-name:network-name</i> .
<b>Destination</b>	Select the destination port for traffic to match this rule. Enter <b>ANY</b> for any destination port or enter a specific port, a list of ports separated with commas, or a range of ports in the form <i>nnnn-nnnnn</i> .

Table 13: Create Policy Fields (*continued*)

Field	Description
<b>Services</b>	Select this check box to open a field where you can select from a list of available services to apply to this policy. The services are applied in the order in which they are selected. There is a restricted set of options that can be selected when applying services. For more information about services, see <a href="#">“Service Chaining” on page 191</a> .
<b>Mirror</b>	Select this check box to open a field where you can select from the list of configured services that you want to mirror in this policy. You can select a maximum of two services to mirror. For more information about mirroring, see <i>Configuring Traffic Analyzers and Packet Capture for Mirroring</i> .

- When you are finished selecting the rules for this policy, click **Save**.

The policy you just defined is displayed in the **Policy** column.

Next, you can associate the policy to a network, see [“Associating a Network to a Policy—Juniper Networks Contrail” on page 123](#).

## Associating a Network to a Policy—Juniper Networks Contrail

- [Associating Network Policies Overview on page 123](#)
- [Associating a Network Policy to a Network on page 123](#)

### Associating Network Policies Overview

Contrail helps you create and manage virtual networks (VNs). By default, all traffic in a VN is isolated to that VN. Traffic can only leave a VN by means of network policies that are defined for the VN.

This procedure shows how to associate a network policy with a network, using the Juniper Networks Contrail interface.

If you did not associate an existing network policy when you created your virtual network, you can use the **Network Policy(s)** field in the **Edit Network** window, or you can use the **Associate Networks** field in the **Edit Policy** window to associate or disassociate network policies with networks. The following procedures demonstrate both methods.

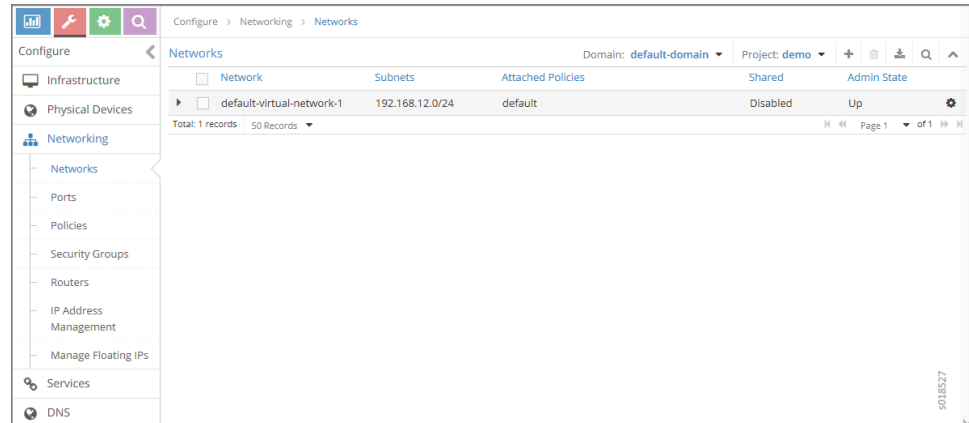
### Associating a Network Policy to a Network

This procedure shows how to attach (associate) a network policy to a network when starting from the **Edit Network** window.

1. Select **Configure > Networking > Networks**; see [Figure 16 on page 124](#).

Make sure your project is the active project in the upper right.

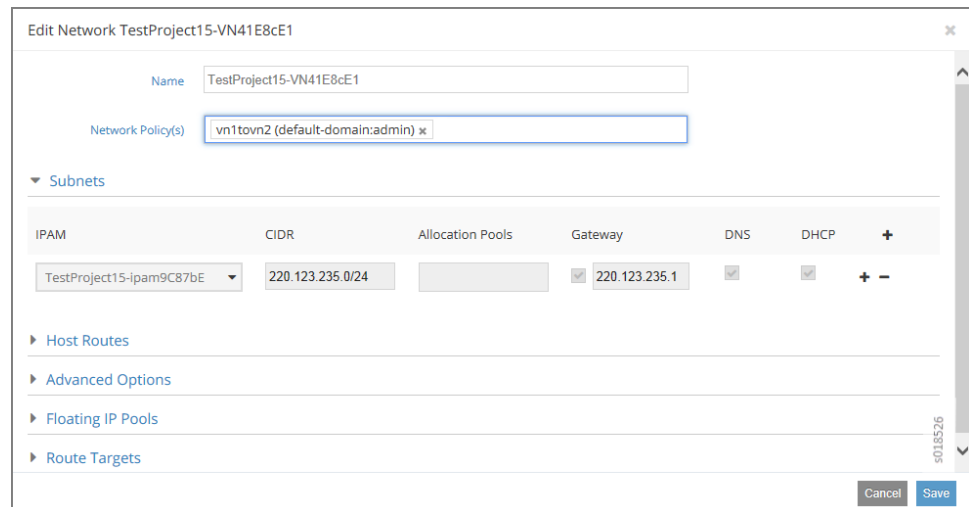
**Figure 16: Configure > Networking > Networks**



2. Select the network you want to associate with a policy, then in the **Action** column, click the gear wheel icon and select **Edit**.

The **Edit Network** window for the selected network is displayed; see [Figure 17 on page 124](#).

**Figure 17: Edit Network**



3. Click the **Network Policy(s)** field to show a list of existing policies, and then select a policy to associate with the selected network.

You can also disassociate a selected policy by clicking the - next to its name when it appears configured in the **Network Policy(s)** field.

4. When you are finished, click **Save**, or click **Cancel** to undo your selections.

#### Related Documentation

- [Creating an Image and Launching a Virtual Machine on page 133](#)

- Creating a Floating IP Address Pool and Allocating it to a Virtual Machine on page 138

## Creating Virtual Networks and Policies in OpenStack Contrail


- Creating a Virtual Network—OpenStack Contrail on page 125
- Deleting a Virtual Network with OpenStack Contrail on page 127
- Creating a Network Policy with OpenStack Contrail on page 129
- Associating a Network to a Policy—OpenStack Contrail on page 131

### Creating a Virtual Network—OpenStack Contrail

Contrail makes creating a virtual network very easy for you. You create networks and network policies at the user dashboard, then associate policies with each network. The following procedure shows how to create a virtual network when using OpenStack.

1. Select **Project > Other > Networking**. The **Networks** window is displayed. See [Figure 18 on page 125](#).

Figure 18: Networks Window



2. Verify that the correct project is displayed in the **Current Project** box, then click **Create Network**. The **Create Network** window is displayed. See [Figure 19 on page 125](#) and [Figure 20 on page 126](#).

Figure 19: Create Network Window

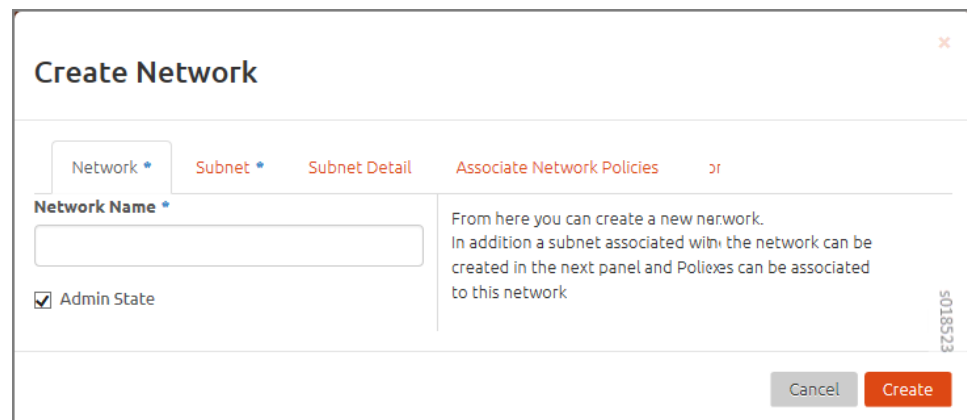


Figure 20: Create Network Window Subnet Tab

3. Click the **Network**, **Subnet**, **Subnet Detail**, and **Associate Network Policies** tabs to complete the fields in the **Create Network** window. See field descriptions in [Table 14 on page 126](#).

Table 14: Create Network Fields

Field	Description
<b>Network Name</b>	Enter a name for the network.
<b>Subnet Name</b>	Enter a name for the subnetwork.
<b>IPAM</b>	<p>Select the IPAM associated with the IP block.</p> <p>For new projects, an IPAM can be added while creating the virtual network. VM instances created in this virtual network are assigned an address from this address block automatically by the system when a VM is launched.</p>
<b>Network Address</b>	Enter the network address in CIDR format.
<b>IP Version*</b>	Select IPv4 or IPv6.
<b>Gateway IP</b>	Optionally, enter an explicit gateway IP address for the IP address block. Check the Disable Gateway box if no gateway is to be used.



Table 14: Create Network Fields (*continued*)

Field	Description
<b>Network Policy</b>	Any policies already created are listed. To select a policy, click the check box for the policy.

- Click the **Subnet Details** tab to specify the Allocation Pool, DNS Name Servers, and Host Routes.
- Click the **Associate Network Policies** tab to associate policies to the network.
- To save your network, click **Create Network**, or click **Cancel** to discard your work and start over.


## Deleting a Virtual Network with OpenStack Contrail

You can delete any of the virtual networks in your system. However, you must first disassociate any virtual machines (instances) that are associated with that network.

To delete a virtual network when using OpenStack Contrail:

- To view virtual machines that are associated with a virtual network, in the OpenStack module, select **Project > Other > Networking**. The **Networks** window is displayed. See [Figure 21 on page 127](#).

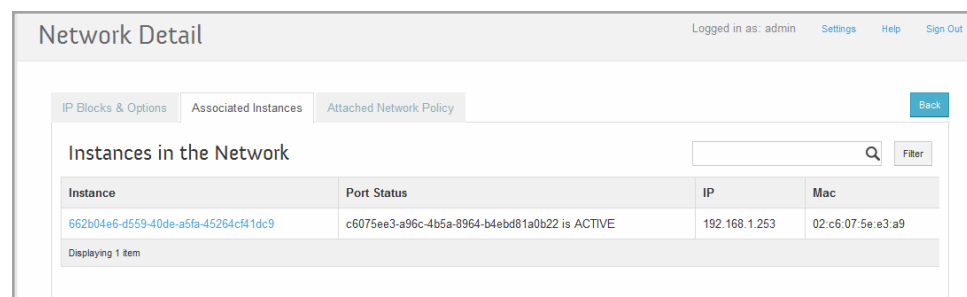
Figure 21: OpenStack Networks



- In the **Networks** window, select the network to be deleted.

The **Network Detail** window appears; see [Figure 22 on page 127](#).

Figure 22: OpenStack Network Detail, Associated Instances Tab



- Click the **Associated Instances** tab to see the instances associated with this network. Make note of the IP addresses of any instances that are associated with this network.

4. On the **Project** tab, select **Instances**.

The **Instances** window appears, displaying the instances associated with the current project; see [Figure 23 on page 128](#).

**Figure 23: Instances**

<input type="checkbox"/>	Instance Name	IP Address	Size	Keypair	Status	Task	Power State	Actions
<input type="checkbox"/>	be2	192.168.2.253	m1.tiny   512MB RAM   1 VCPU   0 Disk	-	Active	None	Running	Create Snapshot
<input type="checkbox"/>	fe1	192.168.1.253	m1.tiny   512MB RAM   1 VCPU   0 Disk	-	Active	None	Running	Create Snapshot

Displaying 2 items

5. On the **Instances** window, click the check box for any instance that is associated with the network that you want to delete, then click **Terminate Instances** to delete the instance.
6. When all instances that are associated with the network to be deleted have been terminated, delete the network.

To delete a network, return to the **Networks** page (see [Figure 21 on page 127](#)), select the network to be deleted, then click **Delete Networks** in the upper right.

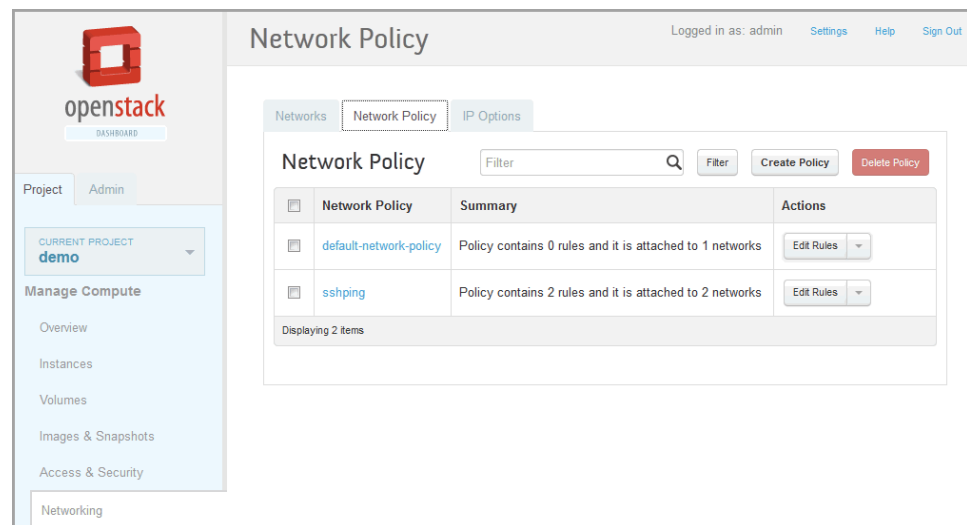
## Creating a Network Policy with OpenStack Contrail

Contrail makes creating network traffic policies very simple. You work from the self-service user interface to define a policy, then define a rule or rules to be applied in that policy. You can define such parameters as the type and direction of traffic for the rule, the source and destination of that traffic, traffic originating from or destined for specific ports, the sequence in which to apply a rule, and so on.

To create a network policy when using OpenStack:

1. On the OpenStack dashboard, make sure your project is displayed in the **Current Project** box, click **Networking**, and then click the **Network Policy** tab to display the **Network Policy** window; see [Figure 24 on page 129](#).

Figure 24: Network Policy



2. Click **Create Policy** at the upper right.

The **Create Network Policy** window is displayed; see [Figure 25 on page 129](#).

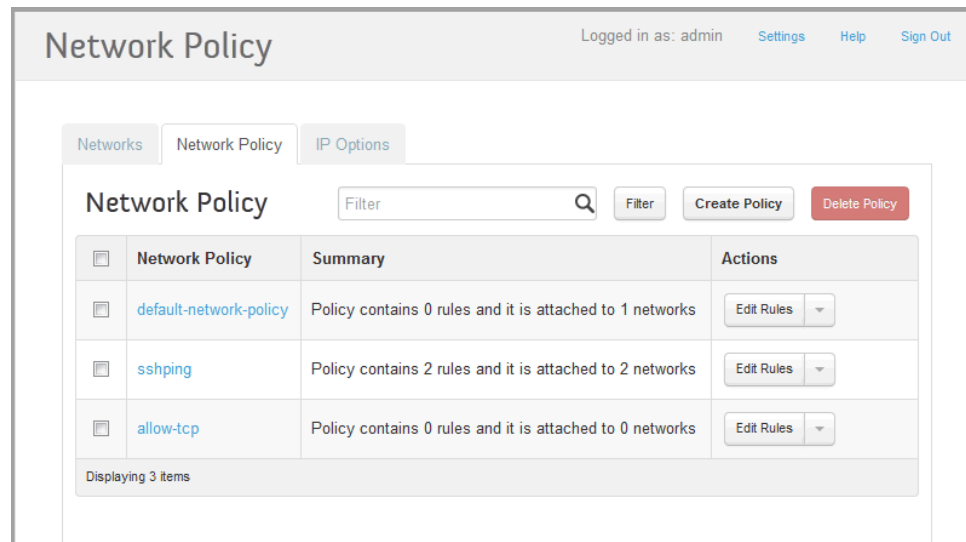
Figure 25: Create Network Policy

3. Enter a name and a description for this policy. Names cannot include spaces.

- When finished, click **Create Policy** on the lower right.

Your policy is created and it appears in the **Network Policy** window; see [Figure 26 on page 130](#).

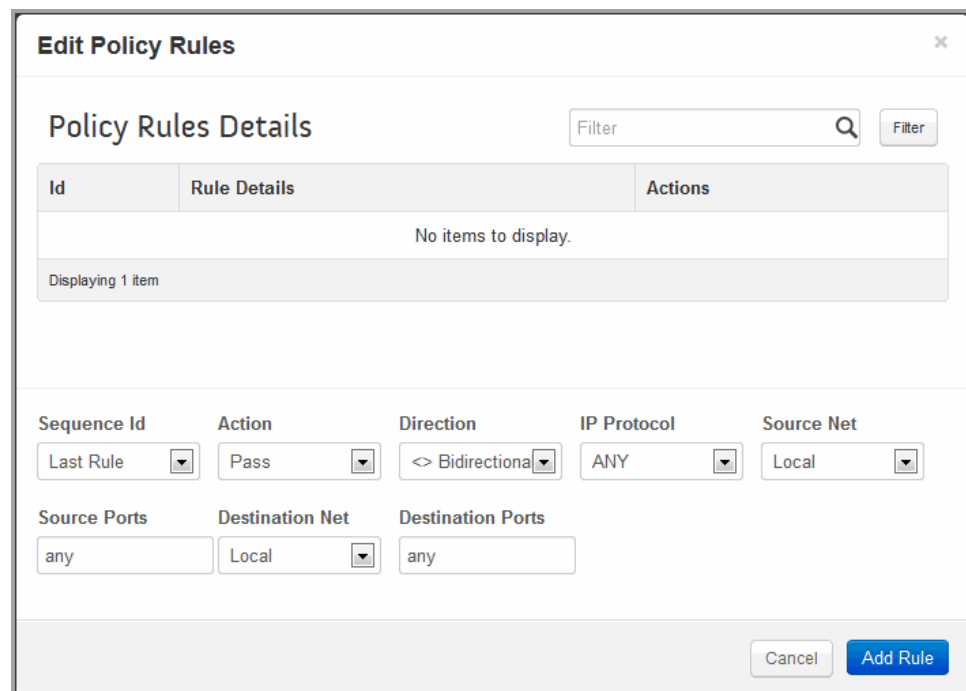
**Figure 26: Network Policy**



- In the **Network Policy** window, select the check box for your new policy, then click **Edit Rules** for that policy.

The **Edit Policy Rules** window is displayed; see [Figure 27 on page 130](#).

**Figure 27: Edit Policy Rules**



6. Define the rules for your policy, using the guidelines in [Table 15 on page 131](#).

**Table 15: Edit Policy Rules Fields**

Field	Description
<b>Policy Rules Details</b>	This section of the window displays any rules that have already been created for this policy.
<b>Id</b>	Displays a sequential number identifier for each rule within a policy.
<b>Rule Details</b>	Displays a description of the rule on this line.
<b>Actions</b>	Available actions for the rule on this line appear in this column. Currently you can use the <b>Delete</b> button in this column to delete a rule.
<b>Sequence Id</b>	This field lets you define the order in which to apply the current rule. Select from a list: <b>Last Rule, First Rule, After Rule</b> .
<b>Action</b>	Define the action to take with traffic that matches the current rule. Select from a list: <b>Pass, Deny</b> .
<b>Direction</b>	Define the direction in which to apply the rule, for example, to traffic moving in and out, or only to traffic moving in one direction. Select from a list: <b>Bidirectional, Unidirectional</b> .
<b>IP Protocol</b>	Select from a list of available protocols (or <b>ANY</b> ): <b>ANY, TCP, UDP, ICMP</b> ,
<b>Source Net</b>	Select the source network for this rule. Choose <b>Local</b> (any network to which this policy is associated), <b>Any</b> (all networks created under the current project) or select from a list of all sources available displayed in the list, in the form: <i>domain-name:project-name:network-name</i> .
<b>Source Ports</b>	Accept traffic from <b>any</b> port or enter a specific port, a list of ports separated with commas, or a range of ports in the form <i>nnnn-nnnnn</i> .
<b>Destination Net</b>	Select the destination network for this rule. Choose <b>Local</b> (any network to which this policy is associated), <b>Any</b> (all networks created under the current project) or select from a list of all destinations available displayed in the list, in the form: <i>domain-name:project-name:network-name</i> .
<b>Destination Ports</b>	Send traffic to <b>any</b> port or enter a specific port, a list of ports separated with commas, or a range of ports in the form <i>nnnn-nnnnn</i> .

7. After you select the rules for this policy, click **Add Rule** on the lower right of the **Edit Policy Rules** window.

Next, you can associate the policy to a network, see [“Associating a Network to a Policy—OpenStack Contrail” on page 131](#).

## Associating a Network to a Policy—OpenStack Contrail

- [Associating Network Policies Overview on page 132](#)
- [Associating a Network Policy to a Network on page 132](#)

## Associating Network Policies Overview

Contrail helps you create and manage virtual networks (VNs). By default, all traffic in a VN is isolated to that VN. Traffic can only leave a VN by means of network policies that are defined for the VN.


This procedure shows how to associate a network policy with a network when using OpenStack.

## Associating a Network Policy to a Network

1. Using the OpenStack Networking module, select the **Project** tab and click **Networking**.

The **Networks** window is displayed; see [Figure 28 on page 132](#).

**Figure 28: Networks Screen**

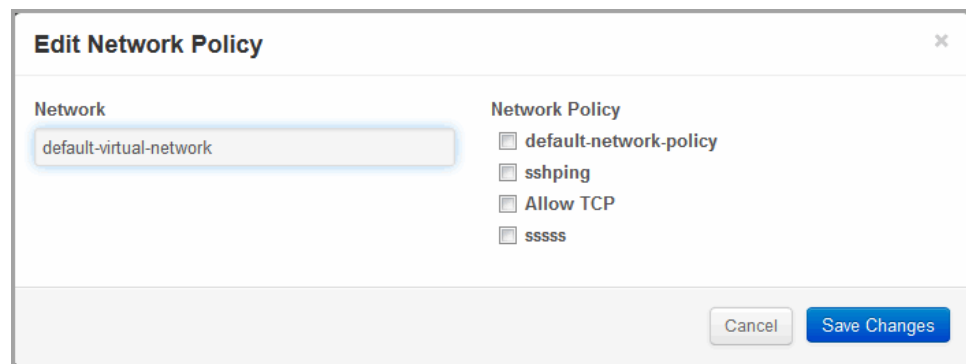


2. Click the check box to select the network you want to associate with a policy, then click the drop-down box in the **Actions** column and select **Edit Policy**.

The **Edit Network Policy** window is displayed; see [Figure 29 on page 132](#).

Available network policies are listed in the **Edit Network Policy** window.

**Figure 29: Edit Network Policy**



3. Click the check box of any policies to be associated with the selected network.
4. When finished, click **Save Changes**.

### Related Documentation

- [Creating an Image and Launching a Virtual Machine on page 133](#)
- [Creating a Floating IP Address Pool and Allocating it to a Virtual Machine on page 138](#)

## Creating an Image and Launching a Virtual Machine

- [Creating an Image on page 133](#)
- [Launching a Virtual Machine \(Instance\) on page 136](#)

### Creating an Image

You can use the OpenStack dashboard to specify an image to upload to the Image Service for a project in your system.

1. In OpenStack, select **Project > Compute > Images**. The Images window is displayed. See [Figure 30 on page 133](#).

Figure 30: 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 31 on page 134](#).

Figure 31: 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

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

Cancel

Create Image

4. Complete the fields to specify your image. [Table 16 on page 134](#) describes each of the fields on the screen.



**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 16: Create An Image Fields

Field	Description
Name	Required field. Enter a name for this image.



Table 16: Create An Image Fields (*continued*)

Field	Description
<b>Description</b>	Enter a description for the image.
<b>Image Source</b>	<p>Select <b>Image File</b> or <b>Image Location</b>.</p> <p>If you select <b>Image File</b>, you are prompted to browse to the local location of the file.</p>
<b>Image Location</b>	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.
<b>Format</b>	<p>Required field. Select the format of the image from a list:</p> <ul style="list-style-type: none"> <li>AKI- Amazon Kernel Image</li> <li>AMI- Amazon Machine Image</li> <li>ARI- Amazon Ramdisk Image</li> <li>ISO- Optical Disk Image</li> <li>QCOW2-QEMU Emulator</li> <li>Raw</li> <li>VDI</li> <li>VHD</li> <li>VMDK</li> </ul>
<b>Architecture</b>	Enter the architecture.
<b>Minimum Disk (GB)</b>	Enter the minimum disk size required to boot the image. If you do not specify a size, the default is 0 (no minimum).
<b>Minimum Ram (MB)</b>	Enter the minimum RAM required to boot the image. If you do not specify a size, the default is 0 (no minimum).
<b>Public</b>	Check the box if this is a public image. Leave unchecked for a private image.
<b>Protected</b>	Check the box for a protected image.

5. When you are finished, click **Create Image**.

## Launching a Virtual Machine (Instance)

After you have created virtual networks for your project, you can create and launch virtual machines. Virtual networks (VNs) are populated with virtual machines (VMs), also called instances. A VM is a simulation of a physical machine, such as a workstation or a server, that runs on a host that supports virtualization. Many VMs can run on the same host, sharing the host's resources. A VM has its own operating system image that can be different from that of other VMs running on the same host.

To use the OpenStack module to define and launch VMs (instances):

1. On the OpenStack dashboard **Project** tab, make sure your project is selected in the left column in the **Current Project** box, then click **Instances**.

The **Instances** pane is displayed, displaying all instances (VMs) currently in the selected project; see [Figure 32 on page 136](#).

**Figure 32: OpenStack Instances**

Instance Name	IP Address	Size	Keypair	Status	Task	Power State	Actions
be1	192.168.2.253	m1.tiny   512MB RAM   1 VCPU   0 Disk	-	Active	None	Running	Create Snapshot
fs1	192.168.1.253	m1.tiny   512MB RAM   1 VCPU   0 Disk	-	Active	None	Running	Create Snapshot

Displaying 2 items

2. To create and launch a new instance, click **Launch Instance** in the upper-right corner.

The **Launch Instance** window is displayed, where you can define and launch a new instance.

Figure 33: Launch Instance, Details Tab

Launch Instance

Details

Access & Security

Networking

Volume Options

Post-Creation

Instance Source

Image

Image

redmine-db

Instance Name

redmine-db-10

Flavor

m1.tiny

Instance Count

1

Compute Hostname

compute-nodea33

Specify the details for launching an instance.

The chart below shows the resources used by this project in relation to the project's quotas.

Flavor Details

Name	m1.tiny
VCPUs	1
Root Disk	0 GB
Ephemeral Disk	0 GB
Total Disk	0 GB
RAM	512 MB

Project Quotas

Number of Instances (4)

99,996 Available

Compute Host to launch Instance on

99,994 Available

Total RAM (9,216 MB)

9,990,784 MB Available

Cancel

Launch

3. Make sure the **Details** tab is active; see [Figure 33 on page 137](#), then define your instance using the fields shown in [Table 17 on page 137](#)

Table 17: Launch Instance Details Tab Fields

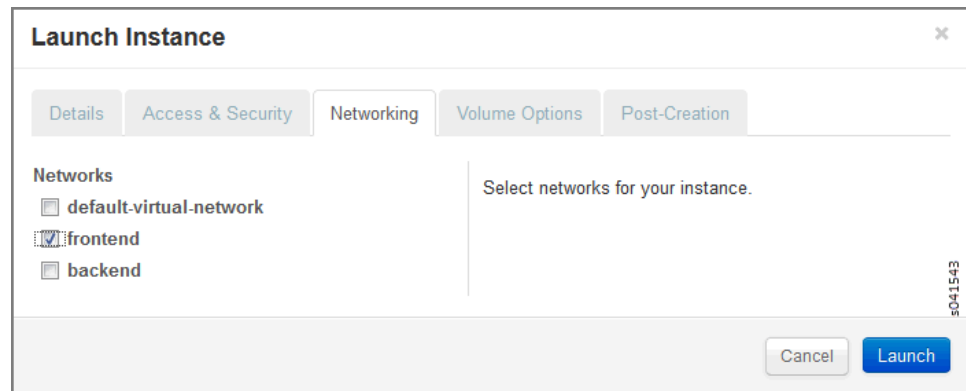
Field	Description
Instance Source	Select from the list the source type: <b>Image</b> or <b>Snapshot</b> .
Image	Select from a list the image to use for this instance. The images represent the operating systems and applications available for this project.
Instance Name	Enter a name for this instance
Flavor	From the list, select the OpenStack Flavor for this instance. Flavors provide general definitions for sizing VMs. The Flavor Details of the Flavor you select are displayed on the right column of this window.
Instance Count	Enter the number of instances you want to launch using the details defined in this window. On the right side column, <b>Project Quotas</b> displays the number of instances currently active and the number still available for this project.

Table 17: Launch Instance Details Tab Fields (*continued*)

Field	Description
<b>Compute Hostname</b>	To launch a VM on a specific compute node, enter the name of the compute node. This functionality is available only to administrators.

- Click the **Networking** tab in the **Launch Instance** window to identify one or more networks to associate with this instance; see [Figure 34 on page 138](#).

Figure 34: Launch Instance, Networking Tab



- After you finish defining this instance, click **Launch** at the lower right.  
Your new VM instance is launched as part of your project.

- Related Documentation**
- [Creating Virtual Networks and Policies in Juniper Networks Contrail on page 117](#)
  - [Creating Virtual Networks and Policies in OpenStack Contrail on page 125](#)

## Creating a Floating IP Address Pool and Allocating it to a Virtual Machine

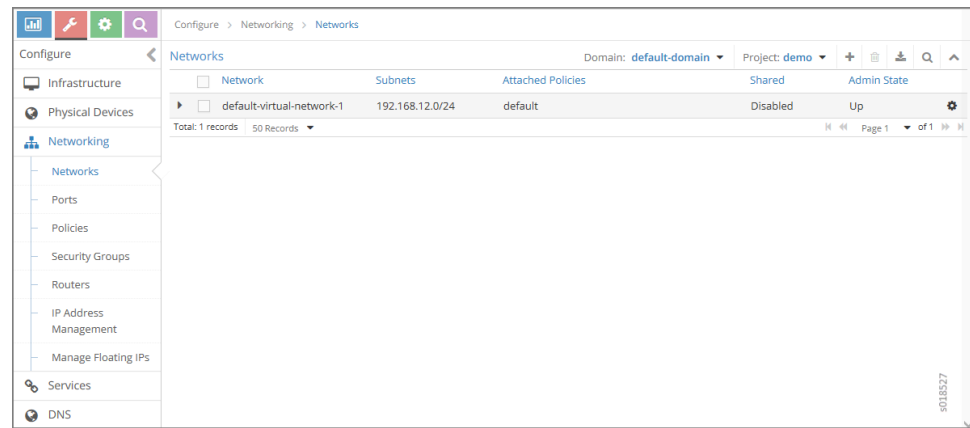
- [Creating a Floating IP Address Pool on page 139](#)
- [Allocating a Floating IP Address to a Virtual Machine on page 140](#)

## Creating a Floating IP Address Pool

A floating IP address is an IP address (typically public) that can be dynamically assigned to a running virtual instance. You can configure floating IP address pools in project networks in Contrail, then allocate floating IP addresses from the pool to virtual machine instances in other virtual networks.

1. Select **Configure > Networking > Networks**; see [Figure 35 on page 139](#). Make sure your project is the active project in the upper right.

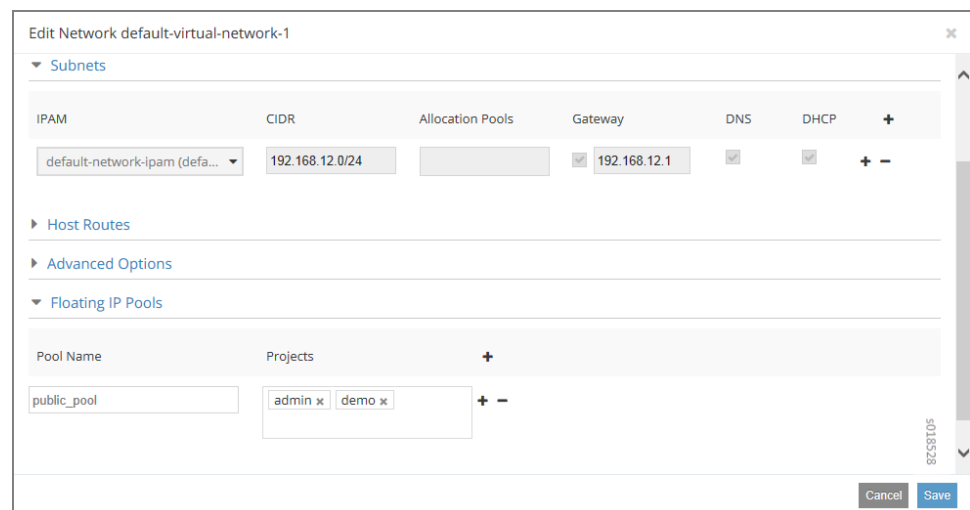
**Figure 35: Configure > Networking > Networks**



2. Click the network you want to associate with a floating IP pool, then in the **Action** column, click the action icon and select **Edit**.

The **Edit Network** window for the selected network is displayed; see [Figure 36 on page 139](#).

**Figure 36: Edit Network**



3. In the **Floating IP Pools** section, click the **Pool Name** field, enter a name for your floating IP pool, and click the + (plus sign) to add the IP pool to the table below the field.

- Multiple floating IP pools can be created at the same time.
  - A floating IP pool can be associated to multiple projects.
4. Click **Save** to create the floating IP address pool, or click **Cancel** to remove your work and start over.

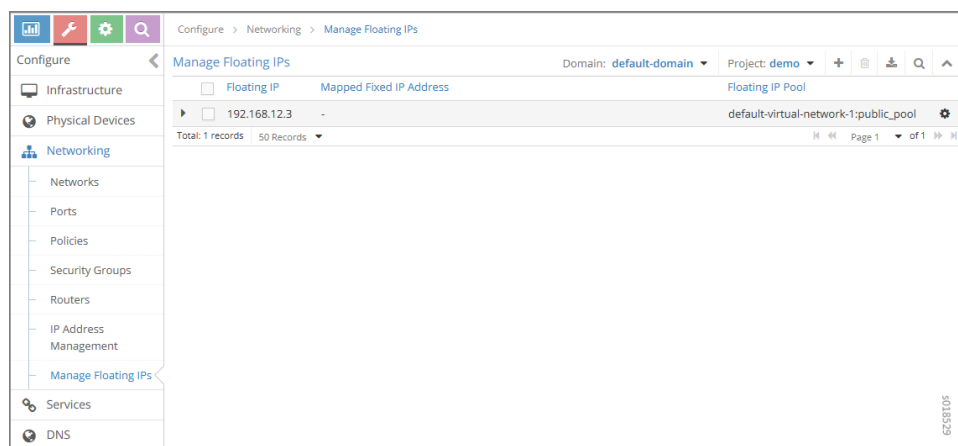
## Allocating a Floating IP Address to a Virtual Machine

To allocate a floating IP address pool to a VM instance:

1. In the Contrail Web user interface, select **Configure > Networking > Manage Floating IPs**. The the Manage Floating IPs window is displayed.

Make sure your project is displayed (active) in the upper right. See [Figure 37 on page 140](#).

**Figure 37: Manage Floating IPs**



2. In the Manage Floating IPs window, click **+**.

The **Allocate Floating IP** window is displayed; see [Figure 38 on page 141](#).

Figure 38: Allocate Floating IP Window

The figure displays two instances of the 'Allocate Floating IP' window. Both windows have a title bar with a close button (X) and a vertical ID 's018530' on the right side.

**Top Window (Dynamic Allocation):**

- Floating IP Pool:** demo:default-virtual-network-1:public\_pool (192.168.12.0/24)
- Allocation Type:** Dynamic
- Number of IP Addresses:** 1
- Buttons:** Cancel, Save

**Bottom Window (Specific IP Allocation):**

- Floating IP Pool:** demo:default-virtual-network-1:public\_pool (192.168.12.0/24)
- Allocation Type:** Specific IP
- IP Address:** 192.168.12.5
- Buttons:** Cancel, Allocate

3. Select the name of the floating IP pool from the Floating IP Pool list . The floating IP pool is shared among multiple projects.
4. Select either **Dynamic** or **Specific IP** from the Allocation Type list.
5. If Dynamic is selected, type the number of IP addresses. If Specific IP is selected, type the specific IP address.
6. If Dynamic is selected, click **Save**. If Specific IP is selected, click **Allocate**.
7. After the floating IP pool has been allocated, you can associate it to or disassociate it from instance addresses. In the Manage Floating IPs window, select the floating IP pool you want, then click the gear wheel icon and select **Associate Port** or **Disassociate**.

If you select Associate Port, the **Associate Floating IP** window is displayed; see [Figure 39 on page 142](#)

Figure 39: Associate Floating IP

8. In the **Port** field, select the UUID of the VM instance to associate with the selected floating IP pool from the Port list, and click **Save**.

#### Related Documentation

- [Creating Projects in OpenStack for Configuring Tenants in Contrail on page 116](#)
- [Creating Virtual Networks and Policies in Juniper Networks Contrail on page 117](#)
- [Creating Virtual Networks and Policies in OpenStack Contrail on page 125](#)
- [Creating an Image and Launching a Virtual Machine on page 133](#)

## Using Security Groups with Virtual Machines (Instances)

- [Security Groups Overview on page 142](#)
- [Creating Security Groups and Adding Rules on page 142](#)

### 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.

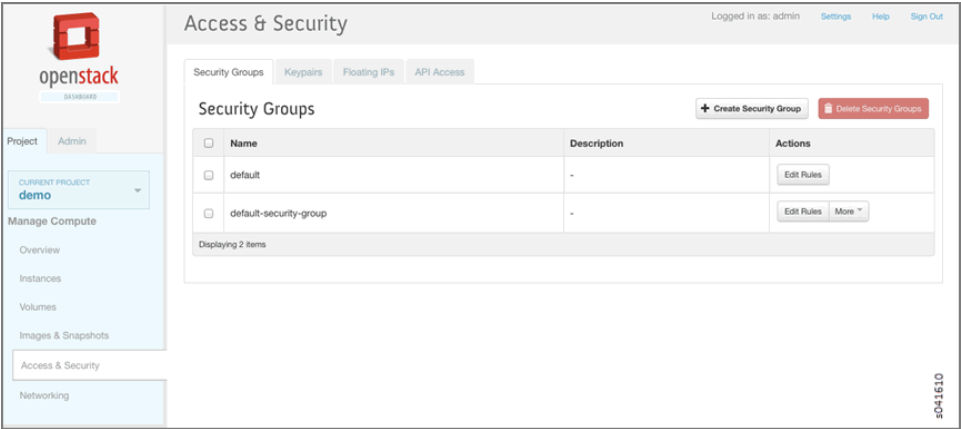


To add rules to a security group:

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 40 on page 143](#).

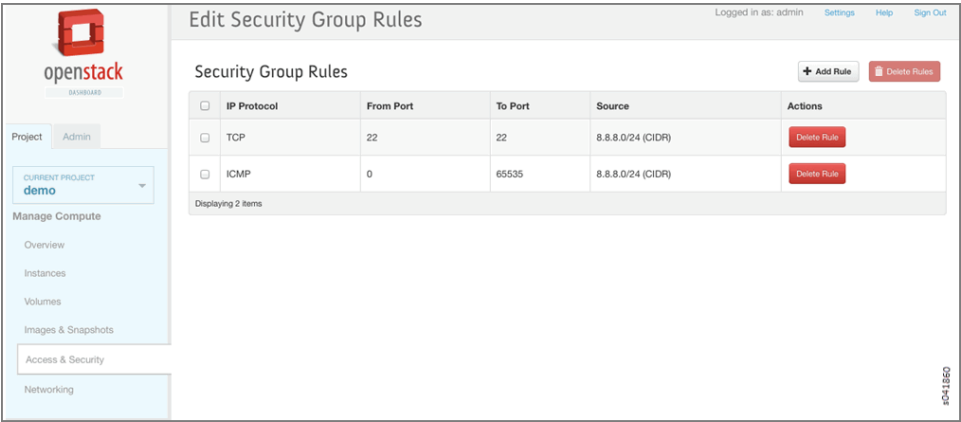
Figure 40: Security Groups



2. Select the **default-security-group** and click **Edit Rules** in the **Actions** column.

The **Edit Security Group Rules** window is displayed; see [Figure 41 on page 143](#). Any rules already associated with the security group are listed.

Figure 41: Edit Security Group Rules



3. Click **Add Rule** to add a new rule; see [Figure 42 on page 144](#).

Figure 42: Add Rule

**Add Rule**

**IP Protocol**

**Type**

**Code**

**Source**

**Description:**  
 Rules define which traffic is allowed to instances assigned to the security group. A security group rule consists of three main parts:

**Protocol:** You must specify the desired IP protocol to which this rule will apply; the options are TCP, UDP, or ICMP.

**Open Port/Port Range:** For TCP and UDP rules you may choose to open either a single port or a range of ports. Selecting the "Port Range" option will provide you with space to provide both the starting and ending ports for the range. For ICMP rules you instead specify an ICMP type and code in the spaces provided.

**Source:** You must specify the source of the traffic to be allowed via this rule. You may do so either in the form of an IP address block (CIDR) or via a source group (Security Group). Selecting a security group as the source will allow any other instance in that security group access to any other instance via this rule.

CH 18 62

Table 18: Add Rule Fields

Column	Description
<b>IP Protocol</b>	Select the IP protocol to apply for this rule: TCP, UDP, ICMP.
<b>From Port</b>	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.
<b>To Port</b>	The port to which traffic is destined that applies to this rule, using the same options as in the <b>From Port</b> field.
<b>Source</b>	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.

- Click **Create Security Group** to create additional security groups.

The **Create Security Group** window is displayed; see [Figure 43 on page 145](#).

Each new security group has a unique 32-bit security group ID and an ACL is associated with the configured rules.

Figure 43: Create Security Group

5. When an instance is launched, there is an opportunity to associate a security group; see [Figure 44 on page 145](#).

In the **Security Groups** list, select the security group name to associate with the instance.

Figure 44: 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**.

## Support for IPv6 Networks in Contrail

In Contrail Release 2.0 and later, support for IPv6 overlay networks is provided.

- [Overview: IPv6 Networks in Contrail on page 146](#)
- [Creating IPv6 Virtual Networks in Contrail on page 146](#)
- [Adding IPv6 Peers on page 147](#)

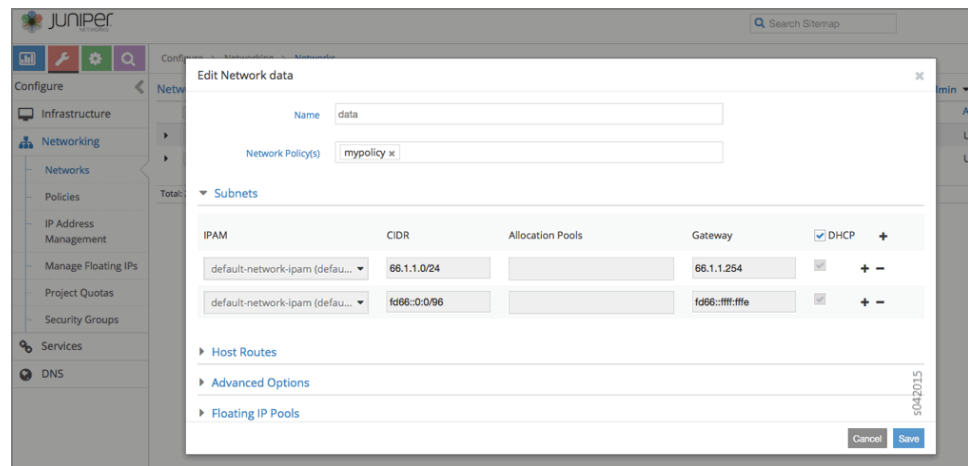
## Overview: IPv6 Networks in Contrail

In Contrail Release 2.0 and later, support for IPv6 overlay networks is provided, including:

- Configuring IPv6 subnets from the Contrail user interface or by using Neutron APIs
- IPv6 address assignment to virtual machine interfaces over DHCPv6
- IPv6 forwarding in overlay networks between virtual machines, and between virtual machines and BGP peers
- IPv6 to-VPN peering with other BGP peers
- IPv6 forwarding in Layer 2-only networks
- IPv6 interface static routes

## Creating IPv6 Virtual Networks in Contrail

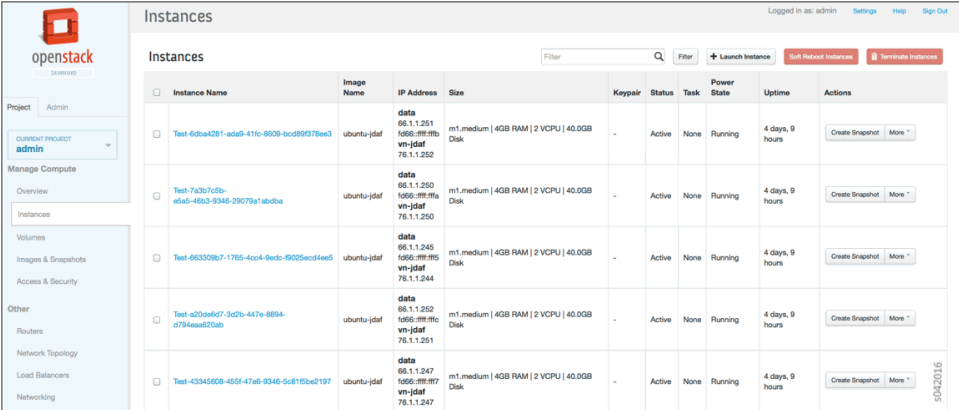
You can create an IPv6 virtual network from the Contrail user interface in the same way you create an IPv4 virtual network. When you create a new virtual network by selecting **Configure > Networking > Networks**, the Edit fields accept IPv6 addresses, as shown in the following image.



## Address Assignments

When virtual machines are launched with an IPv6 virtual network created in the Contrail user interface, the virtual machine interfaces get assigned addresses from all the families configured in the virtual network.

The following is a sample of IPv6 instances with address assignments, as listed in the OpenStack Horizon user interface.



The screenshot shows the OpenStack Horizon 'Instances' page. The table lists five instances, all of which are Ubuntu-based and have IPv6 addresses assigned. The instances are named 'Test-6d8a4261-...', 'Test-7a3b7c5b-...', 'Test-663309b7-...', 'Test-a205a6d7-...', and 'Test-43343608-...'. Each instance has a unique ID, image name, IP address, size, keypair, status, task, power state, uptime, and actions (Create Snapshot, More...).

Instance Name	Image Name	IP Address	Size	Keypair	Status	Task	Power State	Uptime	Actions
Test-6d8a4261-...	ubuntu-jdaff	66.1.1.251 66.1.1.252 75.1.1.252	m1.medium   4GB RAM   2 VCPU   40.0GB Disk	-	Active	None	Running	4 days, 9 hours	Create Snapshot More...
Test-7a3b7c5b-...	ubuntu-jdaff	66.1.1.250 66.1.1.251 75.1.1.250	m1.medium   4GB RAM   2 VCPU   40.0GB Disk	-	Active	None	Running	4 days, 9 hours	Create Snapshot More...
Test-663309b7-...	ubuntu-jdaff	66.1.1.245 66.1.1.246 75.1.1.244	m1.medium   4GB RAM   2 VCPU   40.0GB Disk	-	Active	None	Running	4 days, 9 hours	Create Snapshot More...
Test-a205a6d7-...	ubuntu-jdaff	66.1.1.250 66.1.1.251 75.1.1.251	m1.medium   4GB RAM   2 VCPU   40.0GB Disk	-	Active	None	Running	4 days, 9 hours	Create Snapshot More...
Test-43343608-...	ubuntu-jdaff	66.1.1.247 66.1.1.247 75.1.1.247	m1.medium   4GB RAM   2 VCPU   40.0GB Disk	-	Active	None	Running	4 days, 9 hours	Create Snapshot More...

### Enabling DHCPv6 In Virtual Machines

To allow IPv6 address assignment using DHCPv6, the virtual machine network interface configuration must be updated appropriately.

For example, to enable DHCPv6 for Ubuntu-based virtual machines, add the following line in the `/etc/network/interfaces` file:

```
iface eht0 inet6 dhcp
```

Also, `dhclient -6` can be run from within the virtual machine to get IPv6 addresses using DHCPv6.

## Adding IPv6 Peers

The procedure to add an IPv6 BGP peer in Contrail is similar to adding an IPv4 peer. Select **Configure > Infrastructure > BGP Peers**, include `inet6-vpn` in the Address Family list to allow advertisement of IPv6 addresses.

A sample is shown in the following.



**NOTE:** Additional configuration is required on the peer router to allow inet6-vpn peering.

## Configuring EVPN and VXLAN

Contrail supports Ethernet VPNs (EVPN) and Virtual Extensible Local Area Networks (VXLAN).

EVPN is a flexible solution that uses Layer 2 overlays to interconnect multiple edges (virtual machines) within a data center. Traditionally, the data center is built as a flat Layer 2 network with issues such as flooding, limitations in redundancy and provisioning, and high volumes of MAC address learning, which cause churn during node failures. EVPNs are designed to address these issues without disturbing flat MAC connectivity.

In EVPNs, MAC address learning is driven by the control plane, rather than by the data plane, which helps control learned MAC addresses across virtual forwarders, thus avoiding flooding. The forwarders advertise locally learned MAC addresses to the controllers. The controllers use MP-BGP to communicate with peers. The peering of controllers using BGP for EVPN results in better and faster convergence.

With EVPN, MAC learning is confined to the virtual networks to which the virtual machine belongs, thus isolating traffic between multiple virtual networks. In this manner, virtual networks can share the same MAC addresses without any traffic crossover.

### *Unicast in EVPNs*

Unicast forwarding is based on MAC addresses where traffic can terminate on a local endpoint or is encapsulated to reach the remote endpoint. Encapsulation can be MPLS/UDP, MPLS/GRE, or VXLAN.

### *BUM Traffic in EVPN*

Multicast and broadcast traffic is flooded in a virtual network. The replication tree is built by the control plane, based on the advertisements of end nodes (virtual machines) sent by forwarders. Each virtual network has one distribution tree, a method that avoids maintaining multicast states at fabric nodes, so the nodes are unaffected by multicast. The replication happens at the edge forwarders. Per-group subscription is not provided. Broadcast, unknown unicast, and multicast (BUM) traffic is handled the same way, and gets flooded in the virtual network to which the virtual machine belongs.

### *VXLAN*

VXLAN is an overlay technology that encapsulates MAC frames into a UDP header at Layer 2. Communication is established between two virtual tunnel endpoints (VTEPs). VTEPs encapsulate the virtual machine traffic into a VXLAN header, as well as strip off the encapsulation. Virtual machines can only communicate with each other when they belong to the same VXLAN segment. A 24-bit virtual network identifier (VNID) uniquely identifies the VXLAN segment. This enables having the same MAC frames across multiple VXLAN segments without traffic crossover. Multicast in VXLAN is implemented as Layer 3 multicast, in which endpoints subscribe to groups.

### *Design Details of EVPN and VXLAN*

In Contrail Release 1.03 and later, EVPN is enabled by default. The supported forwarding modes include:

- Fallback bridging—IPv4 traffic lookup is performed using the IP FIB. All non-IPv4 traffic is directed to a MAC FIB.
- Layer 2-only— All traffic is forwarded using a MAC FIB lookup.

You can configure the forwarding mode individually on each virtual network.

EVPN is used to share MAC addresses across different control planes in both forwarding models. The result of a MAC address lookup is a next hop, which, similar to IP forwarding, points to a local virtual machine or a tunnel to reach the virtual machine on a remote server. The tunnel encapsulation methods supported for EVPN are MPLSoGRE, MPLSoUDP, and VXLAN. The encapsulation method selected is based on a user-configured priority.

In VXLAN, the VNID is assigned uniquely for every virtual network carried in the VXLAN header. The VNID uniquely identifies a virtual network. When the VXLAN header is received from the fabric at a remote server, the VNID lookup provides the VRF of the virtual machine. This VRF is used for the MAC lookup from the inner header, which then provides the destination virtual machine.

Non-IP multicast traffic uses the same multicast tree as for IP multicast (255.255.255.255). The multicast is matched against the all-broadcast prefix in the bridging table (FF:FF:FF:FF:FF:FF). VXLAN is not supported for IP/non-IP multicast traffic.

The following table summarizes the traffic and encapsulation types supported for EVPN.

		Encapsulation		
		MPLS-GRE	MPLS-UDP	VXLAN
Traffic Type	IP unicast	Yes	Yes	No
	IP-BUM	Yes	Yes	No
	non IP unicast	Yes	Yes	Yes
	non IP-BUM	Yes	Yes	No

- [Configuring the VXLAN Identifier Mode on page 150](#)
- [Configuring Forwarding on page 151](#)
- [Configuring the VXLAN Identifier on page 153](#)
- [Configuring Encapsulation Methods on page 154](#)

## Configuring the VXLAN Identifier Mode

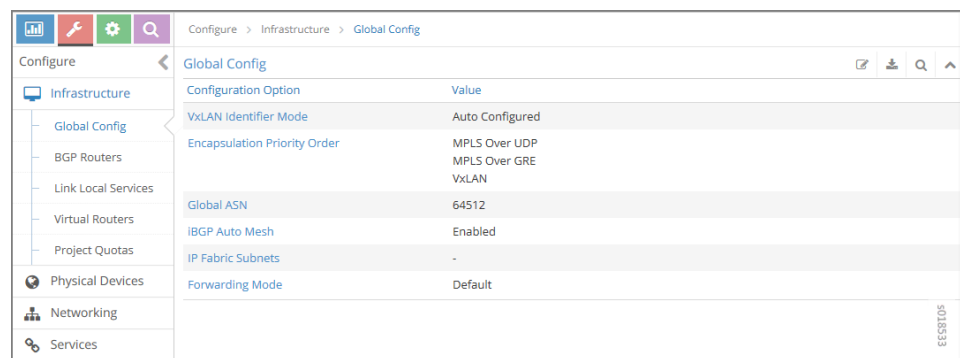
You can configure the global VXLAN identifier mode to select an auto-generated VNID or a user-generated VXLAN ID, either through the Contrail Web UI or by modifying a python file.

To configure the global VXLAN identifier mode:

1. From the Contrail Web UI, select **Configure > Infrastructure > Global Config**.

The Global Config options and values are displayed in the Global Config window.

**Figure 45: Global Config Window for VXLAN ID**



2. Click the edit icon .

The Edit Global Config window is displayed as shown in [Figure 46 on page 151](#).



Figure 46: Edit Global Config Window for VXLAN Identifier Mode

3. Select one of the following:

- **Auto Configured**— The VXLAN identifier is automatically assigned for the virtual network.
- **User Configured**— You must provide the VXLAN identifier for the virtual network.



**NOTE:** When **User Configured** is selected, if you do not provide an identifier, then VXLAN encapsulation *is not used* and the mode falls back to MPLS.

Alternatively, you can set the VXLAN identifier mode by using Python to modify the `/opt/contrail/utils/encap.py` file as follows:

```
python encap.py <add | update | delete> <username> <password> <tenant_name> <config_node_ip>
```

## Configuring Forwarding

In Contrail, the default forwarding mode is enabled for fallback bridging (IP FIB and MAC FIB). The mode can be changed, either through the Contrail Web UI or by using python provisioning commands.

To change the forwarding mode:

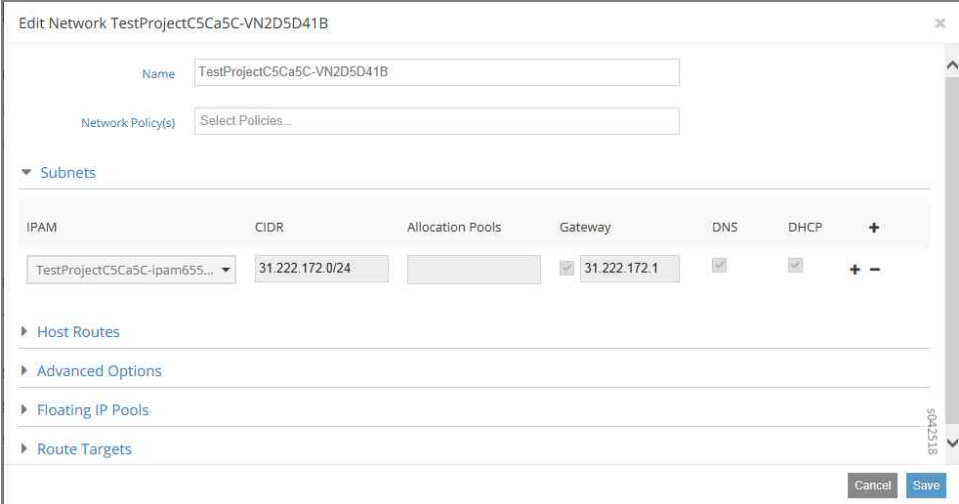
1. From the Contrail Web UI, select **Configure > Networking > Networks**.

2. Select the virtual network that you want to change the forwarding mode for.

3. Click the gear icon  and select **Edit**.

The Edit Network window is displayed as shown in [Figure 47 on page 152](#).

**Figure 47: Edit Network Window**



Under the Advanced Options select the forwarding mode from the following choices:

- Select **Default** to enable the default forwarding mode.
- Select **L2 and L3** to enable IP and MAC FIB (fallback bridging).
- Select **L2 Only** to enable only MAC FIB.
- Select **L3 Only** to enable only IP.



**NOTE:** The full list of forwarding modes are only displayed if you change entries in the `/usr/src/contrail/contrail-web-core/config/config.global.js` file. For example:

1. To make the **L2** selection available locate the following:

```
config.network = {};
config.network.L2_enable = false;
```

2. Change the entry to the following:

```
config.network = {};
config.network.L2_enable = true;
```

3. To make the other selections available, modify the corresponding entries.

4. Save the file and quit the editor.

5. Restart the Contrail Web user interface process (webui).

Alternatively, you can use the following python provisioning command to change the forwarding mode:

```
python provisioning forwarding_mode --project_fq_name 'defaultdomain: admin' --vn_name vn1 --forwarding_mode < l2_l3| l2 >
```

Options:


**l2\_l3** = Enable IP FIB and MAC FIB (fallback bridging)

**l2** = Enable MAC FIB only (Layer 2 only)

## Configuring the VXLAN Identifier

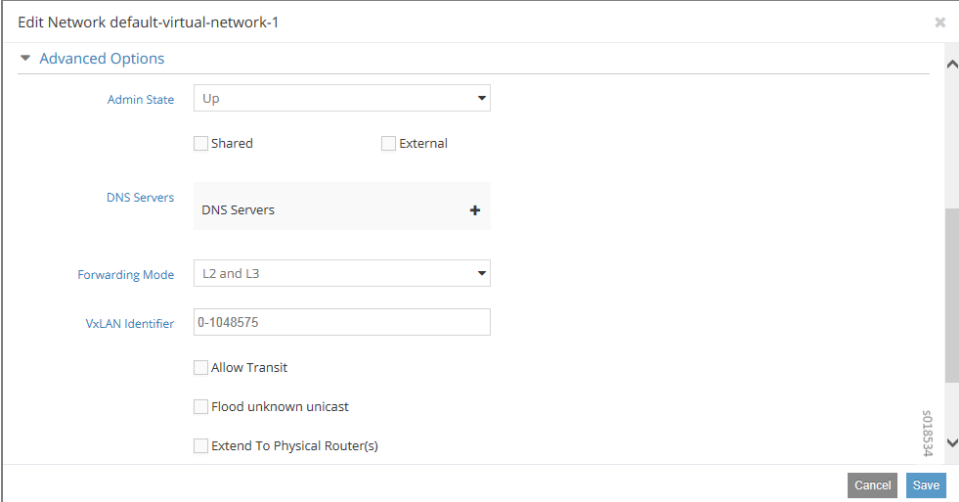
The VXLAN identifier can be set only if the VXLAN network identifier mode has been set to User Configured. You can then set the VXLAN ID by either using the Contrail Web UI or by using Python commands.

To configure the global VXLAN identifier:

1. From the Contrail Web UI, select **Configure > Networking > Networks**.
2. Select the virtual network that you want to change the forwarding mode for.
3. Click the gear icon  and select **Edit**.

The Edit Network window is displayed. Select the **Advanced Options** as shown in [Figure 48 on page 153](#).

**Figure 48: Edit Network Window for VXLAN Identifier**



4. Type the VXLAN identifier.
5. Click **Save**.

Alternatively, you can use the following Python provisioning command to configure the VXLAN identifier:

```
python provisioning_forwarding_mode --project_fq_name 'defaultdomain:admin' --vn_name
vn1 --forwarding_mode < vxlan_id >
```


## Configuring Encapsulation Methods

The default encapsulation mode for EVPN is MPLS over UDP. All packets on the fabric are encapsulated with the label allocated for the virtual machine interface. The label encoding and decoding is the same as for IP forwarding. Additional encapsulation methods supported for EVPN include MPLS over GRE and VXLAN. MPLS over UDP is different from MPLS over GRE only in the method of tunnel header encapsulation.

VXLAN has its own header and uses a VNID label to carry the traffic over the fabric. A VNID is assigned with every virtual network and is shared by all virtual machines in the virtual network. The VNID is mapped to the VRF of the virtual network to which it belongs.

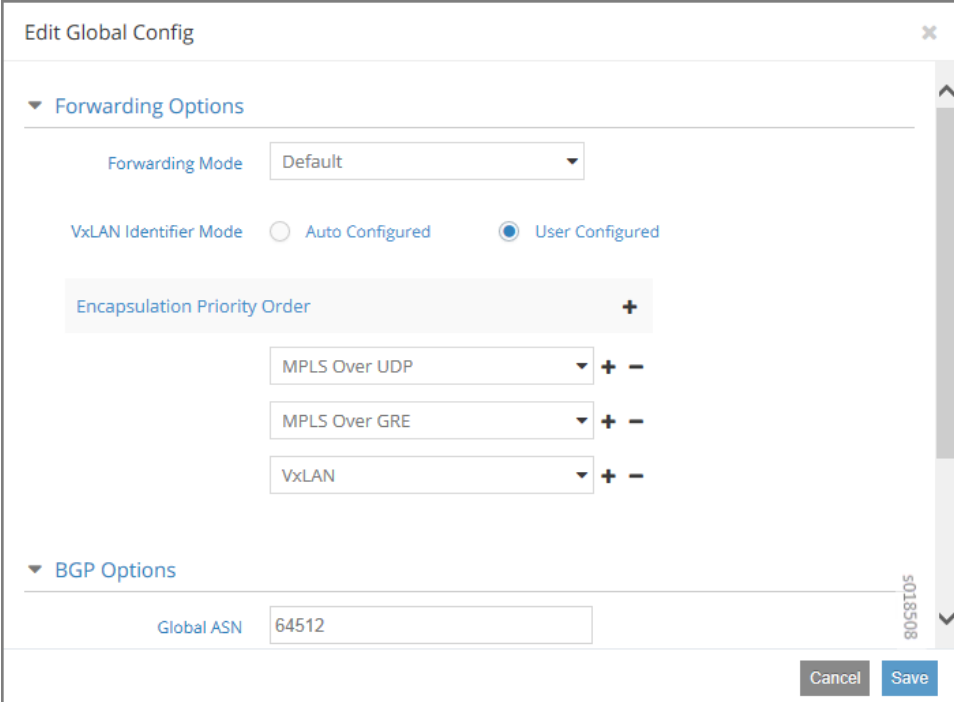
The priority order in which to apply encapsulation methods is determined by the sequence of methods set either from the Contrail Web UI or in the **encap.py** file.

To configure the global VXLAN identifier mode:

- From the Contrail Web UI, select **Configure > Infrastructure > Global Config**.
- The Global Config options are displayed.
- Click the edit icon  .

The Edit Global Config window is displayed as shown in [Figure 49 on page 154](#).

**Figure 49: Edit Global Config Window for Encapsulation Priority Order**



**Edit Global Config**

▼ **Forwarding Options**

Forwarding Mode: Default

VxLAN Identifier Mode: ☐ Auto Configured ☒ User Configured

Encapsulation Priority Order

- MPLS Over UDP
- MPLS Over GRE
- VxLAN

▼ **BGP Options**

Global ASN: 64512

Cancel Save

Under Encapsulation Priority Order select one of the following:

- MPLS over UDP
- MPLS over GRE
- VxLAN

Click the + plus symbol to the right of the first priority to add a second priority or third priority.

Use the following procedure to change the default encapsulation method to VXLAN by editing the `encap.py` file.



**NOTE:** VXLAN is *only* supported for EVPN unicast. It is not supported for IP traffic or multicast traffic. VXLAN priority and presence in the `encap.py` file or configured in the Web UI is ignored for traffic not supported by VXLAN.

To set the priority of encapsulation methods to VXLAN:

1. Modify the `encap.py` file found in the `/opt/contrail/utils/` directory.

The default encapsulation line is:

```
encap_obj=EncapsulationPrioritiesType(encapsulation=['MPLSoUDP','MPLSoGRE'])
```

Modify the line to:

```
encap_obj=EncapsulationPrioritiesType(encapsulation=['VXLAN',
'MPLSoUDP','MPLSoGRE'])
```

2. After the status is modified, execute the following script:

```
python encap_set.py <add|update|delete> <username> <password> <tenant_name>
<config_node_ip>
```

The configuration is applied globally for all virtual networks.



## CHAPTER 8

# Example of Deploying a Multi-Tier Web Application Using Contrail

- [Example: Deploying a Multi-Tier Web Application on page 157](#)
- [Sample Network Configuration for Devices for Simple Tiered Web Application on page 163](#)

### Example: Deploying a Multi-Tier Web Application

---


- [Multi-Tier Web Application Overview on page 157](#)
- [Example: Setting Up Virtual Networks for a Simple Tiered Web Application on page 158](#)
- [Verifying the Multi-Tier Web Application on page 160](#)
- [Sample Addressing Scheme for Simple Tiered Web Application on page 161](#)
- [Sample Physical Topology for Simple Tiered Web Application on page 162](#)
- [Sample Physical Topology Addressing on page 162](#)

### Multi-Tier Web Application Overview

A common requirement for a cloud tenant is to create a tiered web application in leased cloud space. The tenant enjoys the favorable economics of a private IT infrastructure within a shared services environment. The tenant seeks speedy setup and simplified operations.

The following example shows how to set up a simple tiered web application using Contrail. The example has a web server that a user accesses by means of a public floating IP address. The front-end web server gets the content it serves to customers from information stored in a SQL database server that resides on a back-end network. The web server can communicate directly with the database server without going through any gateways. The public (or client) can only communicate to the web server on the front-end network. The client is not allowed to communicate directly with any other parts of the infrastructure. See [Figure 50 on page 158](#).

Figure 50: Simple Tiered Web Use Case



### Example: Setting Up Virtual Networks for a Simple Tiered Web Application

This example provides basic steps for setting up a simple multi-tier network application. Basic creation steps are provided, along with links to the full explanation for each of the creation steps. Refer to the links any time you need more information about completing a step.

1. Working with a system that has the Contrail software installed and provisioned, create a project named **demo**.

For more information; see [“Creating Projects in OpenStack for Configuring Tenants in Contrail” on page 116](#).

2. In the **demo** project, create three virtual networks:

- a. A network named **public** with IP address **10.84.41.0/24**

This is a special use virtual network for floating IP addresses— it is assigned an address block from the public floating address pool that is assigned to each web server. The assigned block is the only address block advertised outside of the data center to clients that want to reach the web services provided.

- b. A network named **frontend** with IP address **192.168.1.0/24**

This network is the location where the web server virtual machine instances are launched and attached. The virtual machines are identified with private addresses that have been assigned to this virtual network.

- c. A network named **backend** with IP address **192.168.2.0/24**

This network is the location where the database server virtual machines instances are launched and attached. The virtual machines are identified with private addresses that have been assigned to this virtual network.



For more information; see “Creating Virtual Networks and Policies in OpenStack Contrail” on page 125 or “Creating Virtual Networks and Policies in Juniper Networks Contrail” on page 117.

3. Create a floating IP pool named **public\_pool** for the **public** network within the **demo** project; see Figure 51 on page 159.

Figure 51: Create Floating IP Pool

Edit Network public

Network Name: public

Network Policy(s): Select Policies...

Address Management: default-network...  + -

IPAM	IP Block
default-network-ipam	10.84.41.0/24

Floating IP Pools: public\_pool  + -

Pool Name: admin

Cancel Save

4. Allocate the floating IP pool **public\_pool** to the **demo** project; see Figure 52 on page 159.

Figure 52: Allocate Floating IP

Allocate Floating IP

Floating IP Pool: public:public\_pool

Cancel Save

5. Verify that the floating IP pool has been allocated; see **Configure > Networking > Allocate Floating IPs**.

For more information; see [“Creating a Floating IP Address Pool and Allocating it to a Virtual Machine” on page 138](#) and [“Allocating a Floating IP Address to a Virtual Machine” on page 140](#).

6. Create a policy that allows any host to talk to any host using any IP address, protocol, and port, and apply this policy between the **frontend** network and the **backend** network.

This now allows communication between the web servers in the front-end network and the database servers in the back-end network.

For more information; see [“Creating a Network Policy with Juniper Networks Contrail” on page 121](#), [“Associating a Network to a Policy—Juniper Networks Contrail” on page 123](#), or [“Creating a Network Policy with OpenStack Contrail” on page 129](#), and [“Associating a Network to a Policy—OpenStack Contrail” on page 131](#).

7. Launch the virtual machine instances that represent the web server and the database server.



**NOTE:** Your installation might not include the virtual machines needed for the web server and the database server. Contact your account team if you need to download the VMs for this setup.

On the **Instances** tab for this project, select **Launch Instance** and for each instance that you launch, complete the fields to make the following associations:

- Web server VM: select **frontend** network and the policy created to allow communication between **frontend** and **backend** networks. Apply the floating IP address pool to the web server.
- Database server VM: select **backend** network and the policy created to allow communication between **frontend** and **backend** networks.

For more information; see [“Launching a Virtual Machine \(Instance\)” on page 136](#).

## Verifying the Multi-Tier Web Application

Verify your web setup.

- To demonstrate this web application setup, go to the client machine, open a browser, and navigate to the address in the **public** network that is assigned to the web server in the **frontend** network.

The result will display the Contrail interface with various data populated, verifying that the web server is communicating with the database server in the **backend** network and retrieving data.

The client machine only has access to the public IP address. Attempts to browse to any of the addresses assigned to the **frontend** network or to the **backend** network should fail.

### Sample Addressing Scheme for Simple Tiered Web Application

Use the information in [Table 19 on page 161](#) as a guide for addressing devices in the simple tiered web example.


**Table 19: Sample Addressing Scheme for Example**

System Name	Address Allocation
System001	10.84.11.100
System002	10.84.11.101
System003	10.84.11.102
System004	10.84.11.103
System005	10.84.11.104
MX80-1	10.84.11.253 10.84.45.1 (public connection)
MX80-2	10.84.11.252 10.84.45.2 (public connection)
EX4200	10.84.11.254 10.84.45.254 (public connection) 10.84.63.259 (public connection)
frontend network	192.168.1.0/24
backend network	192.168.2.0/24
public network (floating address)	10.84.41.0/24

## Sample Physical Topology for Simple Tiered Web Application

Figure 53 on page 162 provides a guideline diagram for the physical topology for the simple tiered web application example.


Figure 53: Sample Physical Topology for Simple Tiered Web Application



## Sample Physical Topology Addressing

Figure 54 on page 163 provides a guideline diagram for addressing the physical topology for the simple tiered web application example.

Figure 54: Sample Physical Topology Addressing



## Sample Network Configuration for Devices for Simple Tiered Web Application

This section shows sample device configurations that can be used to create the “[Example: Deploying a Multi-Tier Web Application](#)” on page 157. Configurations are shown for Juniper Networks devices: two MX80s and one EX4200.

### MX80-1 Configuration

```
version 12.2R1.3;
system {
  root-authentication {
    encrypted-password "xxxxxxxxx"; ## SECRET-DATA
  }
  services {
    ssh {
      root-login allow;
    }
  }
  syslog {
    user * {
```

```
        any emergency;
    }
    file messages {
        any notice;
        authorization info;
    }
}
chassis {
    fpc 1 {
        pic 0 {
            tunnel-services;
        }
    }
}
interfaces {
    ge-1/0/0 {
        unit 0 {
            family inet {
                address 10.84.11.253/24;
            }
        }
    }
    ge-1/1/0 {
        description "IP Fabric interface";
        unit 0 {
            family inet {
                address 10.84.45.1/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 127.0.0.1/32;
            }
        }
    }
}
routing-options {
    static {
        route 0.0.0.0/0 next-hop 10.84.45.254;
    }
    route-distinguisher-id 10.84.11.253;
    autonomous-system 64512;
    dynamic-tunnels {
        setup1 {
            source-address 10.84.11.253;
            gre;
            destination-networks {
                10.84.11.0/24;
            }
        }
    }
}
protocols {
```

```

bgp {
  group mx {
    type internal;
    local-address 10.84.11.253;
    family inet-vpn {
      unicast;
    }
    neighbor 10.84.11.252;
  }
  group contrail-controller {
    type internal;
    local-address 10.84.11.253;
    family inet-vpn {
      unicast;
    }
    neighbor 10.84.11.101;
    neighbor 10.84.11.102;
  }
}
routing-instances {
  customer-public {
    instance-type vrf;
    interface ge-1/1/0.0;
    vrf-target target:64512:10000;
    routing-options {
      static {
        route 0.0.0.0/0 next-hop 10.84.45.254;
      }
    }
  }
}
}

```

#### *MX80-2 Configuration*

```

version 12.2R1.3;
system {
  root-authentication {
    encrypted-password "xxxxxxxxx"; ## SECRET-DATA
  }
  services {
    ssh {
      root-login allow;
    }
  }
  syslog {
    user * {
      any emergency;
    }
    file messages {
      any notice;
      authorization info;
    }
  }
}
chassis {

```

```
fpc 1 {
  pic 0 {
    tunnel-services;
  }
}
interfaces {
  ge-1/0/0 {
    unit 0 {
      family inet {
        address 10.84.11.252/24;
      }
    }
  }
  ge-1/1/0 {
    description "IP Fabric interface";
    unit 0 {
      family inet {
        address 10.84.45.2/24;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 127.0.0.1/32;
      }
    }
  }
}
routing-options {
  static {
    route 0.0.0.0/0 next-hop 10.84.45.254;
  }
  route-distinguisher-id 10.84.11.252;
  autonomous-system 64512;
  dynamic-tunnels {
    setup1 {
      source-address 10.84.11.252;
      gre;
      destination-networks {
        10.84.11.0/24;
      }
    }
  }
}
protocols {
  bgp {
    group mx {
      type internal;
      local-address 10.84.11.252;
      family inet-vpn {
        unicast;
      }
      neighbor 10.84.11.253;
    }
  }
}
```



```

group contrail-controller {
    type internal;
    local-address 10.84.11.252;
    family inet-vpn {
        unicast;
    }
    neighbor 10.84.11.101;
    neighbor 10.84.11.102;
}
}
}
routing-instances {
    customer-public {
        instance-type vrf;
        interface ge-1/1/0.0;
        vrf-target target:64512:10000;
        routing-options {
            static {
                route 0.0.0.0/0 next-hop 10.84.45.254;
            }
        }
    }
}
}

```

#### *EX4200 Configuration*

```

system {
    host-name EX4200;
    time-zone America/Los_Angeles;
    root-authentication {
        encrypted-password "xxxxxxxxxxxx"; ## SECRET-DATA
    }
    login {
        class read {
            permissions [ clear interface view view-configuration ];
        }
        user admin {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "xxxxxxxxxxxx"; ## SECRET-DATA
            }
        }
        user user {
            uid 2002;
            class read;
            authentication {
                encrypted-password "xxxxxxxxxxxx"; ## SECRET-DATA
            }
        }
    }
}
services {
    ssh {
        root-login allow;
    }
    telnet;
}

```

```
netconf {
    ssh;
}
web-management {
    http;
}
}
syslog {
    user * {
        any emergency;
    }
    file messages {
        any notice;
        authorization info;
    }
    file interactive-commands {
        interactive-commands any;
    }
}
}
chassis {
    aggregated-devices {
        ethernet {
            device-count 64;
        }
    }
}
}
```

## CHAPTER 9

# Configuring Services

- [Configuring DNS Servers on page 169](#)
- [Configuring Discovery Service on page 178](#)
- [Support for Multicast on page 182](#)
- [Using Static Routes with Services on page 184](#)
- [Configuring Metadata Service on page 188](#)

## Configuring DNS Servers

---

- [DNS Overview on page 169](#)
- [Defining Multiple Virtual Domain Name Servers on page 170](#)
- [IPAM and Virtual DNS on page 170](#)
- [DNS Record Types on page 171](#)
- [Configuring DNS Using the Interface on page 172](#)
- [Configuring DNS Using Scripts on page 177](#)

## DNS Overview

Domain Name System (DNS) is the standard protocol for resolving domain names into IP addresses so that traffic can be routed to its destination. DNS provides the translation between human-readable domain names and their IP addresses. The domain names are defined in a hierarchical tree, with a root followed by top-level and next-level domain labels.

A DNS server stores the records for a domain name and responds to queries from clients based on these records. The server is authoritative for the domains for which it is configured to be the name server. For other domains, the server can act as a caching server, fetching the records by querying other domain name servers.

The following are the key attributes of domain name service in a virtual world:


- It should be possible to configure multiple domain name servers to provide name resolution service for the virtual machines spawned in the system.
- It should be possible to configure the domain name servers to form DNS server hierarchies required by each tenant.

- The hierarchies can be independent and completely isolated from other similar hierarchies present in the system, or they can provide naming service to other hierarchies present in the system.
- DNS records for the virtual machines spawned in the system should be updated dynamically when a virtual machine is created or destroyed.
- The service should be scalable to handle an increase in servers and the resulting increased numbers of virtual machines and DNS queries handled in the system.

## Defining Multiple Virtual Domain Name Servers

Contrail provides the flexibility to define multiple virtual domain name servers under each domain in the system. Each virtual domain name server is an authoritative server for the DNS domain configured. [Figure 55 on page 170](#) shows examples of virtual DNS servers defined in **default-domain**, providing the name service for the DNS domains indicated.


Figure 55: DNS Servers Examples



## IPAM and Virtual DNS

Each IP address management (IPAM) service in the system can refer to one of the virtual DNS servers configured. The virtual networks and virtual machines spawned are associated with the DNS domain specified in the corresponding IPAM. When the VMs are configured with DHCP, they receive the domain assignment in the DHCP **domain-name** option. Examples are shown in [Figure 56 on page 171](#)

Figure 56: IPAM and Virtual DNS



## DNS Record Types


DNS records can be added statically. DNS record types **A**, **CNAME**, **PTR**, and **NS** are currently supported in the system. Each record includes the type, class (IN), name, data, and TTL values. See [Table 20 on page 171](#) for descriptions of the record types.

Table 20: DNS Record Types Supported

DNS Record Type	Description
<b>A</b>	Used for mapping hostnames to IPv4 addresses. Name refers to the name of the virtual machine, and data is the IPv4 address of the virtual machine.
<b>CNAME</b>	Provides an alias to a name. Name refers to the name of the virtual machine, and data is the new name (alias) for the virtual machine.
<b>PTR</b>	A pointer to a record, it provides reverse mapping from an IP address to a name. Name refers to the IP address, and data is the name for the virtual machine. The address in the PTR record should be part of a subnet configured for a VN within one of the IPAMs referring to this virtual DNS server.
<b>NS</b>	Used to delegate a subdomain to another DNS server. The DNS server could be another virtual DNS server defined in the system or the IP address of an external DNS server reachable via the infrastructure. Name refers to the subdomain being delegated, and data is the name of the virtual DNS server or IP address of an external server.

[Figure 57 on page 172](#) shows an example usage for the DNS record type of **NS**.

Figure 57: Example Usage for NS Record Type



## Configuring DNS Using the Interface

DNS can be configured by using the user interface or by using scripts. The following procedure shows how to configure DNS through the Juniper Networks Contrail interface.

1. Access **Configure > DNS > Servers** to create or delete virtual DNS servers and records.

The **Configure DNS Records** page appears; see [Figure 58 on page 172](#).

Figure 58: Configure DNS Records

Configure > DNS > Servers

Search

Configure DNS Records

default-domain admin

Create Delete

Virtual DNS Name	DNS Domain Name	Next DNS Server
No Data Found		

DNS Records Associated IPAMs

DNS Records of {{dnsname}}

Add Record Delete

Name	Type : Data	TTL (secs)	Class
------	-------------	------------	-------

2. To add a new DNS server, click the **Create** button.

Enter DNS server information in the **Add DNS** window; see [Figure 59 on page 173](#)

**Figure 59: Add DNS**

Complete the fields for the new server; see [Table 21 on page 173](#).

**Table 21: Add DNS Fields**

Field	Description
Server Name	Enter a name for this server.
Domain Name	Enter the name of the domain for this server.
Time To Live	Enter the TTL in seconds.
Next DNS Server	Select from a list the name of the next DNS server to process DNS requests if they cannot be processed at this server, or <b>None</b> .
Load Balancing Order	Select the load-balancing order from a list— <b>Random</b> , <b>Fixed</b> , <b>Round Robin</b> . When a name has multiple records matching, the configured record order determines the order in which the records are sent in the response. Select <b>Random</b> to have the records sent in random order. Select <b>Fixed</b> to have records sent in the order of creation. Select <b>Round Robin</b> to have the record order cycled for each request to the record.
OK	Click <b>OK</b> to create the record.
Cancel	Click <b>Cancel</b> to clear the fields and start over.

- To add a new DNS record, from the **Configure DNS Records** page, click the **Add Record** button in the lower right portion of the screen.

The **Add DNS Record** window appears; see [Figure 60 on page 174](#).

Figure 60: Add DNS Record

4. Complete the fields for the new record; see [Table 22 on page 174](#).

Table 22: Add DNS Record Fields

Field	Description
Record Name	Enter a name for this record.
Type	Select the record type from a list— <b>A</b> , <b>CNAME</b> , <b>PTR</b> , <b>NS</b> .
IP Address	Enter the IP address for the location for this record.
Class	Select the record class from a list— <b>IN</b> is the default.
Time To Live	Enter the <b>TTL</b> in seconds.
OK	Click <b>OK</b> to create the record.
Cancel	Click <b>Cancel</b> to clear the fields and start over.

5. To associate an IPAM to a virtual DNS server, from the **Configure DNS Records** page, select the **Associated IPAMs** tab in the lower right portion of the screen and click the **Edit** button.

The **Associate IPAMs to DNS** window appears; see [Figure 61 on page 175](#).



Figure 61: Associate IPAMs to DNS

Complete the IPAM associations, using the field descriptions in [Table 23 on page 175](#).

Table 23: Associate IPAMs to DNS Fields

Field	Description
<b>Associate to All IPAMs</b>	Select this box to associate the selected DNS server to all available IPAMs.
<b>Available IPAMs</b>	This column displays the currently available IPAMs.
<b>Associated IPAMs</b>	This column displays the IPAMs currently associated with the selected DNS server.
>>	Use this button to associate an available IPAM to the selected DNS server, by selecting an available IPAM in the left column and clicking this button to move it to the Associated IPAMs column. The selected IPAM is now associated with the selected DNS server.
<<	Use this button to disassociate an IPAM from the selected DNS server, by selecting an associated IPAM in the right column and clicking this button to move it to the left column (Available IPAMs). The selected IPAM is now disassociated from the selected DNS server.
<b>OK</b>	Click <b>OK</b> to commit the changes indicated in the window.
<b>Cancel</b>	Click <b>Cancel</b> to clear all entries and start over.

- Use the **IP Address Management** page (**Configure > Networking > IP Address Management**); see [Figure 62 on page 176](#)) to configure the DNS mode for any DNS server and to associate an IPAM to DNS servers of any mode or to tenants' IP addresses.

**Figure 62: Configure IP Address Management**

7. To associate an IPAM to a virtual DNS server or to tenant's IP addresses, at the **IP Address Management** page, select the network associated with this IPAM, then click the **Action** button in the last column, and click **Edit**.

The **Edit IP Address Management** window appears; see [Figure 63 on page 176](#).

**Figure 63: DNS Server**

8. In the first field, select the **DNS Method** from a list (**None**, **Default DNS**, **Tenant DNS**, **Virtual DNS**; see [Table 24 on page 176](#).

**Table 24: DNS Modes**

DNS Mode	Description
<b>None</b>	Select <b>None</b> when no DNS support is required for the VMs.
<b>Default</b>	In default mode, DNS resolution for VMs is performed based on the name server configuration in the server infrastructure. The subnet default gateway is configured as the DNS server for the VM, and the DHCP response to the VM has this DNS server option. DNS requests sent by a VM to the default gateway are sent to the name servers configured on the respective compute nodes. The responses are sent back to the VM.

Table 24: DNS Modes (*continued*)

DNS Mode	Description
<b>Tenant</b>	Configure this mode when a tenant wants to use its own DNS servers. Configure the list of servers in the IPAM. The server list is sent in the DHCP response to the VM as DNS servers. DNS requests sent by the VMs are routed the same as any other data packet based on the available routing information.
<b>Virtual DNS</b>	Configure this mode to support virtual DNS servers (VDNS) to resolve the DNS requests from the VMs. Each IPAM can have a virtual DNS server configured in this mode.

9. Complete the remaining fields on this page, and click **OK** to commit the changes, or click **Cancel** to clear the fields and start over.

## Configuring DNS Using Scripts

DNS can be configured via the user interface or by using scripts that are available in the `opt/contrail/utls` directory. The scripts are described in [Table 25 on page 177](#).



**CAUTION:** Be aware of the following cautions when using scripts to configure DNS:

- DNS doesn't allow special characters in the names, other than - (dash) and . (period). Any records that include special characters in the name will be discarded by the system.
- The IPAM DNS mode and association should only be edited when there are *no* virtual machine instances in the virtual networks associated with the IPAM.

Table 25: DNS Scripts

Action	Script
Add a virtual DNS server	Script: <code>add_virtual_dns.py</code>  Sample usage: <code>python add_virtual_dns.py --api_server_ip 10.204.216.21 --api_server_port 8082 --name vdns1 --domain_name default-domain --dns_domain juniper.net --dyn_updates --record_order random --ttl 1200 --next_vdns default-domain:vdns1</code>
Delete a virtual DNS server	Script: <code>del_virtual_dns_record.py</code>  Sample usage: <code>python del_virtual_dns.py --api_server_ip 10.204.216.21 --api_server_port 8082 --fq_name default-domain:vdns1</code>
Add a DNS record	Script: <code>add_virtual_dns_record.py</code>  Sample usage: <code>python add_virtual_dns_record.py --api_server_ip 10.204.216.21 --api_server_port 8082 --name rec1 --vdns_fqname default-domain:vdns1 --rec_name one --rec_type A --rec_class IN --rec_data 1.2.3.4 --rec_ttl 2400</code>

Table 25: DNS Scripts (*continued*)

Action	Script
Delete a DNS record	Script: <code>del_virtual_dns_record.py</code>  Sample usage: <code>python del_virtual_dns_record.py --api_server_ip 10.204.216.21 --api_server_port 8082 --fq_name default-domain:vdns1:rec1</code>
Associate a virtual DNS server with an IPAM	Script: <code>associate_virtual_dns.py</code>  Sample usage: <code>python associate_virtual_dns.py --api_server_ip 10.204.216.21 --api_server_port 8082 --ipam_fqname default-domain:demo:ipam1 --vdns_fqname default-domain:vdns1</code>
Disassociate a virtual DNS server with an IPAM	Script: <code>disassociate_virtual_dns.py</code>  Sample usage: <code>python disassociate_virtual_dns.py --api_server_ip 10.204.216.21 --api_server_port 8082 --ipam_fqname default-domain:demo:ipam1 --vdns_fqname default-domain:vdns1</code>

## Configuring Discovery Service

The Contrail Discovery Service publishes the IP address and port of the multiple components of the configuration node. The system runs multiple instances of each process for high availability and load balancing purposes.

- [Contrail Discovery Service Introduction on page 178](#)
- [Discovery Service Registration and Publishing on page 179](#)
- [Discovery Service Subscription on page 179](#)
- [Discovery Service REST API on page 180](#)
- [Discovery Service Heartbeats on page 182](#)
- [Discovery Service Internal Databases on page 182](#)
- [Discovery Service Client Library on page 182](#)
- [Discovery Service Debugging on page 182](#)

## Contrail Discovery Service Introduction

The following ports are used by the discovery service.

- API port : 5998 TCP
- Hearbeat port: 5998 TCP

To display the publishers, connect to the `http://discovery-server-ip:5998/services` URL.

To display the subscribers, connect to the `http://discovery-server-ip:5998/clients` URL.

The Contrail Discovery Service uses the following configuration file and log file:

```
/etc/contrail/discovery.conf
/var/log/contrail/discovery.log
```

## Discovery Service Registration and Publishing

The Discovery Service publishers send registration requests to the discovery server using a REST API.

The Discovery Service publishers send periodic heartbeat to the discovery server. The default interval for the heartbeat is 5 seconds.

If three successive heartbeats are missed, the Discovery Service is marked down.

The Discovery Service status is maintained internally. It indicates if the service is up or down based on the received heartbeat messages.

The discovery server currently supports three policies for selecting what information to return. The three policies are:

**Load Balance**—The service is returned based on the in-use count (how many subscribers are currently using the service).

**Round Robin**—The service is assigned based on a timestamp. The earliest (oldest) publisher is selected for the next assignment.

**Fixed**—An ordered list of available servers is always sent. If a service goes offline and comes back again, that service moves to the bottom of the list.

The three policies are configured in the `/etc/contrail/discovery.cfg` file under the service type section.

The response to a publish request is a cookie that must be sent back in the heartbeats.

## Discovery Service Subscription

Clients that need service send requests to the discovery server using a REST API.

The client can specify how many instances of a service to be returned. The default is 1. If the requested number of instances is 0, the information about all of the publishers of that service type is returned. Use this to display all the providers of a particular service.

A client is identified by a token (uuid). The token is typically sent as part of a subscription request. The client information is removed from the discovery server database when the time to live (TTL) expires.

A response to a client includes a TTL value. When the TTL expires, the client refreshes the information by sending another subscription request. The TTL sent to the client is a random value, in the range of 5 to 30 minutes.

If a service is overloaded and a new one is started, the new clients are automatically assigned a new service instance. To spray the new servers to the existing subscribers, use the `discovery_cli.py` file to reassign them on demand.

Clients find the discovery service by using the configured IP address and port.

## Discovery Service REST API

A REST API is available for registering and publishing the Contrail Discovery Server.

The following values are defined in the Contrail Discovery Server file:

- **POST: /publish or POST /publish/<publisher-id>**
- Content type: application/json or application/xml
- Body: information to be published (service type and data)

The following example shows the REST API for registering and publishing the discovery service

```
JSON simple:
{
  "control-node": {"ip_addr": "192.168.2.0", "port":1682 }
}
```

```
JSON verbose:
{
  "service-type" : "foobar",
  "foobar" : {"ip_addr": "192.168.2.0", "port":1682 }
}
```

```
XML simple:
<foobar2>
  <ip-addr>1.1.1.1</ip-addr>
  <port>4289</port>
</foobar2>
```

```
XML verbose:
<publish>
  <foobar2>
    <ip-addr>1.1.1.1</ip-addr>
    <port>4289</port>
  </foobar2>
  <oper-state>down</oper-state>
  <service-type>foobar2</service-type>
</publish>
```

```
JSON Response: {"cookie": c76716813f4b}
```

```
XML Response: <response><cookie>c76716813f4b</cookie></response>
```

The following fields are allowed in the body of the file:

service-type—Name of the service to publish

admin-state—Up or down state

remote-addr—IP address of the client

remote-version—Version number of the client

remote-name—Hostname of the client

oper-state—Each published service can set the oper-state up or down based on its internal state. You can display the reason the oper-state is up or down using the port 5998 URL.

A REST API is available for subscribing to the Contrail Discovery Server.

The following values are defined in the Contrail Discovery Server **discovery\_cli.py** file:

- POST http://discovery-server-ip:5998/subscribe
- Content-Type: application/json or application/xml
- Body: Service type, instance count, client ID

The following example shows the REST API for subscribing to the discovery service.

```
JSON: {
  "service": "control-node",
  "instances": 1,
  "client": "6c3f48bf-1098-46e8-8117-5cc745b45983",
  "remote-addr" : '1.1.1.'
}
XML:
<control-node>
  <instances>1</instances>
  <client>UUID</client>
  <remote-addr>1.1.1.1</remote-addr>
</control-node>

Response: TTL, List of <service type, Blob>
JSON: {
  "Apiservice": [{"ip_addr": "10.84.13.34", "port": "8082"}],
  "ttl": 357
}
XML:
<response>
  <ttl>300</ttl>
  <control-node>
    <ip_addr>192.168.2.0</ip_addr>
    <port>1682</port>
  </control-node>
</response>
```

The following fields are allowed in the body of the file:

Service Type—This is a string denoting what service is being requested (Apiservice). The instance count is the number of servers needed.

Client ID—This is a unique ID for the subscriber. Typically it is constructed from the UUID and the name of the subscriber.



**NOTE:** The subscription response includes a list of the services.

## Discovery Service Heartbeats

A cookie is returned in response to a request to publish API. The cookie is sent in the heartbeat message to the discovery server. If three heartbeat messages are missed, the discovery server marks the service down and it is no longer assigned to the subscribers.

The heartbeat responses from the discovery server are either 200 Ok or 401. The 401 response is sent if the discovery server does not recognize the cookie. This could happen if the discovery server is restarted with the `reset_config` option. In this case, the client should plan on republishing the information.

## Discovery Service Internal Databases

The database is maintained in Cassandra. There is a persistent copy so that the discovery service can maintain state across restarts.

## Discovery Service Client Library

Python and C++ client libraries are available that allow publishing and subscription of services.

## Discovery Service Debugging

To see a list of Discovery Service publishers, connect to the `http://discovery-server-ip:5998/services` URL.

To see list of Discovery Service subscribers, connect to the `http://discovery-server-ip:5998/clients` URL.

To see the log messages for the Discovery Service, display the `/var/log/contrail/discovery.log` file.

### Related Documentation

- *Configuring Load Balancing as a Service in Contrail*

---

## Support for Multicast

This section describes how the Contrail Controller supports broadcast and multicast.

- [Subnet Broadcast on page 183](#)
- [All-Broadcast/Limited-Broadcast and Link-Local Multicast on page 183](#)
- [Host Broadcast on page 184](#)



## Subnet Broadcast

Multiple subnets can be attached to a virtual network when it is spawned. Each of the subnets has one subnet broadcast route installed in the unicast routing table assigned to that virtual network. The recipient list for the subnet broadcast route includes all of the virtual machines that belong to that subnet. Packets originating from any VM in that subnet are replicated to all members of the recipient list, except the originator. Because the next hop is the list of recipients, it is called a composite next hop.

If there is no virtual machine spawned under a subnet, the subnet routing entry discards the packets received. If all of the virtual machines in a subnet are turned off, the routing entry points to discard. If the IPAM is deleted, the subnet route corresponding to that IPAM is deleted. If the virtual network is turned off, all of the subnet routes associated with the virtual network are removed.

### *Subnet Broadcast Example*

The following configuration is made:

Virtual network name – **vn1**

Unicast routing instance – **vn1.uc.inet**

Subnets (IPAM) allocated – **1.1.1.0/24; 2.2.0.0/16; 3.3.0.0/16**

Virtual machines spawned – **vm1 (1.1.1.253); vm2 (1.1.1.252); vm3 (1.1.1.251); vm4 (3.3.1.253)**

The following subnet route additions are made to the routing instance **vn1.uc.inet.0**:

**1.1.1.255** -> forward to NH1 (composite next hop)

**2.2.255.255** -> DROP

**3.3.255.255** -> forward to NH2

The following entries are made to the next-hop table:

NH1 – **1.1.1.253; 1.1.1.252; 1.1.1.251**

NH2 – **3.3.1.253**

If traffic originates for **1.1.1.255** from **vm1 (1.1.1.253)**, it will be forwarded to **vm2 (1.1.1.252)** and **vm3 (1.1.1.251)**. The originator **vm1 (1.1.1.253)** will not receive the traffic even though it is listed as a recipient in the next hop.

## All-Broadcast/Limited-Broadcast and Link-Local Multicast

The address group **255.255.255.255** is used with all-broadcast (limited-broadcast) and multicast traffic. The route is installed in the multicast routing instance. The source address is recorded as ANY, so the route is **ANY/255.255.255.255 (\*G)**. It is unique per routing instance, and is associated with its corresponding virtual network. When a virtual network is spawned, it usually contains multiple subnets, in which virtual machines are added. All of the virtual machines, regardless of their subnets, are part of the recipient list for **ANY/255.255.255.255**. The replication is sent to every recipient except the originator.

Link-local multicast also uses the all-broadcast method for replication. The route is deleted when all virtual machines in this virtual network are turned off or the virtual network itself is deleted.

#### *All-Broadcast Example*

The following configuration is made:

Virtual network name – **vn1**

Unicast routing instance – **vn1.uc.inet**

Subnets (IPAM) allocated – **1.1.1.0/24; 2.2.0.0/16; 3.3.0.0/16**

Virtual machines spawned – **vm1 (1.1.1.253); vm2 (1.1.1.252); vm3 (1.1.1.251); vm4 (3.3.1.253)**

The following subnet route addition is made to the routing instance **vn1.uc.inet.0**:

**255.255.255.255/\* -> NH1**

The following entries are made to the next-hop table:

**NH1 – 1.1.1.253; 1.1.1.252; 1.1.1.251; 3.3.1.253**

If traffic originates for **1.1.1.255** from **vm1 (1.1.1.253)**, the traffic is forwarded to **vm2 (1.1.1.252)**, **vm3 (1.1.1.251)**, and **vm4 (3.3.1.253)**. The originator **vm1 (1.1.1.253)** will not receive the traffic even though it is listed as a recipient in the next hop.

## Host Broadcast

The host broadcast route is present in the host routing instance so that the host operating system can send a subnet broadcast/all-broadcast (limited-broadcast). This type of broadcast is sent to the fabric by means of a **vhost** interface. Additionally, any subnet broadcast/all-broadcast received from the fabric will be handed over to the host operating system.

## Using Static Routes with Services

---

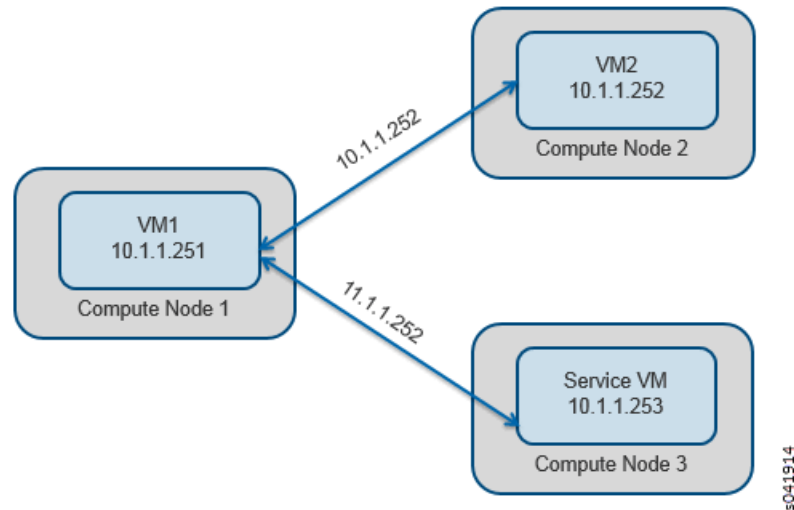
- [Static Routes for Service Instances on page 184](#)
- [Configuring Static Routes on a Service Instance on page 185](#)
- [Configuring Static Routes on Service Instance Interfaces on page 186](#)
- [Configuring Static Routes as Host Routes on page 188](#)

## Static Routes for Service Instances

Static routes can be configured in a virtual network to direct traffic to a service virtual machine.

The following figure shows a virtual network with subnet 10.1.1.0/24. All of the traffic from a virtual machine that is directed to subnet 11.1.1.0/24 can be configured to be routed by

means of a service machine, by using the static route 11.1.1.252 configured on the service virtual machine interface.



### Configuring Static Routes on a Service Instance

To configure static routes on a service instance, first enable the static route option in the service template to be used for the service instance.

To enable the static route option in a service template:

1. Go to **Configure > Services > Service Templates** and click **Create**.
2. At **Add Service Template**, complete the fields for **Name**, **Service Mode**, and **Image Name**.
3. Select the **Interface Types** to use for the template, then for each interface type that might have a static route configured, click the check box under the **Static Routes** column to enable the static route option for that interface.

The following figure shows a service template in which the left and right interfaces of service instances have the static routes option enabled. Now a user can configure

a static route on a corresponding interface on a service instance that is based on the service template shown.

**Add Service Template**

Name:

Service Mode:

Image Name:

Interface Types	Shared IP	Static Routes	+
Management	<input type="checkbox"/>	<input type="checkbox"/>	+ -
Left	<input type="checkbox"/>	<input checked="" type="checkbox"/>	+ -
Right	<input type="checkbox"/>	<input checked="" type="checkbox"/>	+ -

► [Advanced options](#)

Cancel Save

s041915

## Configuring Static Routes on Service Instance Interfaces

To configure static routes on a service instance interface:

1. Go to **Configure > Services > Service Instances** and click **Create**.
2. At **Create Service Instances**, complete the fields for **Instance Name** and **Services Template**.
3. Select the virtual network for each of the interfaces
4. Click the **Static Routes** dropdown menu under each interface field for which the static routes option is enabled to open the **Static Routes** menu and configure the static routes in the fields provided.



**NOTE:** If the **Auto Configured** option is selected, traffic destined to the static route subnet is load balanced across service instances.

The following figure shows a configuration to apply a service instance between VN1 (10.1.1.0/24) and VN2 (11.1.1.0/24). The left interface of the service instance is configured with VN1 and the right interface is configured to be VN2 (11.1.1.0/24). The static route

11.1.1.0/24 is configured on the left interface, so that all traffic from VN1 that is destined to VN2 reaches the left interface of the service instance.

The screenshot shows the 'Create Service Instances' dialog box. The 'Instance Name' is 'nat'. The 'Services Template' is 'nat - [in-network (management, left, right)]'. Under 'Interface 1', 'Management' is selected with 'Auto Configured'. Under 'Interface 2', 'Left' is selected with 'vn1'. A 'Static Routes' section for Interface 2 shows a table with one entry: Prefix '11.1.1.0/24' and Next hop 'Interface 2'. Under 'Interface 3', 'Right' is selected with 'vn2'. A 'Static Routes' section for Interface 3 is empty. The dialog has 'Cancel' and 'Save' buttons at the bottom right. A vertical label 's041916' is on the right edge.

The following figure shows static route 10.1.1.0/24 configured on the right interface, so that all traffic from VN2 that is destined to VN1 reaches the right interface of the service virtual machine.

The screenshot shows the 'Create Service Instances' dialog box. The 'Instance Name' is 'nat'. The 'Services Template' is 'nat - [in-network (management, left, right)]'. Under 'Interface 1', 'Management' is selected with 'Auto Configured'. Under 'Interface 2', 'Left' is selected with 'vn1'. A 'Static Routes' section for Interface 2 shows a table with one entry: Prefix '11.1.1.0/24' and Next hop 'Interface 2'. Under 'Interface 3', 'Right' is selected with 'vn2'. A 'Static Routes' section for Interface 3 shows a table with one entry: Prefix '10.1.1.0/24' and Next hop 'Interface 3'. The dialog has 'Cancel' and 'Save' buttons at the bottom right. A vertical label 's041917' is on the right edge.

When the static routes are configured for both the left and the right interfaces, all inter-virtual network traffic is forwarded through the service instance.

## Configuring Static Routes as Host Routes

You can also use static routes for host routes for a virtual machine, by using the classless static routes option in the DHCP server response that is sent to the virtual machine.

The routes to be sent in the DHCP response to the virtual machine can be configured for each virtual network as it is created.

To configure static routes as host routes:

1. Go to **Configure > Network > Networks** and click **Create**.
2. At **Create Network**, click the **Host Routes** option and add the host routes to be sent to the virtual machines.

An example is shown in the following figure.

The screenshot shows the 'Create Network' dialog box. At the top, there's a section for 'Address Management' with a dropdown menu set to 'ipam1', and fields for 'IP Block' and 'Gateway' with '+' and '-' buttons. Below this is a table with three columns: 'IPAM', 'IP Block', and 'Gateway'. The first row contains 'ipam1', '1.2.3.0/24', and '1.2.3.254'. Below the table are three expandable sections: 'Route Targets', 'Floating IP Pools', and 'Host Routes'. The 'Host Routes' section is expanded, showing a table with two columns: 'IPAM' and 'Route Prefix'. There are two rows in this table, both with 'ipam1' in the 'IPAM' column. The first row has '1.1.1.0/24' in the 'Route Prefix' column, and the second row has '2.2.2.0/24' in the 'Route Prefix' column. Each row has '+' and '-' buttons to its right. At the bottom right of the dialog are 'Cancel' and 'Save' buttons. A small vertical text 's041918' is visible on the right side of the dialog.

## Configuring Metadata Service

OpenStack enables virtual machines to access metadata by sending an HTTP request to the link-local address 169.254.169.254. The metadata request from the virtual machine is proxied to Nova with additional HTTP header fields that Nova uses to identify the source instance, then responds with appropriate metadata.

In Contrail, the vRouter acts as the proxy, by trapping the metadata requests, adding the necessary header fields, and sending the requests to the Nova API server.

The metadata service is configured by setting the **linklocal-services** property on the **global-vrouter-config** object.

Use the following elements to configure the **linklocal-services** element for metadata service:

- **linklocal-service-name** = **metadata**
- **linklocal-service-ip** = **169.254.169.254**
- **linklocal-service-port** = **80**
- **ip-fabric-service-ip** = *[server-ip-address]*
- **ip-fabric-service-port** = *[server-port]*

The **linklocal-services** properties can be set from the Contrail UI (**Configure > Infrastructure > Link Local Services**) or by using the following command:

```
python /opt/contrail/utils/provision_linklocal.py --admin_user <user> --admin_password  
<passwd> --linklocal_service_name metadata --linklocal_service_ip 169.254.169.254  
--linklocal_service_port 80 --ipfabric_service_ip --ipfabric_service_port 8775
```





## CHAPTER 10

# Configuring Service Chaining

- [Service Chaining on page 191](#)
- [Service Chaining MX Series Configuration on page 195](#)
- [ECMP Load Balancing in the Service Chain on page 196](#)
- [Customized Hash Field Selection for ECMP Load Balancing on page 197](#)
- [Using the Contrail Heat Template on page 201](#)
- [Service Chain Route Reorigination on page 204](#)
- [Service Instance Health Check on page 223](#)

## Service Chaining

---

Contrail Controller supports chaining of various Layer 2 through Layer 7 services such as firewall, NAT, IDP, and so on.


- [Service Chaining Basics on page 191](#)
- [Service Chaining Configuration Elements on page 192](#)

## Service Chaining Basics

Services are offered by instantiating service virtual machines to dynamically apply single or multiple services to virtual machine (VM) traffic. It is also possible to chain physical appliance-based services.


[Figure 64 on page 191](#) shows the basic service chain schema, with a single service. The service VM spawns the service, using the convention of left interface (left IF) and right interface (right IF). Multiple services can also be chained together.

**Figure 64: Service Chaining**



When you create a service chain, the Contrail software creates tunnels across the underlay network that span through all services in the chain. [Figure 65 on page 192](#) shows two end points and two compute nodes, each with one service instance and traffic going to and from one end point to the other.

**Figure 65: Contrail Service Chain**



The following are the modes of services that can be configured.

*Transparent or bridge mode*

Used for services that do not modify the packet. Also known as bump-in-the-wire or Layer 2 mode. Examples include Layer 2 firewall, IDP, and so on.

*In-network or routed mode*

Provides a gateway service where packets are routed between the service instance interfaces. Examples include NAT, Layer 3 firewall, load balancer, HTTP proxy, and so on.

*In-network-nat mode*

Similar to in-network mode, however, return traffic does not need to be routed to the source network. In-network-nat mode is particularly useful for NAT service.

## Service Chaining Configuration Elements

Service chaining requires the following configuration elements in the solution:

- Service template
- Service instance

- Service policy

### *Service Template*

Service templates are always configured in the scope of a domain, and the templates can be used on all projects within a domain. A template can be used to launch multiple service instances in different projects within a domain.

The following are the parameters to be configured for a service template:

- Service template name
- Domain name
- Service mode
  - Transparent
  - In-Network
  - In-Network NAT
- Image name (for virtual service)
  - If the service is a virtual service, then the name of the image to be used must be included in the service template. In an OpenStack setup, the image must be added to the setup by using Glance.
- Interface list
  - Ordered list of interfaces---this determines the order in which Interfaces will be created on the service instance.
  - Most service templates will have management, left, and right interfaces. For service instances requiring more interfaces, "other" interfaces can be added to the interface list.
  - Shared IP attribute, per interface
  - Static routes enabled attribute, per interface
- Advanced options
  - Service scaling— use this attribute to enable a service instance to have more than one instance of the service instance virtual machine.
  - Flavor—assign an OpenStack flavor to be used while launching the service instance. Flavors are defined in OpenStack Nova with attributes such as assignments of CPU cores, memory, and disk space.

### *Service Instance*

A service instance is always maintained within the scope of a project. A service instance is launched using a specified service template from the domain to which the project belongs.

The following are the parameters to be configured for a service instance:

- Service instance name
- Project name
- Service template name
- Number of virtual machines that will be spawned
  - Enable service scaling in the service template for multiple virtual machines
- Ordered virtual network list
  - Interfaces listed in the order specified in the service template
  - Identify virtual network for each interface
  - Assign static routes for virtual networks that have static route enabled in the service template for their interface
    - Traffic that matches an assigned static route is directed to the service instance on the interface created for the corresponding virtual network

#### *Service Policy*

The following are the parameters to be configured for a service policy:

- Policy name
- Source network name
- Destination network name
- Other policy match conditions, for example direction and source and destination ports
- Policy configured in “routed/in-network” or “bridged/” mode
- An action type called **apply\_service** is used:

Example: 'apply\_service': [DomainName:ProjectName:ServiceInstanceName]

#### **Related Documentation**

- [Example: Creating an In-Network or In-Network-NAT Service Chain on page 227](#)
- [Example: Creating a Service Chain With the CLI on page 239](#)
- [ECMP Load Balancing in the Service Chain on page 196](#)

## Service Chaining MX Series Configuration

This topic shows how to extend service chaining to the MX Series routers.

To configure service chaining for MX Series routers, extend the virtual networks to the MX Series router and program routes so that traffic generated from a host connected to the router can be routed through the service.

1. The following configuration snippet for an MX Series router has a left virtual network called **enterprise** and a right virtual network called **public**. The configuration creates two routing instances with loopback interfaces and route targets.

```
routing-instances {
  enterprise {
    instance-type vrf;
    interface lo0.1;
    vrf-target target:100:20000;
  }
  public {
    instance-type vrf;
    interface lo0.2;
    vrf-target target:100:10000;
  }
  routing-options {
    static {
      route 0.0.0.0/0 next-hop 10.84.20.1
    }
  }
  interface xe-0/0/0.0;
}
```

2. The following configuration snippet shows the configuration for the loopback interfaces.

```
interfaces {
  lo0 {
    unit 1 {
      family inet {
        address 2.1.1.100/32;
      }
    }
    unit 2 {
      family inet {
        address 200.1.1.1/32;
      }
    }
  }
}
```

3. The following configuration snippet shows the configuration to enable BGP. The **neighbor 10.84.20.39** and **neighbor 10.84.20.40** are control nodes.

```
protocols {
  bgp {
    group demo_contrail {
      type internal;
      description "To Contrail Control Nodes & other MX";
      local-address 10.84.20.252;
    }
  }
}
```

```

    keep all;
    family inet-vpn {
        unicast;
    }
    neighbor 10.84.20.39;
    neighbor 10.84.20.40;
}
}

```


4. The final step is to add **target:100:10000** to the public virtual network and **target:100:20000** to the enterprise virtual network, using the Contrail Juniper Networks interface.

A full MX Series router configuration for Contrail can be seen in “[Sample Network Configuration for Devices for Simple Tiered Web Application](#)” on page 163.

## ECMP Load Balancing in the Service Chain

Traffic flowing through a service chain can be load-balanced by distributing traffic streams to multiple service virtual machines (VMs) that are running identical applications. This is illustrated in [Figure 66 on page 196](#), where the traffic streams between VM-A and VM-B are distributed between Service VM-1 and Service VM-2. If Service VM-1 goes down, then all streams that are dependent on Service VM-1 will be moved to Service VM-2.

Figure 66: Load Balancing a Service Chain



s041830

The following are the major features of load balancing in the service chain:

- Load balancing can be configured at every level of the service chain.
- Load balancing is supported in routed and bridged service chain modes.
- Load balancing can be used to achieve high availability—if a service VM goes down, the traffic passing through that service VM can be distributed through another service VM.
- A load balanced traffic stream always follows the same path through the chain of service VM.

- Related Documentation**
- [Service Chaining on page 191](#)
  - [Example: Creating a Service Chain With the CLI on page 239](#)
  - [Customized Hash Field Selection for ECMP Load Balancing on page 197](#)

---

## Customized Hash Field Selection for ECMP Load Balancing

### Overview: Custom Hash Feature

Starting with Contrail Release 3.0, it is possible to configure the set of fields used to hash upon during equal-cost multipath (ECMP) load balancing.

Earlier versions of Contrail had this set of fields fixed to the standard 5-tuple set of: source L3 address, destination L3 address, L4 protocol, L4 SourcePort, and L4 DestinationPort.

With the custom hash feature, users can configure an exact subset of fields to hash upon when choosing the forwarding path among a set of eligible ECMP candidates.

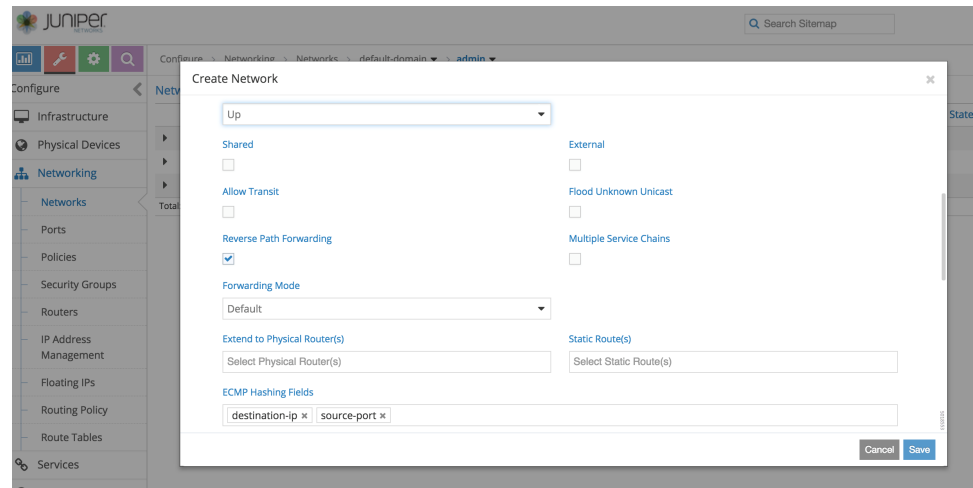
The custom hash configuration can be applied in the following ways:

- globally
- per virtual network (VN)
- per virtual network interface (VNI)

VNI configurations take precedence over VN configurations, and VN configurations take precedence over global level configuration (if present).

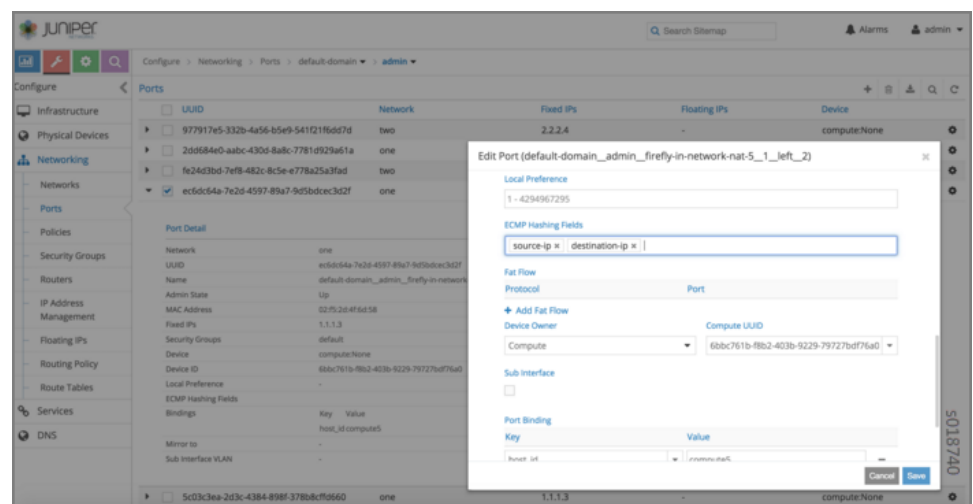
Custom hash is useful whenever packets originating from a particular source and addressed to a particular destination must go through the same set of service instances during transit. This might be required if source, destination, or transit nodes maintain a certain state based on the flow, and the state behavior could also be used for subsequent new flows, between the same pair of source and destination addresses. In such cases, subsequent flows must follow the same set of service nodes followed by the initial flow.

You can use the Contrail UI to identify specific fields in the network upon which to hash at the **Configure > Networking > Network, Create Network** window, in the **ECMP Hashing Fields** section as shown in the following figure.



If the hashing fields are configured for a virtual network, all traffic destined to that VN will be subject to the customized hash field selection during forwarding over ECMP paths by vRouters. This may not be desirable in all cases, as it could potentially skew all traffic to the destination network over a smaller set of paths across the IP fabric.

A more practical scenario is one in which flows between a source and destination must go through the same service instance in between, where one could configure customized ECMP fields for the virtual machine interface (VMI) of the service instance. Then, each service chain route originating from that VMI would get the desired ECMP field selection applied as its path attribute, and eventually get propagated to the ingress vRouter node. See the following example.





## Using ECMP Hash Fields Selection

Custom hash fields selection is most useful in scenarios where multiple ECMP paths exist for a destination. Typically, the multiple ECMP paths point to ingress service instance nodes, which could be running anywhere in the Contrail cloud.

### Configuring ECMP Hash Fields Over Service Chains

Use the following steps to create customized hash fields with ECMP over service chains.

1. Create the virtual networks needed to interconnect using service chaining, with ECMP load-balancing.
2. Create a service template and enable scaling.
3. Create a service instance, and using the service template, configure by selecting:
  - the desired number of instances for scale-out
  - the left and right virtual network to connect
  - the shared address space, to make sure that instantiated services come up with the same IP address for left and right, respectively

This configuration enables ECMP among all those service instances during forwarding.

4. Create a policy, then select the service instance previously created and apply the policy to to the desired VMIs or VNs.
5. After the service VMs are instantiated, the ports of the left and right interfaces are available for further configuration. At the Contrail UI Ports section under Networking, select the left port (VMI) of the service instance and apply the desired ECMP hash field configuration.



**NOTE:** Currently the ECMP field selection configuration for the service instance left or right interface must be applied by using the Ports (VMIs) section under Networking and explicitly configuring the ECMP fields selection for each of the instantiated service instances' VMIs. This must be done for all service interfaces of the group, to ensure the end result is as expected, because the load balance attribute of only the best path is carried over to the ingress vRouter. If the load balance attribute is not configured, it is not propagated to the ingress vRouter, even if other paths have that configuration.

When the configuration is finished, the vRouters get programmed with routing tables with the ECMP paths to the various service instances. The vRouters are also programmed with the desired ECMP hash fields to be used during load balancing of the traffic.

## Sample Flows

This section provides sample flows with and without ECMP custom hash field selection.

### Sample Traffic Flow Path Without Custom ECMP Hash Fields

The following is an example of a traffic flow path without using a customized ECMP hash fields selection configuration. The flow is configured with standard 5-tuple flow fields.

```
tcpdump -i eth0 'port 1023 and tcp[tcpflags] & (tcp-syn) != 0 and tcp[tcpflags]
& (tcp-ack) == 0'
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 65535 bytes
14:55:10.115122 IP 2.2.2.5.18337 > 2.2.2.100.1023: Flags [S], seq 2276852196, win
29200, options [mss 1398,sackOK,TS val 25208882 ecr 0,nop,wscale 7], length 0
14:55:10.132753 IP 2.2.2.4.21193 > 2.2.2.100.1023: Flags [S], seq 4161487314, win
29200, options [mss 1398,sackOK,TS val 25208886 ecr 0,nop,wscale 7], length 0
14:55:10.152053 IP 2.2.2.5.24230 > 2.2.2.100.1023: Flags [S], seq 2466454857, win
29200, options [mss 1398,sackOK,TS val 25208892 ecr 0,nop,wscale 7], length 0
14:55:11.146029 IP 2.2.2.5.24230 > 2.2.2.100.1023: Flags [S], seq 2466454857, win
29200, options [mss 1398,sackOK,TS val 25209142 ecr 0,nop,wscale 7], length 0
14:55:13.147616 IP 2.2.2.5.24230 > 2.2.2.100.1023: Flags [S], seq 2466454857, win
29200, options [mss 1398,sackOK,TS val 25209643 ecr 0,nop,wscale 7], length 0
14:55:13.164367 IP 2.2.2.3.25582 > 2.2.2.100.1023: Flags [S], seq 2259034580, win
29200, options [mss 1398,sackOK,TS val 25209644 ecr 0,nop,wscale 7], length 0
14:55:13.179939 IP 2.2.2.5.24895 > 2.2.2.100.1023: Flags [S], seq 2174031724, win
29200, options [mss 1398,sackOK,TS val 25209648 ecr 0,nop,wscale 7], length 0
14:55:14.168282 IP 2.2.2.5.24895 > 2.2.2.100.1023: Flags [S], seq 2174031724, win
29200, options [mss 1398,sackOK,TS val 25209898 ecr 0,nop,wscale 7], length 0
14:55:16.172384 IP 2.2.2.5.24895 > 2.2.2.100.1023: Flags [S], seq 2174031724, win
29200, options [mss 1398,sackOK,TS val 25210399 ecr 0,nop,wscale 7], length 0
14:55:16.189864 IP 2.2.2.5.22952 > 2.2.2.100.1023: Flags [S], seq 3099816842, win
29200, options [mss 1398,sackOK,TS val 25210401 ecr 0,nop,wscale 7], length 0
14:55:16.205142 IP 2.2.2.4.16487 > 2.2.2.100.1023: Flags [S], seq 3961114202, win
29200, options [mss 1398,sackOK,TS val 25210405 ecr 0,nop,wscale 7], length 0
14:55:17.196763 IP 2.2.2.4.16487 > 2.2.2.100.1023: Flags [S], seq 3961114202, win
29200, options [mss 1398,sackOK,TS val 25210655 ecr 0,nop,wscale 7], length 0
14:55:19.200623 IP 2.2.2.4.16487 > 2.2.2.100.1023: Flags [S], seq 3961114202, win
29200, options [mss 1398,sackOK,TS val 25211156 ecr 0,nop,wscale 7], length 0
14:55:19.215809 IP 2.2.2.3.18914 > 2.2.2.100.1023: Flags [S], seq 3157557440, win
29200, options [mss 1398,sackOK,TS val 25211158 ecr 0,nop,wscale 7], length 0
14:55:19.228405 IP 2.2.2.7.15569 > 2.2.2.100.1023: Flags [S], seq 3850648420, win
29200, options [mss 1398,sackOK,TS val 25211161 ecr 0,nop,wscale 7], length 0
14:55:20.223482 IP 2.2.2.7.15569 > 2.2.2.100.1023: Flags [S], seq 3850648420, win
29200, options [mss 1398,sackOK,TS val 25211412 ecr 0,nop,wscale 7], length 0
14:55:22.232068 IP 2.2.2.7.15569 > 2.2.2.100.1023: Flags [S], seq 3850648420, win
29200, options [mss 1398,sackOK,TS val 25211913 ecr 0,nop,wscale 7], length 0
14:55:22.247325 IP 2.2.2.4.28388 > 2.2.2.100.1023: Flags [S], seq 3609240658, win
29200, options [mss 1398,sackOK,TS val 25211915 ecr 0,nop,wscale 7], length 0
```

### Sample Traffic Flow Path With Custom ECMP Hash Fields

The following is an example of a traffic flow path using a customized ECMP hash fields selection configuration, for **source-ip** and **destination-ip** only.

```
tcpdump -i eth0 'port 1023 and tcp[tcpflags] & (tcp-syn) != 0 and tcp[tcpflags]
& (tcp-ack) == 0'
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 65535 bytes
15:57:18.680853 IP 2.2.2.4.21718 > 2.2.2.100.1023: Flags [S], seq 2052086108, win
29200, options [mss 1398,sackOK,TS val 26141024 ecr 0,nop,wscale 7], length 0
15:57:18.696114 IP 2.2.2.4.13585 > 2.2.2.100.1023: Flags [S], seq 2039627277, win
29200, options [mss 1398,sackOK,TS val 26141028 ecr 0,nop,wscale 7], length 0
15:57:18.714846 IP 2.2.2.4.16414 > 2.2.2.100.1023: Flags [S], seq 3252526560, win
```

```

29200, options [mss 1398,sackOK,TS val 26141033 ecr 0,nop,wscale 7], length 0
15:57:18.731281 IP 2.2.2.4.32499 > 2.2.2.100.1023: Flags [S], seq 1389133175, win
29200, options [mss 1398,sackOK,TS val 26141037 ecr 0,nop,wscale 7], length 0
15:57:18.747051 IP 2.2.2.4.6081 > 2.2.2.100.1023: Flags [S], seq 427936299, win
29200, options [mss 1398,sackOK,TS val 26141041 ecr 0,nop,wscale 7], length 0
15:57:19.740204 IP 2.2.2.4.6081 > 2.2.2.100.1023: Flags [S], seq 427936299, win
29200, options [mss 1398,sackOK,TS val 26141291 ecr 0,nop,wscale 7], length 0
15:57:21.743951 IP 2.2.2.4.6081 > 2.2.2.100.1023: Flags [S], seq 427936299, win
29200, options [mss 1398,sackOK,TS val 26141792 ecr 0,nop,wscale 7], length 0
15:57:21.758532 IP 2.2.2.4.13800 > 2.2.2.100.1023: Flags [S], seq 3020971712, win
29200, options [mss 1398,sackOK,TS val 26141794 ecr 0,nop,wscale 7], length 0
15:57:21.772646 IP 2.2.2.4.23894 > 2.2.2.100.1023: Flags [S], seq 3373734307, win
29200, options [mss 1398,sackOK,TS val 26141797 ecr 0,nop,wscale 7], length 0
15:57:22.764469 IP 2.2.2.4.23894 > 2.2.2.100.1023: Flags [S], seq 3373734307, win
29200, options [mss 1398,sackOK,TS val 26142047 ecr 0,nop,wscale 7], length 0
15:57:24.768511 IP 2.2.2.4.23894 > 2.2.2.100.1023: Flags [S], seq 3373734307, win
29200, options [mss 1398,sackOK,TS val 26142548 ecr 0,nop,wscale 7], length 0
15:57:24.784119 IP 2.2.2.4.21858 > 2.2.2.100.1023: Flags [S], seq 2212369297, win
29200, options [mss 1398,sackOK,TS val 26142550 ecr 0,nop,wscale 7], length 0
15:57:24.797149 IP 2.2.2.4.29440 > 2.2.2.100.1023: Flags [S], seq 2007897735, win
29200, options [mss 1398,sackOK,TS val 26142554 ecr 0,nop,wscale 7], length 0
15:57:25.792816 IP 2.2.2.4.29440 > 2.2.2.100.1023: Flags [S], seq 2007897735, win
29200, options [mss 1398,sackOK,TS val 26142804 ecr 0,nop,wscale 7], length 0
15:57:27.797538 IP 2.2.2.4.29440 > 2.2.2.100.1023: Flags [S], seq 2007897735, win
29200, options [mss 1398,sackOK,TS val 26143305 ecr 0,nop,wscale 7], length 0
15:57:27.814002 IP 2.2.2.4.23452 > 2.2.2.100.1023: Flags [S], seq 1659332655, win
29200, options [mss 1398,sackOK,TS val 26143307 ecr 0,nop,wscale 7], length 0

```

## Using the Contrail Heat Template

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 on page 201](#)
- [Heat Architecture on page 202](#)
- [Support for Heat Version 2 Resources on page 202](#)
- [Heat Version 2 with Service Chaining and Port Tuple Sample Workflow on page 203](#)
- [Example: Creating a Service Template Using Heat on page 203](#)

## 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** in Contrail Release 3.0.2 and greater:

<https://github.com/Juniper/contrail-heat/tree/master/generated/resources>

The following contains a list of new example templates:

[https://github.com/Juniper/contrail-heat/tree/master/contrail\\_heat/new\\_templates](https://github.com/Juniper/contrail-heat/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 chains using the Heat Version 1 templates.

## Heat Version 2 with Service Chaining 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. Define an environment file to give input to the Heat template.

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 using the following command:

```
heat stack- create stack1 -f service_template.yaml -e service_template.env
```

**Related Documentation**

- *Service Chain Version 2 with Port Tuple*

---

## Service Chain Route Reorigination

- [Overview: Service Chaining in Contrail on page 205](#)
- [Route Aggregation on page 206](#)

- [Routing Policy on page 213](#)
- [Control for Route Reorigination on page 221](#)

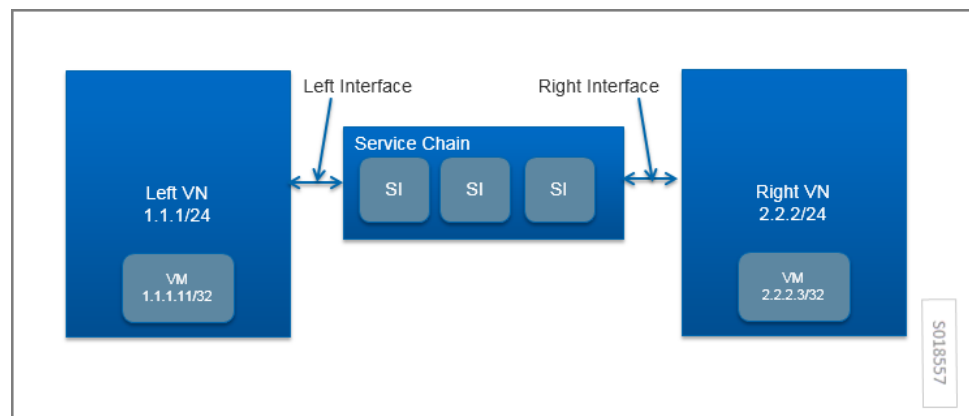
## Overview: Service Chaining in Contrail

In Contrail, the service chaining feature allows the operator to insert dynamic services to control the traffic between two virtual networks. The service chaining works on a basic rule of next-hop stitching.

In [Figure 67 on page 205](#), the service chain is inserted between the Left VN and the Right VN. The service chain contains one or more service instances to achieve a required network policy.

In the example, the route for the VM in the Right VN is added to the routing table for the Left VN, with the next hop modified to ensure that the traffic is sent by means of the left interface of the service chain. This is an example of route reorigination.

**Figure 67: Route Reorigination**



Using reorigination of routes for service chaining (for example, putting the route for the right network in the left routing table) requires the following features:

- **Route aggregation**

For scaling purposes, it is useful to publish an aggregated route as the service chain route, rather than publishing every route of each VM (/32). This reduces the memory footprint for the route table in the gateway router and also reduces route exchanges between control nodes and the gateway router. The route can be aggregated to the default route (0/0), to the VN subnet prefix, or to any arbitrary route prefix.

- **Path attribute modification for reoriginated routes**

There are cases where the **BgpPath** attribute for the service chain route needs to be modified. An example is the case of service chain failover, in which there are two service chains with identical services that are connected between the same two VNs. The operator needs to control which service chain is used for traffic between two networks, in addition to ensuring redundancy and high availability by providing failover support. Path attribute modification for reoriginated routes is implemented by means of routing

policy, by providing an option to alter the MED (multi-exit discriminator) or **local-pref** of the reoriginated service chain route.

- **Control to enable and disable reorigination of the route**

In some scenarios, the operator needs a control to stop reorigination of the route as the service chain route, for example, when static routes are configured on service VM interfaces. Control to enable or disable reorigination of the route is implemented by tagging the routes with the **no-reoriginate** community. Routes with the **no-reoriginate** community tag are skipped for route reorigination.

## Route Aggregation

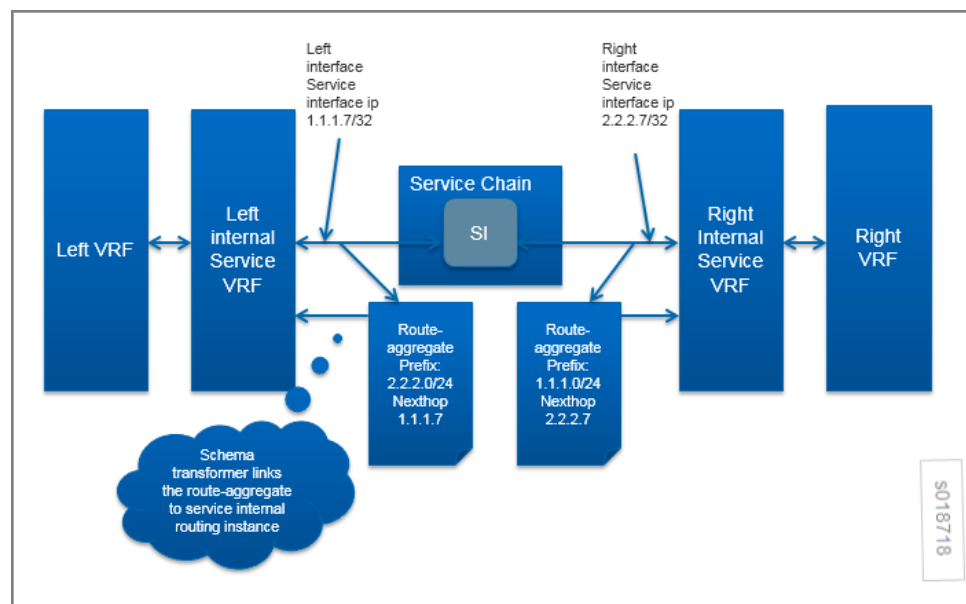
The route aggregation configuration object contains a list of prefixes to aggregate. The next-hop field in the route aggregate object contains the address of the route whose next hop is stitched as a next hop of the aggregate route.

Route aggregation is configured on the service instance. The operator can attach multiple route aggregation objects to a service instance. For example, if routes from the Right VN need to be aggregated and reoriginated in the route table of the Left VN, the route aggregate object is created with a prefix of the Right VN's subnet prefix and attached to the left interface of the service instance.

If the service chain has multiple service instances, the route aggregate object is attached to the left interface of the left-most service instance and to the right interface of the right-most service instance.

The relationships are shown in [Figure 68 on page 206](#).

**Figure 68: Route Aggregate Relationships**





The schema transformer sets the next-hop field of the route aggregate object to the service chain interface address. The schema transformer also links the route aggregate object to the internal routing instance created for the service instance.

Using the configuration as described, the Contrail control service reads the route aggregation object on the routing instance. When the first, more specific route or contributing route is launched (when the first VM is launched on the right VN), the aggregate route is published. Similarly, the aggregated route is deleted when the last, more specific route or contributing route is deleted (when the last VM is deleted in the right VN). The aggregated route is published when the next hop for the aggregated route gets resolved.

By default, in BGP or XMPP route exchanges, the control node will not publish contributing routes of an aggregate route.

### Schema for Route Aggregation

- [Route Aggregate Object on page 207](#)
- [Service Instance Link to Route Aggregate Object on page 207](#)
- [Routing Instance Link to Route Aggregate Object on page 208](#)

#### Route Aggregate Object

The following is the schema for route aggregate objects. Multiple prefixes can be specified in a single route aggregate object.

```
<xsd:element name="route-aggregate" type="ifmap:IdentityType"/>
<xsd:complexType name="RouteListType">
  <xsd:element name="route" type="xsd:string" maxOccurs="unbounded"/>
</xsd:complexType>

<xsd:element name='aggregate-route-entries' type='RouteListType'/>
<!--#IFMAP-SEMANTICS-IDL
  Property('aggregate-route-entries', 'route-aggregate') -->

<xsd:element name='aggregate-route-nexthop' type='xsd:string'/>
<!--#IFMAP-SEMANTICS-IDL
  Property('aggregate-route-nexthop', 'route-aggregate') -->
```

#### Service Instance Link to Route Aggregate Object

The following is the schema for the service instance link to route aggregation objects. The operator can link multiple route aggregate objects to a single service interface.

```
<xsd:element name="route-aggregate" type="ifmap:IdentityType"/>
<xsd:complexType name="RouteListType">
  <xsd:element name="route" type="xsd:string" maxOccurs="unbounded"/>
</xsd:complexType>

<xsd:element name='aggregate-route-entries' type='RouteListType'/>
<!--#IFMAP-SEMANTICS-IDL
  Property('aggregate-route-entries', 'route-aggregate') -->

<xsd:element name='aggregate-route-nexthop' type='xsd:string'/>
<!--#IFMAP-SEMANTICS-IDL
  Property('aggregate-route-nexthop', 'route-aggregate') -->
```

```
<xsd:simpleType name="ServiceInterfaceType">
  <xsd:restriction base="xsd:string">
    <xsd:pattern value="management|left|right|other[0-9]*"/>
  </xsd:restriction>
</xsd:simpleType>
<xsd:complexType name='ServiceInterfaceTag'>
  <xsd:element name="interface-type" type="ServiceInterfaceType"/>
</xsd:complexType>

<xsd:element name="route-aggregate-service-instance"
type="ServiceInterfaceTag"/>
<!--#IFMAP-SEMANTICS-IDL
  Link('route-aggregate-service-instance',
    'bgp:route-aggregate', 'service-instance', ['ref']) -->
```

### ***Routing Instance Link to Route Aggregate Object***

The following is the schema for the routing instance link to the route aggregation object. A routing instance can be linked to multiple route aggregate objects to perform route aggregation for multiple route prefixes.

```
<xsd:element name="route-aggregate-routing-instance"/>
<!--#IFMAP-SEMANTICS-IDL
  Link('route-aggregate-routing-instance',
    'route-aggregate', 'routing-instance', ['ref']) -->
```

---

## **Configuring and Troubleshooting Route Aggregation**

- [Configure Route Aggregate Object on page 208](#)
- [Configuring a Service Instance on page 209](#)
- [Create a Virtual Network and Network Policy on page 210](#)
- [Validate the Route Aggregate Object in the API Server on page 211](#)
- [Validate the Route Aggregate Object in the Control Node on page 212](#)

### ***Configure Route Aggregate Object***

You can use the Contrail UI **Create Route Aggregate** screen to name the route aggregate object and identify the routes to aggregate. See [Figure 69 on page 209](#).

Figure 69: Create Route Aggregate

The screenshot shows a 'Create Route Aggregate' window. The 'Name' field is filled with 'left-to-right'. Below it, the 'Aggregate Route Entries' section is expanded, revealing a 'Route' field containing '1.1.1.0/24'. A '+ Route Entry' button is located below the route list. At the bottom right of the window are 'Cancel' and 'Save' buttons. A vertical label '5018719' is positioned on the right side of the dialog.

**Example VNC Script to Create a Route Aggregate Object**

You can use a VNC script to create a route aggregate object, as in the following example:

```
from vnc_api.vnc_api import *
vnc_lib = VncApi("admin", "<password>", "admin")
project=vnc_lib.project_read(fq_name=["default-domain", "admin"])
route_aggregate=RouteAggregate(name="left_to_right", parent_obj=project)
route_list=RouteListType(["<ip address>"])
route_aggregate.set_aggregate_route_entries(route_list)
vnc_lib.route_aggregate_create(route_aggregate)
```

**Configuring a Service Instance**

Create a service instance with the route aggregate object linked to the aggregate left network subnet prefix in the right virtual network. See the example in [Figure 70 on page 210](#).

Figure 70: Create Service Instance

**Create Service Instance**

si-aggregate st-with-aggregate - [transparent (left, right)...

▼ **Interface Details**

Interface Type: left Virtual Network: Auto Configured

Interface Type: right Virtual Network: Auto Configured

▼ **Advanced Options**

Routing Policy

▼ **Route Aggregate**

Interface Type: right Route Aggregate: left-to-right

Cancel Save

s018720

### Create a Virtual Network and Network Policy

Create a left and right virtual network with the subnets 1.1.1/24 and 2.2.2/24, respectively. Create a network policy to apply a service chain between the left VN and the right VN. See the following example.

**Create Policy**

Policy Name: service-chain-policy

Policy Rules

Action	Protocol	Source	Ports	Direction	Destination	Ports	Log	Services	Mirror
PASS	ANY	left	ANY	right	right	ANY	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Service Instance: si-aggregate

+ Add Rule

Cancel Save

s018721

Attach the network policy to create the service chain between the left and right VNs. See the following example.

#### Validate the Route Aggregate Object in the API Server

Validate the route aggregate object in the API server configuration database. Verify the routing instance reference and the service instance reference for the aggregate object. The **aggregate\_route\_nexthop** field in the route aggregate object is initialized by the schema transformer to the service chain address. See the following example.

```
{
  - route-aggregate: {
    - fq_name: {
      "default-domain",
      "admin",
      "left-to-right"
    },
    uuid: "872b1fbd-b36c-4165-8723-7e10806d7716",
    parent_uuid: "6861d89d-a02f-4215-b329-1864084c8a75",
    aggregate_route_nexthop: "1.1.1.3",
    - routing_instance_refs: {
      - {
        - to: {
          "default-domain",
          "admin",
          "right",
          "service-ace7ae00-56e3-42d1-96ec-7fe7708d97f-default-domain_admin_si-aggregate"
        },
        href: "http://nodes27.englab.juniper.net:8082/routing-instance/d291a95a-1a5a-4fce-94c8-4abd0968d992",
        attr: null,
        uuid: "d291a95a-1a5a-4fce-94c8-4abd0968d992"
      }
    },
    parent_href: "http://nodes27.englab.juniper.net:8082/project/6861d89d-a02f-4215-b329-1864084c8a75",
    parent_type: "project",
    perms2: {(-)},
    href: "http://nodes27.englab.juniper.net:8082/route-aggregate/872b1fbd-b36c-4165-8723-7e10806d7716",
    + id_perms: {(-)},
    - aggregate_route_entries: {
      - route: {
        "1.1.1.0/24"
      }
    },
    display_name: "left-to-right",
    - service_instance_refs: {
      - {
        - to: {
          "default-domain",
          "admin",
          "si-aggregate"
        },
        href: "http://nodes27.englab.juniper.net:8082/service-instance/62accf30-8cc8-4148-b7b8-975573b0d950",
        attr: {
          interface_type: "right"
        },
        uuid: "62accf30-8cc8-4148-b7b8-975573b0d950"
      }
    },
    name: "left-to-right"
  }
}
```

### Validate the Route Aggregate Object in the Control Node

Validate the instance configurations of the route aggregate by checking the control node introspect for the service instance internal routing instance. For example:

[http://<control-node>:8083/Snh\\_ShowBgpInstanceConfigReq?search\\_string=default-domain:admin:right:service-ace7ae00-56e3-42d1-96ec-7fe77088d97f-default-domain\\_admin\\_si-aggregate](http://<control-node>:8083/Snh_ShowBgpInstanceConfigReq?search_string=default-domain:admin:right:service-ace7ae00-56e3-42d1-96ec-7fe77088d97f-default-domain_admin_si-aggregate)

See the following example.

service_chain_infos					static_routes aggregate_routes	
family	routing_instance	chain_address	prefixes	service_instance	static_routes	aggregate_routes
inet	default-domain:admin:left:left	1.1.1.3	prefixes 1.1.1.0/24	default-domain:admin:si-aggregate	prefix 1.1.1.0/24	next-hop 1.1.1.3

To check the state of the route aggregate object on the control node, point your browser to:

[http://<control-node>:8083/Snh\\_ShowRouteAggregateReq](http://<control-node>:8083/Snh_ShowRouteAggregateReq)

See the following example.

service_chain_infos					static_routes aggregate_routes	
family	routing_instance	chain_address	prefixes	service_instance	static_routes	aggregate_routes
inet	default-domain:admin:left:left	1.1.1.3	prefixes 1.1.1.0/24	default-domain:admin:si-aggregate	prefix 1.1.1.0/24	next-hop 1.1.1.3

You can also check the route table for the aggregate route in the right VN BGP able. For example:

[http://<control-node>:8083/Snh\\_ShowRouteReq?x=default-domain:admin:right:right.inet.0](http://<control-node>:8083/Snh_ShowRouteReq?x=default-domain:admin:right:right.inet.0)

See the following example.

routes									
prefix	last_modified	paths							
1.1.1.0/24	2016-Feb-18 05:00:29.215876	paths							
protocol	last_modified	local_preference	local_as	peer_as	peer_router_id	source_as	path	next_hop	label
Aggregate	2016-Feb-18 05:00:29.215876	100	0	0	-	-	-	10.204.216.23	22

## Routing Policy

Contrail uses routing policy infrastructure to manipulate the route and path attribute dynamically. Contrail also supports attaching the import routing policy on the service instances.


The routing policy contains list terms. A term can be a terminal rule, meaning that upon a match on the specified term, no further terms are evaluated and the route is dropped or accepted, based on the action in that term.

If the term is not a terminal rule, subsequent terms are evaluated for the given route.

The list terms are structured as in the following example.

```
Policy {
  Term-1
  Term-2
}
```

The matches and actions of the policy term lists operate similarly to the Junos language match and actions operations. A visual representation is the following.



Each term is represented as in the following:

```
from {
  match-condition-1
  match-condition-2
  ..
}
then {
  action
  update-action-1
  update-action-2
  ..
}
```

The term should not contain an **any** match condition, for example, an empty **from** should not be present.

If an **any** match condition is present, all routes are considered as matching the term.

However, the **then** condition can be empty or the action can be unspecified.

### Applying Routing Policy

---

The routing policy evaluation has the following key points:

- If the term of a routing policy consists of multiple match conditions, a route must satisfy all match conditions to apply the action specified in the term.
- If a term in the policy does not specify a match condition, all routes are evaluated against the match.
- If a match occurs but the policy does not specify an accept, reject, or next term action, one of the following occurs:
  - The next term, if present, is evaluated.
  - If no other terms are present, the next policy is evaluated.
  - If no other policies are present, the route is accepted. The default routing policy action is "accept".
- If a match does not occur with a term in a policy, and subsequent terms in the same policy exist, the next term is evaluated.
- If a match does not occur with any terms in a policy, and subsequent policies exist, the next policy is evaluated.
- If a match does not occur by the end of a policy or all policies, the route is accepted.

A routing policy can consist of multiple terms. Each term consists of match conditions and actions to apply to matching routes.

Each route is evaluated against the policy as follows:

1. The route is evaluated against the first term. If it matches, the specified action is taken. If the action is to accept or reject the route, that action is taken and the evaluation of the route ends. If the next term action is specified or if no action is specified, or if the route does not match, the evaluation continues as described above to subsequent terms.
2. Upon hitting the last non-terminal term of the given routing policy, the route is evaluated against the next policy, if present, in the same manner as described in step 1.

### ***Match Condition: From***

The match condition **from** contains a list of match conditions to be satisfied for applying the action specified in the term. It is possible that the term doesn't have any match condition. This indicates that all routes match this term and action is applied according to the action specified in the term.



The following table describes the match conditions supported by Contrail.

Match Condition	User Input	Description
Prefix	List of prefixes to match	<p>Each prefix in the list is represented as prefix and match type, where the prefix match type can be:</p> <ul style="list-style-type: none"> <li>• <b>exact</b></li> <li>• <b>orlonger</b></li> <li>• <b>longer</b></li> </ul> <p>Example: 1.1.0.0/16 <b>orlonger</b></p> <p>A route matches this condition if its prefix matches any of the prefixes in the list.</p>
Community	Community string to match	<p>Represented as either a well-known community string with <b>no export</b> or <b>no reoriginate</b>, or a string representation of a community (64512:11).</p>
Protocol	Array of path source or path protocol to match	<p>BGP   XMPP   StaticRoute   ServiceChain   Aggregate. A path is considered as matching this condition if the path protocol is one of protocols in the list.</p>

#### ***Routing Policy Action and Update Action***

The policy action contains two parts, action and update action.

The following table describes **action** as supported by Contrail.

Action	Terminal?	Description
Reject	Yes	Reject the route that matches this term. No more terms are evaluated after hitting this term.
Accept	Yes	Accept the route that matches this term. No more terms are evaluated after hitting this term. The route is updated using the update specified in the policy action.
Next Term	No	This is the default action taken upon matching the policy term. The route is updated according to the update specified in the policy action. Next terms present in the routing policy are processed on the route. If there are no more terms in the policy, the next routing policy is processed, if present.

The update action section specifies the route modification to be performed on the matching route.

The following table describes **update action** as supported by Contrail.


Update Action	User Input	Description
community	List of community	As part of the policy update, the following actions can be taken for community: <ul style="list-style-type: none"> <li>• Add a list of community to the existing community.</li> <li>• Set a list of community.</li> <li>• Remove a list of community (if present) from the existing community.</li> </ul>
MED	Update the MED of the BgpPath	Unsigned integer representing the MED
local-pref	Update the local-pref of the BgpPath	Unsigned integer representing local-pref

### Routing Policy Configuration

Routing policy is configured on the service instance. Multiple routing policies can be attached to a single service instance interface.

When the policy is applied on the left interface, the policy is evaluated for all the routes that are reoriginated in the left VN for routes belonging to the right VN. Similarly, the routing policy attached to the right interface influences the route reorigination in the right VN, for routes belonging to the left VN.

The following figure illustrates a routing policy configuration.



The policy sequence number specified in the routing policy link data determines the order in which the routing policy is evaluated. The routing policy link data on the service instance

also specifies whether the policy needs to be applied to the left service interface, to the right service interface, or to both interfaces.

It is possible to attach the same routing policy to both the left and right interfaces for a service instance, in a different order of policy evaluation. Consequently, the routing policy link data contains the sequence number for policy evaluation separately for the left and right interfaces.

The schema transformer links the routing policy object to the internal routing instance created for the service instance. The transformer also copies the routing policy link data to ensure the same policy order.

### Configuring and Troubleshooting Routing Policy

This section shows how to create a routing policy for service chains and how to validate the policy.

- [Create Routing Policy on page 217](#)
- [Configure Service Instance on page 218](#)
- [Configure the Network Policy for the Service Chain on page 218](#)

#### Create Routing Policy

First, create the routing policy. See the following example.

**Create Routing Policy**

**Name**  
failover

**Term(s)**  
from: { prefix 2.2.2.0/24 orlonger } then: { local-preference 200 }

**From**  
prefix 2.2.2.0/24 orlonger

**Then**  
local-preference 200

Cancel Save

s018729

### Configure Service Instance

Create a service instance and attach the routing policy to both the left and right interfaces. The order of the policy is calculated by the UI, based on the order of the policy specified in the list.

### Configure the Network Policy for the Service Chain

At **Edit Policy**, create a policy for the service chain, see the following example.

### Using a VNC Script to Create Routing Policy

The following example shows use of a VNC API script to create a routing policy.

```
from vnc_api.vnc_api import *
vnc_lib = VncApi("admin", "<password>", "admin")
project=vnc_lib.project_read(fq_name=["default-domain", "admin"])
routing_policy=RoutingPolicy(name="vnc_3", parent_obj=project)
policy_term=PolicyTermType()
policy_statement=PolicyStatementType()

match_condition=TermMatchConditionType(protocol=["bgp"], community="22:33")
prefix_match=PrefixMatchType(prefix="1.1.1.0/24", prefix_type="orlonger")
match_condition.set_prefix([prefix_match])

term_action=TermActionListType(action="accept")
```

```

action_update=ActionUpdateType(local_pref=101, med=10)
add_community=ActionCommunityType()
comm_list=CommunityListType(["11:22"])
add_community.set_add(comm_list)
action_update.set_community(add_community)
term_action.set_update(action_update)

policy_term.set_term_action_list(term_action)
policy_term.set_term_match_condition(match_condition)

policy_statement.add_term(policy_term)
routing_policy.set_routing_policy_entries(policy_statement)
vnc_lib.routing_policy_create(routing_policy)

```

### Verify Routing Policy in API Server

You can verify the service instance references and the routing instance references for the routing policy by looking in the API server configuration database. See the following example.

```

-----
- routing_policy_entries: {
  - term: {
    - {
      - term_match_condition: {
        - prefix: {
          - {
            prefix_type: "orlonger",
            prefix: "2.2.2.0/24"
          }
        }
      },
      - term_action_list: {
        action: "accept",
        - update: {
          local_pref: 200
        }
      }
    }
  }
},
+ id_perms: {...},
- routing_instance_refs: [
  - {
    - to: [
      "default-domain",
      "admin",
      "right",
      "service-ace7ae00-56e3-42d1-96ec-7fe77088d97f-default-domain_admin_ha-chain"
    ],
    href: "http://nodea27.englab.juniper.net:8082/routing-instance/32b7eed4-57ce-4c44-bbb0-513f78db6068",
    - attr: {
      sequence: "1"
    },
    uuid: "32b7eed4-57ce-4c44-bbb0-513f78db6068"
  },
  - {
    - to: [
      "default-domain",
      "admin",
      "left",
      "service-ace7ae00-56e3-42d1-96ec-7fe77088d97f-default-domain_admin_ha-chain"
    ],
    href: "http://nodea27.englab.juniper.net:8082/routing-instance/6ad868d1-a412-4765-b8c4-f93ec5d9f4b2",
    - attr: {
      sequence: "1"
    },
    uuid: "6ad868d1-a412-4765-b8c4-f93ec5d9f4b2"
  }
],
- service_instance_refs: [
  - {
    - to: [
      "default-domain",
      "admin",
      "ha-chain"
    ],
    href: "http://nodea27.englab.juniper.net:8082/service-instance/983bb90b-b3f4-4d6c-be54-33a474eee7de",
    - attr: {
      left_sequence: "1",
      right_sequence: "1"
    },
    uuid: "983bb90b-b3f4-4d6c-be54-33a474eee7de"
  }
],
name: "failover"
-----
s018732

```

## Verify Routing Policy in the Control Node

You can verify the routing policy in the control node.

Point your browser to:

[http://<control-node>:8083/Snh\\_ShowRoutingPolicyReq?search\\_string=failover](http://<control-node>:8083/Snh_ShowRoutingPolicyReq?search_string=failover)

See the following example.

routing_policies						
name	generation	ref_count	terms			deleted
default-domain:admin:failover	0	2	<div><div>terminal</div><div>matches</div><div>actions</div></div> <div><div>true</div><div>prefix [2.2.2.0/24 orlonger]</div><div>accept local-pref 200</div></div>			false
default-domain:default-project:default-routing-policy	0	0	terms			false

## Verify Routing Policy Configuration in the Control Node

You can verify the routing policy configuration in the control node.

Point your browser to:

[http://<control-node>:8083/Snh\\_ShowBgpRoutingPolicyConfigReq?search\\_string=failover](http://<control-node>:8083/Snh_ShowBgpRoutingPolicyConfigReq?search_string=failover)

See the following example.

ShowBgpRoutingPolicyConfigResp			
routing_policies			
name		terms	
default-domain:admin:failover		terms	
		match	action
		from { prefix 2.2.2.0/24 orlonger }	then { local-preference 200 accept }

5018733

## Verify Routing Policy Configuration on the Routing Instance

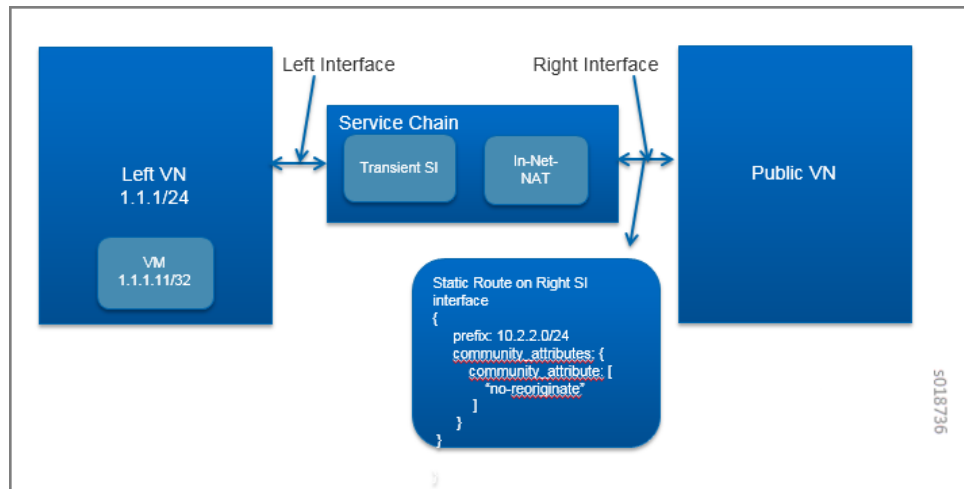
You can verify the routing policy configuration on the internal routing instance.

Point your browser to:

[http://<control-node>:8083/Snh\\_ShowBgpInstanceConfigReq?search\\_string=<name-of-internal-vrf>](http://<control-node>:8083/Snh_ShowBgpInstanceConfigReq?search_string=<name-of-internal-vrf>)



When the control node is reoriginating the route, it skips the routes that are tagged with the BGP community.



### Configuring and Troubleshooting Reorigination Control

The community attribute on the static routes for the interface static route of the service instance is specified during creation of the service instance. See the following example.



Use the following example to verify that the service instance configuration object in the API server has the correct community set for the static route. See the following example.

```
{
  - service-instance: {
    + virtual_machine_back_refs: [...],
    + fq_name: [...],
    uuid: "a6e1e71f-f828-43de-a493-b193bdb73ded",
    parent_type: "project",
    parent_uuid: "634f90d9-da62-4c2f-a238-7cc1c1a055a5",
    parent_bref: "http://nodeq2:8082/project/634f90d9-da62-4c2f-a2",
    - service_instance_properties: {
      right_virtual_network: "default-domain:admin:twig",
      - interface_list: [
        - {
          virtual_network: "default-domain:admin:fifo"
        },
        - {
          virtual_network: "default-domain:admin:twig",
          - static_routes: {
            - route: [
              - {
                prefix: "10.2.2.0/24",
                next_hop: null,
                - community_attributes: {
                  - community_attribute: [
                    "no-reoriginate"
                  ],
                },
                next_hop_type: null
              }
            ]
          }
        }
      ],
      left_virtual_network: "default-domain:admin:fifo",
      - scale_out: {
        max_instances: 1
      }
    },
    ...
  },
  ...
}
```

s018738

## Service Instance Health Check

In Contrail Release 3.0 and greater, a service instance health check can be used to determine the liveliness of a service provided by a virtual machine (VM).

- [Health Check Overview on page 223](#)
- [Health Check Object Configuration on page 224](#)
- [Creating a Health Check with the Contrail User Interface on page 225](#)
- [Using the Health Check on page 226](#)
- [Health Check Process on page 226](#)

## Health Check Overview

The service instance health check is used to determine the liveliness of a service provided by a VM, checking whether the service is operationally up or down. The vRouter agent uses ping and an HTTP URL to the link-local address to check the liveliness of the interface.

If the health check determines that a service is no longer operational, it removes the routes for the VM, thereby disabling packet forwarding to the VM.

The service instance health check is used with service template Version 2.

## Health Check Object Configuration

Table 26 on page 224 shows the configurable properties of the health check object.

Table 26: Health Check Configurable Parameters

Field	Description
- enabled	Indicates that health check is enabled. The default is <b>False</b> .
- health-check-type	Indicates the health check type: <b>link-local</b> or <b>end-to-end</b> . The default is <b>link-local</b> .
- monitor-type	The protocol type to be used: <b>PING</b> or <b>HTTP</b> .
- delay	The delay, in seconds, to repeat the health check.
- timeout	The number of seconds to wait for a response.
- max-retries	The number of retries to attempt before declaring an instance health down.
- http-method	When the monitor protocol is HTTP, the type of HTTP method used, such as GET, PUT, POST, and so on.
- url-path	When the monitor protocol is HTTP, the URL to be used. For all other cases, such as ICMP, the destination IP address.
- expected-codes	When the monitor protocol is HTTP, the expected return code for HTTP operations.

### Health Check Modes

The following modes are supported for the service instance health check:

- **link-local**—A local check for the service VM on the vRouter where the VM is running. In this case, the source IP of the packet is the service chain IP.
- **end-to-end**—A remote address or URL is provided for a service health check through a chain of services. The destination of the health check probe is allowed to be outside the service instance. However, the health check probe must be reachable through the interface of the service instance where the health check is attached. The end-to-end health check probe is transmitted all the way to the actual destination outside the service instance. The response to the health check probe is received and processed by the service health check to evaluate the status.

Restrictions include:

- This check is applicable for a chain where the services are not scaled out.
- When this mode is configured, a new health check IP is allocated and used as the source IP of the packet.

- The health check IP is allocated per **virtual-machine-interface** of the service VM where the health check is attached.
- The agent relies on the **service-health-check-ip** flag to use as the source IP.

## Creating a Health Check with the Contrail User Interface

To create a health check with the Contrail Web UI:

1. Navigate to **Configure > Services > Health Check Service**, and click to open the **Create** screen. See [Figure 71 on page 225](#).

Figure 71: Create Health Check Screen

The screenshot shows a 'Create' dialog box for a 'Health Check Service'. It has two tabs: 'Health Check Service' (selected) and 'Permissions'. The form contains the following fields:

- Name:** A text input field containing 'ext\_hc\_service'.
- Protocol:** A dropdown menu with 'PING' selected.
- Monitor Target:** A text input field containing '8.8.8.8'.
- Delay (secs):** A text input field containing '3'.
- Timeout (secs):** A text input field containing '5'.
- Retries:** A text input field containing '2'.
- Health Check Type:** A dropdown menu with 'End-To-End' selected.

At the bottom right, there are 'Cancel' and 'Save' buttons. A small vertical text '5018766' is visible on the right edge of the dialog box.

2. Complete the fields to define the permissions for the health check, see [Table 27 on page 225](#).

Table 27: Create Health Check Fields

Field	Description
Name	Enter a name for the health check service you are creating.
Protocol	Select from the list the protocol to use for the health check, PING, HTTP, and so on.
Monitor Target	Select from the list the address of the target to be monitored by the health check.
Delay (secs)	The delay, in seconds, to repeat the health check.
Timeout (secs)	The number of seconds to wait for a response.
Retries	The number of retries to attempt before declaring an instance health down.

Table 27: Create Health Check Fields (*continued*)

Field	Description
Health Check Type	Select from the list the type of health check—link-local or end-to-end.

## Using the Health Check

A REST API can be used to create a health check object and define its associated properties, then a link is added to the VM interface.

The health check object can be linked to multiple VM interfaces. Additionally, a VM interface can be associated with multiple health check objects. The following is an example:

```

HealthCheckObject 1 ----- VirtualMachineInterface 1 -----
HealthCheckObject 2
|
|
VirtualMachineInterface 2

```

## Health Check Process

The Contrail vRouter agent is responsible for providing the health check service. The agent spawns a Python script to monitor the status of a service hosted on a VM on the same compute node, and the script updates the status to the vRouter agent.

The vRouter agent acts on the status provided by the script to withdraw or restore the exported interface routes. It is also responsible for providing a link-local metadata IP for allowing the script to communicate with the destination IP from the underlay network, using appropriate NAT translations. In a running system, this information is displayed in the vRouter agent introspect at:

[http://<compute-node-ip>:8085/Snh\\_HealthCheckSandeshReq?uuid=](http://<compute-node-ip>:8085/Snh_HealthCheckSandeshReq?uuid=)



**NOTE:** Running health check creates flow entries to perform translation from underlay to overlay. Consequently, in a heavily loaded environment with a full flow table, it is possible to observe false failures.

# Examples: Configuring Service Chaining

- [Example: Creating an In-Network or In-Network-NAT Service Chain on page 227](#)
- [Example: Creating a Transparent Service Chain on page 235](#)
- [Example: Creating a Service Chain With the CLI on page 239](#)

## Example: Creating an In-Network or In-Network-NAT Service Chain

This section provides an example of creating an **in-network** service chain and an **in-network-nat** service chain using the Juniper Networks Contrail user interface. This service chain example also shows scaling of service instances.


- [Creating an In-Network or In-Network-NAT Service Chain on page 227](#)

### Creating an In-Network or In-Network-NAT Service Chain

To create an **in-network** or **in-network-nat** service chain:

1. Create a left and a right virtual network. Select **Configure > Networking > Networks** and create **left\_vn** and **right\_vn**; see [Figure 72 on page 227](#).

Figure 72: Create Networks



2. Configure a service template for an in-network service template for NAT. Navigate to **Configure > Services > Service Templates** and click the **Create** button on **Service Templates**. The **Add Service Template** window appears; see [Figure 73 on page 228](#).

Figure 73: Add Service Template

**Add Service Template**

Name: nat-template

Service Mode: In-Network

Image Name: nat-service

Interface Types	Shared IP	Static Routes	+
Management	<input type="checkbox"/>	<input type="checkbox"/>	+ -
Left	<input checked="" type="checkbox"/>	<input type="checkbox"/>	+ -
Right	<input type="checkbox"/>	<input type="checkbox"/>	+ -

▼ Advanced options

Service Scaling: ☒

Instance Flavor: m1.medium(RAM:4096, CPU cores:2, Disk:...) ▼

Cancel Save

s041902

Table 28: Add Service Template Fields

Field	Description
<b>Name</b>	Enter a name for the service template.
<b>Service Mode</b>	Select the service mode: <b>In-Network</b> (for firewall service), <b>In-Network-NAT</b> (for NAT service), or <b>Transparent</b> .
<b>Service Scaling</b>	If you will be using multiple virtual machines for a single service instance to scale out the service, select the <b>Service Scaling</b> check box. When scaling is selected, you can choose to use the same IP address for a particular interface on each virtual machine interface or to allocate new addresses for each virtual machine. For a NAT service, the left (inner) interface should have the same IP address, and the right (outer) interface should have a different IP address.
<b>Image Name</b>	Select from a list of available images the image for the service.  <b>NOTE:</b> Only images that have been tagged as public in Glance will appear in the drop-down list.

Table 28: Add Service Template Fields (*continued*)


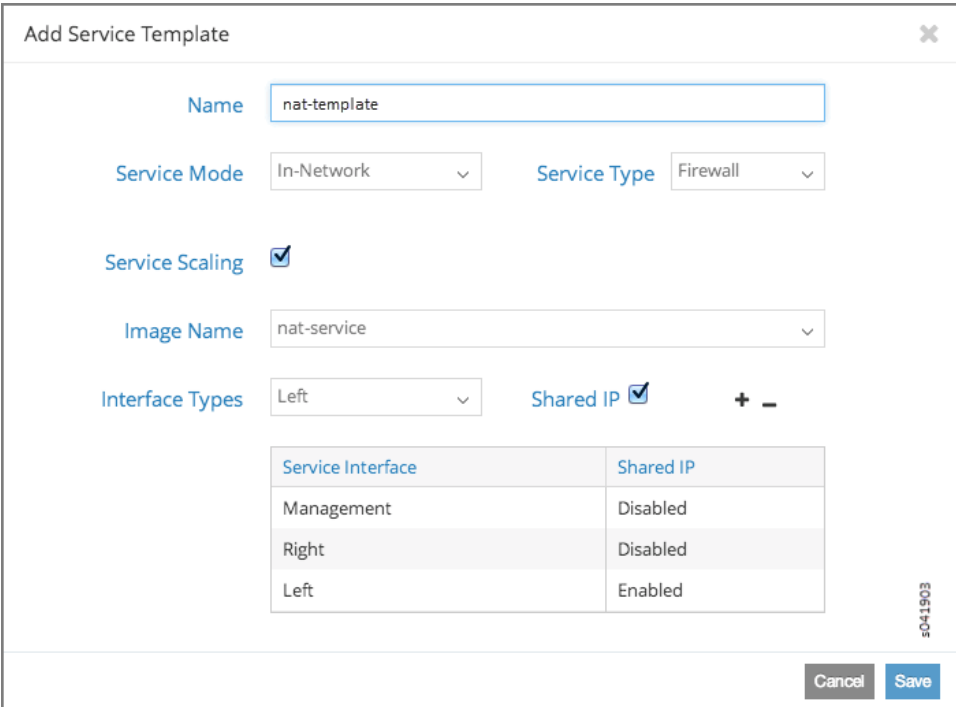
Field	Description
<b>Interface Types</b>	<p>Select the interface type or types for this service:</p> <ul style="list-style-type: none"> <li>For firewall or NAT services, both <b>Left Interface</b> and <b>Right Interface</b> are required.</li> <li>For an analyzer service, only a <b>Left Interface</b> is required.</li> <li>For Juniper Networks virtual images, <b>Management Interface</b> is also required, in addition to any left or right requirement.</li> </ul>
	<p>3. On <b>Add Service Template</b>, complete the following for the in-network service template:</p> <ul style="list-style-type: none"> <li><b>Name:</b> nat-template</li> <li><b>Service Mode:</b> In-Network</li> <li><b>Service Scaling:</b> Select from Advanced</li> <li><b>Image Name:</b> nat-service</li> <li><b>Interface Types:</b> Select Left Interface and Right Interface. For Juniper Networks virtual images, select Management Interface as the first interface.</li> <li>The Left Interface will be automatically marked for sharing the same IP address</li> </ul> <p>4. If multiple instances are to be launched for a particular service instance, select the <b>Service Scaling</b> check box, which enables the <b>Shared IP</b> feature. <a href="#">Figure 74 on page 230</a> shows the <b>Left</b> interface selected, with the <b>Shared IP</b> check box selected, so the left interface will share the IP address.</p>
	<p> <b>NOTE:</b> The <b>Shared IP</b> for <b>Service Scaling</b> is an internal infrastructure feature used only for service scaling, it cannot be used for other features.</p>

Figure 74: Add Service Template Shared IP



**Add Service Template**

Name: nat-template

Service Mode: In-Network Service Type: Firewall

Service Scaling: ☒

Image Name: nat-service

Interface Types: Left

Shared IP: ☒ + -

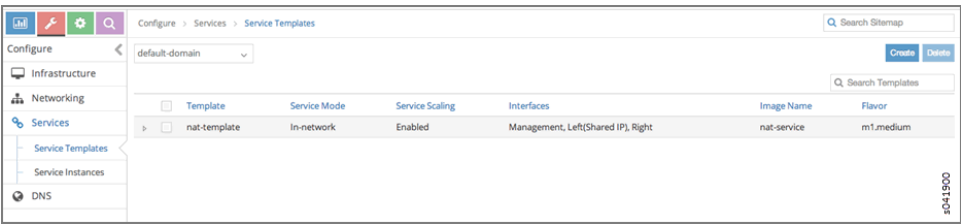
Service Interface	Shared IP
Management	Disabled
Right	Disabled
Left	Enabled

Cancel Save

- Click **Save**.

The service template is created and appears on the **Service Templates** screen, see [Figure 75 on page 230](#).

Figure 75: Service Templates



Configure > Services > Service Templates

default-domain

Create Delete

Search Sitemap

Search Templates

Template	Service Mode	Service Scaling	Interfaces	Image Name	Flavor
nat-template	In-network	Enabled	Management, Left(Shared IP), Right	nat-service	m1.medium

- Create the service instance. Navigate to **Configure > Services > Service Instances**, and click **Create**, then select the template to use and select the corresponding left, right, or management networks; see [Figure 76 on page 231](#).



Figure 76: Create Service Instances

Table 29: Create Service Instances Fields

Field	Description
<b>Instance Name</b>	Enter a name for the service instance.
<b>Services Template</b>	Select from a list of available service templates the service template to use for this instance.
<b>Number of Instances</b>	If scaling is enabled, enter a value in the <b>Number of Instances</b> field to define the number of instances of service virtual machines to launch.
<b>Interface List and Virtual Networks</b>	An ordered list of interfaces as defined in the Service Template. If you are using the <b>Management Interface</b> , select <b>Auto Configured</b> . The software will use an internally-created virtual network. For <b>Left Interface</b> , select <b>left_vn</b> and for <b>Right Interface</b> , select <b>right_vn</b> .

- If static routes are enabled for specific interfaces, open the **Static Routes** field below each enabled interface and enter the static route address details; see [Figure 77 on page 232](#).

Figure 77: Create Service Instances

**Create Service Instances**

Instance Name:

Services Template: nat-ecmp-template - [in-network (management, left, right)]

Number of instances: 1

Interface 1: Management Auto Configured

Interface 2: Left vn10 (admin)

Static Routes

Prefix	Next hop	
10.204.80.0/28	Interface 2	+ -

Interface 3: Right vn10 (admin)

Cancel Save

- The console for the service instances can be viewed. At **Configure > Services > Service Instances**, click the arrow next to the name of the service instance to reveal the details panel for that instance, then click **View Console** to see the console details; see [Figure 78 on page 232](#) and [Figure 79 on page 233](#).

Figure 78: Service Instance Details

Virtual Machine	Status	Power State	Networks
fw-instance_1	ACTIVE	RUNNING	svc-vn-mgmt:250.250.1.252 svc-vn-left:250.250.2.253 svc-vn-right:250.250.3.253

[View Console](#)

Figure 79: Service Instance Console

0.204.216.36:5999/vnc\_auto.html?token=9eada783-24e7-4808-9325-4ad257bf3762

Connected (unencrypted) to: QEMU (instance-0000000b)

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet	250.250.1.253/24	
ge-0/0/0	up	up			
ip-0/0/0	up	up			
lsq-0/0/0	up	up			
lt-0/0/0	up	up			
mt-0/0/0	up	up			
sp-0/0/0	up	up			
sp-0/0/0.0	up	up	inet	10.0.0.1	--> 10.0.0.16
sp-0/0/0.16383	up	up	inet	10.0.0.6	--> 0/0
				128.0.0.1	--> 128.0.1.16
				128.0.0.6	--> 0/0
ge-0/0/1	up	up			
ge-0/0/1.0	up	up	inet	1.1.1.253/24	
ge-0/0/2	up	up			
ge-0/0/2.0	up	up	inet	2.2.2.253/24	
dsc	up	up			
gre	up	up			
ipip	up	up			
lo0	up	up			
lo0.16384	up	up	inet	127.0.0.1	--> 0/0
lo0.16385	up	up	inet	10.0.0.1	--> 0/0
---(more)---					

5041919

9. Configure the network policy. Navigate to **Configure > Networking > Policies**.
  - Name the policy and associate it with the networks created earlier: **left\_vn** and **right\_vn**.
  - Set source network as **left\_vn** and destination network as **right\_vn**.
  - Select **Apply Service** and select the service (**nat-ecmp**).

Figure 80: Create Policy

Create Policy

Policy Name  
fw-policy

Policy Rules

Action	Protocol	Source Network	Source Ports	Direction	Destination Network	Destination Ports	Apply Service	Mirror to	
PAS	ANY	left_vn	Source	<>	right_vn	Destination	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
fw-instance x									

Cancel Save

5041870

10. Associate the policy with both the **left\_vn** and the **right\_vn**. Navigate to **Configure > Networking > Network**.
  - On the right side of **left\_vn**, click the gear icon to enable **Edit Network**.
  - In the **Edit Network** dialog box for **left\_vn**, select **nat-policy** in the **Network Policy(s)** field.
  - Repeat the same process for the **right\_vn**.

Figure 81: Edit Network

Network Name:

Network Policy(s):

Address Management:  IP Block:  Gateway:

IPAM	IP Block	Gateway
default-domain:default-project:default-network-ipam	1.1.1.0/24	1.1.1.254

Route Targets

Floating IP Pools

Host Routes

Advanced Options

Cancel Save

11. Launch virtual machines (from OpenStack) and test the traffic through the service chain by doing the following:
  - a. Navigate to **Configure > Networking > Policies**.
  - b. Launch **left\_vm** in virtual network **left\_vn**.
  - c. Launch **right\_vm** in virtual network **right\_vn**.
  - d. Ping from **left\_vm** to **right\_vm** IP address (**2.2.2.252** in [Figure 82 on page 234](#)).
  - e. A **TCPDUMP** on the **right\_vm** should show that packets are NAT-enabled and have the source IP set to **2.2.2.253**.

Figure 82: Launch Instances

Instances

Launch Instance Terminate Instances

Instance Name	IP Address	Size	Keypair	Status	Task	Power State	Actions
nat-instance_1	250.250.1.253	m1.medium   4GB RAM   2 VCPU   40GB Disk	-	Active	None	Running	Create Snapshot More
right_vm	2.2.2.252	m1.tiny   512MB RAM   1 VCPU   0 Disk	-	Active	None	Running	Create Snapshot More
left_vm	1.1.1.252	m1.tiny   512MB RAM   1 VCPU   0 Disk	-	Active	None	Running	Create Snapshot More

Displaying 3 items

- Related Documentation**
- [Service Chaining on page 191](#)
  - [Example: Creating a Transparent Service Chain on page 235](#)
  - [ECMP Load Balancing in the Service Chain on page 196](#)

## Example: Creating a Transparent Service Chain

This section provides an example of creating a transparent mode service chain using the Juniper Networks Contrail user interface. Also called bridge mode, transparent mode is used for services that do not modify the packet, such as Layer 2 firewall, Intrusion Detection and Prevention (IDP), and so on. The following service chain example also shows scaling of service instances.


- [Creating a Transparent Mode Service Chain on page 235](#)

### Creating a Transparent Mode Service Chain

To create a transparent mode service chain:

1. Create a left and a right virtual network. Select **Configure > Networking > Networks** and create **left\_vn** and **right\_vn**; see [Figure 83 on page 235](#).

**Figure 83: Create Networks**



2. Configure a service template for a transparent mode. Navigate to **Configure > Services > Service Templates** and click the **Create** button on **Service Templates**. The **Add Service Template** window appears; see [Figure 84 on page 236](#).

Figure 84: Add Service Template

**Add Service Template**

Name:

Service Mode:

Image Name:

Interface Types	Shared IP	Static Routes	+
Management	<input type="checkbox"/>	<input type="checkbox"/>	+ -
Left	<input checked="" type="checkbox"/>	<input type="checkbox"/>	+ -
Right	<input checked="" type="checkbox"/>	<input type="checkbox"/>	+ -

**Advanced options**

Service Scaling: ☒

Instance Flavor:

Cancel Save

Table 30: Add Service Template Fields

Field	Description
Name	Enter a name for the service template.
Service Mode	Select the service mode: <b>In-Network</b> or <b>Transparent</b> .
Service Scaling	If you will be using multiple virtual machines for a single service instance to scale out the service, select the <b>Service Scaling</b> check box. When scaling is selected, you can choose to use the same IP address for a particular interface on each virtual machine interface or to allocate new addresses for each virtual machine. For a NAT service, the left (inner) interface should have the same IP address, and the right (outer) interface should have a different IP address.
Image Name	Select from a list of available images the image for the service.
Interface Types	Select the interface type or types for this service: <ul style="list-style-type: none"> <li>For firewall or NAT services, both <b>Left Interface</b> and <b>Right Interface</b> are required.</li> <li>For an analyzer service, only <b>Left Interface</b> is required.</li> <li>For Juniper Networks virtual images, <b>Management Interface</b> is also required, in addition to any left or right requirement.</li> </ul>

3. On **Add Service Template**, complete the following for the transparent mode service template:

- **Name:** firewall-template
- **Service Mode:** Transparent
- **Service Scaling:** Select this.
- **Image Name:** vsrx-bridge
- **Interface Types:** Select Left Interface, Right Interface, and Management Interface.

If multiple instances are to be launched for a particular service instance, select the **Service Scaling** check box, which enables the **Shared IP** feature.

4. Click **Save**.

5. Create the service instance. Navigate to **Configure > Services > Service Instances**, and click **Create**, then select the template to use and select the corresponding left, right, or management networks; see [Figure 85 on page 237](#).

Figure 85: Create Service Instances

Table 31: Create Service Instances Fields


Field	Description
Instance Name	Enter a name for the service instance.
Services Template	Select from a list of available service templates the service template to use for this instance.
Left Network	Select from a list of available virtual networks the network to use for the left interface. For transparent mode, select <b>Auto Configured</b> .
Right Network	Select from a list of available virtual networks the network to use for the right interface. For transparent mode, select <b>Auto Configured</b>

Table 31: Create Service Instances Fields (*continued*)

**Management Network** If you are using the **Management Interface**, select **Auto Configured**. The software will use an internally-created virtual network. For transparent mode, select **Auto Configured**


6. If scaling is enabled, enter a value in the **Number of Instances** field to define the number of instances of service virtual machines to launch; see [Figure 86 on page 238](#).

Figure 86: Service Instance Details



7. Next, configure the network policy. Navigate to **Configure > Networking > Policies**.
  - Name the policy **fw-policy**.
  - Set source network as **left\_vn** and destination network as **right\_vn**.
  - Check **Apply Service** and select the service (**fw-instance**).

Figure 87: Create Policy

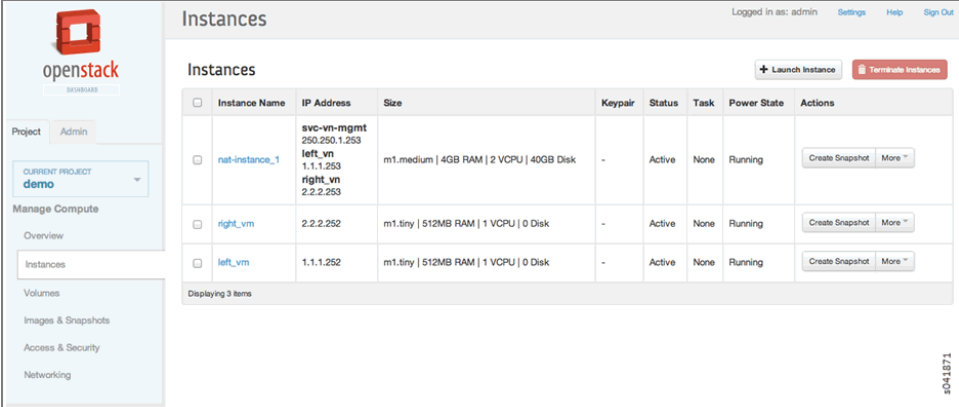


8. Next, associate it to the networks created earlier – **left\_vn** and **right\_vn**. Navigate to **Configure > Networking > Policies**.
  - On the right side of **left\_vn**, click the gear icon to enable **Edit Network**.
  - In the **Edit Network** dialog box for **left\_vn**, select **nat-policy** in the **Network Policy(s)** field.
  - Repeat the process for the **right\_vn**.
9. Next, launch virtual machines (from OpenStack) and test the traffic through the service chain by doing the following:
  - a. Navigate to **Configure > Networking > Policies**.
  - b. Launch **left\_vm** in virtual network **left\_vn**.



- c. Launch `right_vm` in virtual network `right_vn`.
- d. Ping from `left_vm` to `right_vm` IP address (2.2.2.252 in Figure 88 on page 239).
- e. A `TCPDUMP` on the `right_vm` should show that packets have the source IP set to 2.2.2.253.

Figure 88: Launch Instances



Instance Name	IP Address	Size	Keypair	Status	Task	Power State	Actions
<code>nat-instance_1</code>	250.250.1.253 <code>left_vn</code> 1.1.1.253 <code>right_vn</code> 2.2.2.253	m1.medium   4GB RAM   2 VCPU   40GB Disk	-	Active	None	Running	Create Snapshot More
<code>right_vm</code>	2.2.2.252	m1.tiny   512MB RAM   1 VCPU   0 Disk	-	Active	None	Running	Create Snapshot More
<code>left_vm</code>	1.1.1.252	m1.tiny   512MB RAM   1 VCPU   0 Disk	-	Active	None	Running	Create Snapshot More

Related Documentation

- [Service Chaining on page 191](#)

## Example: Creating a Service Chain With the CLI

This section provides syntax and examples for creating service chaining objects for Contrail Controller.

- [CLI for Creating a Service Chain on page 239](#)
- [CLI for Creating a Service Template on page 240](#)
- [CLI for Creating a Service Instance on page 240](#)
- [CLI for Creating a Service Policy on page 240](#)
- [Example: Creating a Service Chain with VSRX and In-Network or Routed Mode on page 241](#)

### CLI for Creating a Service Chain

All of the commands needed to create service chaining objects are located in `/opt/contrail/utils`.

## CLI for Creating a Service Template

The following commands are used to create a service template:

```
./service-template.py add    [--svc_type {firewall, analyzer}]  
  
                             [--image_name IMAGE_NAME]  
  
                             template_name  
  
./service-template.py del    template_name
```

## CLI for Creating a Service Instance

The following commands are used to create a service instance:

```
./service-instance.py add    [--proj_name PROJ_NAME]  
  
                             [--mgmt_vn MGMT_VN]  
  
                             [--left_vn LEFT_VN]  
  
                             [--right_vn RIGHT_VN]  
  
                             instance_name  
  
                             template_name  
  
./service-instance.py del    [--proj_name PROJ_NAME]  
  
                             instance_name  
  
                             template_name
```

---

## CLI for Creating a Service Policy

The following commands are used to create a service policy:

```
./service-policy.py add      --svc_list SVC_LIST [SVC_LIST ...]  
  
                             --vn_list VN_LIST [VN_LIST ...]  
  
                             [--proj_name PROJ_NAME]  
  
                             policy_name  
  
./service-policy.py del      [--proj_name PROJ_NAME]  
  
                             policy_name
```

---

## Example: Creating a Service Chain with VSRX and In-Network or Routed Mode

The following example creates a VSRX firewall service in a virtual network named **test**, using a project named **demo** and a template, an instance, and a policy, all named **test**.

1. Add images to Glance (OpenStack image service).

- a. Download the following images:

```
precise-server-cloudimg-amd64-disk1.img
```

```
junos-vsrx-12.1-nat.img
```

- b. Add the images to Glance, using the names **ubuntu** and **vsrx**.

```
(source /etc/contrail/openstackrc; glance add name='ubuntu' is_public=true
container_format=ovf disk_format=qcow2 <
precise-server-cloudimg-amd64-disk1.img)
```

```
(source /etc/contrail/openstackrc; glance add name='vsrx' is_public=true
container_format=ovf disk_format=qcow2 < junos-vsrx-12.1-dhcp.img)
```

2. Create a service template of type **firewall** and named **vsrx**.

```
./service-template.py add test_template --svc_type firewall --image_name vsrx
```

3. Create virtual networks.

```
VN1
```

```
VN2
```

4. Create a service template.

```
./service-template.py add --svc_scaling ecmp-template
```

5. Create a service instance.

```
./service-instance.py add --proj_name admin --left_vn VN1 --right_vn VN2
--max_instances 3 ecmp-instance ecmp-template
```

6. Create a service policy.

```
./service-policy.py add proj_name admin --svc_list ecmp-instance --vn_list VN1 VN2
ecmp-policy
```

7. Create virtual machines and attach them to virtual networks.

```
VM1 (attached to VN1)—use ubuntu image
```

```
VM2 (attached to VN2)—use ubuntu image
```

8. Launch the instances **VM1** and **VM2**.

9. Send ping traffic from **VM1** to **VM2**.

10. Send traffic from **VM1** in **VN1** to **VM2** in **VN2**.

11. You can use the Contrail Juniper Networks interface to monitor the ping traffic flows. Select **Monitor > Infrastructure > Virtual Routers** and select an individual vRouter. Click

through to view the vRouter details, where you can click the **Flows** tab to view the flows.

**Related Documentation** • [Service Chaining on page 191](#)

## PART 4

# Monitoring and Troubleshooting the Network Using Contrail Analytics

- [Understanding Contrail Analytics on page 245](#)
- [Configuring Contrail Analytics on page 267](#)
- [Using Contrail Analytics to Monitor and Troubleshoot the Network on page 279](#)



# Understanding Contrail Analytics

- [Contrail Analytics Overview on page 245](#)
- [Contrail Alerts on page 246](#)
- [Underlay Overlay Mapping in Contrail on page 249](#)

## Contrail Analytics Overview

---

Contrail is a distributed system of compute nodes, control nodes, configuration nodes, database nodes, web UI nodes, and analytics nodes.

The analytics nodes are responsible for the collection of system state information, usage statistics, and debug information from all of the software modules across all of the nodes of the system. The analytics nodes store the data gathered across the system in a database that is based on the Apache Cassandra open source distributed database management system. The database is queried by means of an SQL-like language and representational state transfer (REST) APIs.

System state information collected by the analytics nodes is aggregated across all of the nodes, and comprehensive graphical views allow the user to get up-to-date system usage information easily.

Debug information collected by the analytics nodes includes the following types:

- System log (syslog) messages—informational and debug messages generated by system software components.
- Object log messages—records of changes made to system objects such as virtual machines, virtual networks, service instances, virtual routers, BGP peers, routing instances, and the like.
- Trace messages—records of activities collected locally by software components and sent to analytics nodes only on demand.

Statistics information related to flows, CPU and memory usage, and the like is also collected by the analytics nodes and can be queried at the user interface to provide historical analytics and time-series information. The queries are performed using REST APIs.

Analytics data is written to a database in Contrail. The data expires after the default time-to-live (TTL) period of 48 hours. This default TTL time can be changed as needed by changing the value of the `database_ttl` value in the file `testbed.py`.

**Related  
Documentation**

- [Contrail Alerts on page 246](#)
- [Analytics Scalability on page 267](#)
- [High Availability for Analytics on page 268](#)
- [Ceilometer Support in a Contrail Cloud on page 271](#)
- [Underlay Overlay Mapping in Contrail on page 249](#)
- [Monitoring the System on page 280](#)
- [Debugging Processes Using the Contrail Introspect Feature on page 282](#)
- [Monitor > Infrastructure > Dashboard on page 286](#)
- [Monitor > Infrastructure > Control Nodes on page 288](#)
- [Monitor > Infrastructure > Virtual Routers on page 295](#)
- [Monitor > Infrastructure > Analytics Nodes on page 306](#)
- [Monitor > Infrastructure > Config Nodes on page 311](#)
- [Monitor > Networking on page 314](#)
- [\*Understanding Flow Sampling\*](#)
- [Query > Flows on page 322](#)
- [Query > Logs on page 329](#)
- [System Log Receiver in Contrail Analytics on page 269](#)
- [Example: Debugging Connectivity Using Monitoring for Troubleshooting on page 334](#)

---

## Contrail Alerts

Starting with Contrail 3.0 and greater, Contrail alerts are provided on a per-user visible entity (UVE) basis.

Contrail analytics raise or clear alerts using Python-coded rules that examine the contents of the UVE and the configuration of the object. Some rules are built in. Others can be added using Python *stevedore* plugins.

This topic describes Contrail alerts capabilities.

### Alert API Format

The Contrail alert analytics API provides the following:

- Read access to the alerts as part of the UVE GET APIs.
- Alert acknowledgement using POST requests.
- UVE and alert streaming using server-sent events (SSEs).



For example:

GET `http://<analytics-ip>:8081/analytics/uves/control-node/a6s40?flat`

```
{
  NodeStatus: {...},
  ControlCpuState: {...},
  UVEAlarms: {
    alarms: [
      {
        description: [
          {
            value: "0 != 2",
            rule: "BgpRouterState.num_up_bgp_peer !=
BgpRouterState.num_bgp_peer"
          }
        ],
        ack: false,
        timestamp: 1442995349253178,
        token: "eyJ0aw1lc3RhbnXAiOiAxNDQyOTk1MzQ5MTUzMTc4LCAiaHR0cF9wb3J0Ijog
NTk5NSwgImhvc3RfaXAiOiAiMTAuODQuMTMuNDAlfQ==",
        type: "BgpConnectivity",
        severity: 4
      }
    ]
  },
  BgpRouterState: {...}
}
```

In the example:

- Alerts are raised on a per-UE basis and can be retrieved by a GET on a UE.
- An **ack** indicates if the alert has been acknowledged or not.
- A **token** is used by clients when requesting acknowledgements

## Analytics APIs for Alerts

The following examples show the API to use to display alerts and alarms and to acknowledge alarms.

- To retrieve a list of alerts raised against the control node named **aXXsYY**.

```
GET
http://<analytics-ip>:<rest-api-port>/analytics/uves/control-node/aXXsYY&filter=UVEAlarms
```

This is available for all UVE table types.

- To retrieve a list of all alarms in the system.

```
GET http://<analytics-ip>:<rest-api-port>/analytics/alarms
```

- To acknowledge an alarm.

```
POST http://<analytics-ip>:<rest-api-port>/analytics/alarms/acknowledge
Body: {"table": <object-type>,"name": <key>,"type": <alarm type>,"token":
<token>}
```

Acknowledged and unacknowledged alarms can be queried specifically using the following URL query parameters along with the GET operations listed previously.

```
ackFilt=True
ackFilt=False
```

## Analytics APIs for SSE Streaming

The following examples show the API to use to retrieve all or portions of SE streams.

- To retrieve an SSE-based stream of UVE updates for the control node alarms.

```
GET
http://<analytics-ip>:<rest-api-port>/analytics/uve-stream?tablefilt=control-node
```

This is available for all UVE table types. If the **tablefilt** URL query parameter is not provided, all UVEs are retrieved.

- To retrieve only the alerts portion of the SSE-based stream of UVE updates instead of the entire content.

```
GET
http://<analytics-ip>:<rest-api-port>/analytics/alarms-stream?tablefilt=control-node
```

This is available for all UVE table types. If the **tablefilt** URL query parameter is not provided, all UVEs are retrieved.

## Built-in Node Alerts

The following built-in node alerts can be retrieved using the APIs listed in *Analytics APIs for Alerts*.

```
control node: {
  PartialSysinfoControl: "Basic System Information is absent for this node in
  BgpRouterState.build_info",
  ProcessStatus: "NodeMgr reports abnormal status for process(es) in
  NodeStatus.process_info",
  XmppConnectivity: "Not enough XMPP peers are up in
  BgpRouterState.num_up_bgp_peer",
  BgpConnectivity: "Not enough BGP peers are up in
  BgpRouterState.num_up_bgp_peer",
  AddressMismatch: "Mismatch between configured IP Address and operational IP
  Address",
  ProcessConnectivity: "Process(es) are reporting non functional components in
  NodeStatus.process_status"
},

vrouter: {
  PartialSysinfoCompute: "Basic System Information is absent for this node in
  VrouterAgent.build_info",
  ProcessStatus: "NodeMgr reports abnormal status for process(es) in
  NodeStatus.process_info",
  ProcessConnectivity: "Process(es) are reporting non functional components in
  NodeStatus.process_status",
  VrouterInterface: "VrouterAgent has interfaces in error state in
  VrouterAgent.error_intf_list",
  VrouterConfigAbsent: "Vrouter is not present in Configuration",
},

config node: {
  PartialSysinfoConfig: "Basic System Information is absent for this node in
  ModuleCpuState.build_info",
  ProcessStatus: "NodeMgr reports abnormal status for process(es) in
```

```
NodeStatus.process_info",
ProcessConnectivity: "Process(es) are reporting non functional components in
NodeStatus.process_status"
},

analytics node: {
ProcessStatus: "NodeMgr reports abnormal status for process(es) in
NodeStatus.process_info"
PartialSysinfoAnalytics: "Basic System Information is absent for this node in
CollectorState.build_info",
ProcessConnectivity: "Process(es) are reporting non functional components in
NodeStatus.process_status"
},

database node: {
ProcessStatus: "NodeMgr reports abnormal status for process(es) in
NodeStatus.process_info",
ProcessConnectivity: "Process(es) are reporting non functional components in
NodeStatus.process_status"
},
```

**Related  
Documentation**

- [Monitoring the System on page 280](#)
- [Debugging Processes Using the Contrail Introspect Feature on page 282](#)
- [Monitor > Infrastructure > Dashboard on page 286](#)
- [Monitor > Infrastructure > Control Nodes on page 288](#)
- [Monitor > Infrastructure > Virtual Routers on page 295](#)
- [Monitor > Infrastructure > Analytics Nodes on page 306](#)
- [Monitor > Infrastructure > Config Nodes on page 311](#)
- [Monitor > Networking on page 314](#)
- [Understanding Flow Sampling](#)
- [Query > Flows on page 322](#)
- [Query > Logs on page 329](#)
- [Example: Debugging Connectivity Using Monitoring for Troubleshooting on page 334](#)

---

## Underlay Overlay Mapping in Contrail

---

- [Overview: Underlay Overlay Mapping using Contrail Analytics on page 250](#)
- [Underlay Overlay Analytics Available in Contrail on page 250](#)
- [Architecture and Data Collection on page 251](#)
- [New Processes/Services for Underlay Overlay Mapping on page 251](#)
- [External Interfaces Configuration for Underlay Overlay Mapping on page 252](#)
- [Physical Topology on page 252](#)
- [SNMP Configuration on page 252](#)
- [Link Layer Discovery Protocol \(LLDP\) Configuration on page 253](#)

- [IPFIX and sFlow Configuration on page 253](#)
- [Sending pRouter Information to the SNMP Collector in Contrail on page 254](#)
- [pRouter UVEs on page 255](#)
- [Contrail User Interface for Underlay Overlay Analytics on page 256](#)
- [Viewing Topology to the Virtual Machine Level on page 256](#)
- [Viewing the Traffic of any Link on page 257](#)
- [Trace Flows on page 257](#)
- [Search Flows and Map Flows on page 258](#)
- [Overlay to Underlay Flow Map Schemas on page 259](#)
- [Module Operations for Overlay Underlay Mapping on page 262](#)
- [SNMP Collector Operation on page 262](#)
- [Topology Module Operation on page 263](#)
- [IPFIX and sFlow Collector Operation on page 263](#)
- [Troubleshooting Underlay Overlay Mapping on page 264](#)
- [Script to add pRouter Objects on page 265](#)

## Overview: Underlay Overlay Mapping using Contrail Analytics

Today's cloud data centers consist of large collections of interconnected servers that provide computing and storage capacity to run a variety of applications. The servers are connected with redundant TOR switches, which in turn, are connected to spine routers. The cloud deployment is typically shared by multiple tenants, each of whom usually needs multiple isolated networks. Multiple isolated networks can be provided by overlay networks that are created by forming tunnels (for example, gre, ip-in-ip, mac-in-mac) over the underlay or physical connectivity.

As data flows in the overlay network, Contrail can provide statistics and visualization of the traffic in the underlay network.

## Underlay Overlay Analytics Available in Contrail

Starting with Contrail Release 2.20, you can view a variety of analytics related to underlay and overlay traffic in the Contrail Web user interface. The following are some of the analytics that Contrail provides for statistics and visualization of overlay underlay traffic.

- View the topology of the underlay network.

A user interface view of the physical underlay network with a drill down mechanism to show connected servers (contrail computes) and virtual machines on the servers.

- View the details of any element in the topology.

You can view details of a pRouter, vRouter, or virtual machine link between two elements. You can also view traffic statistics in a graphical view corresponding to the selected element.

- View the underlay path of an overlay flow.

Given an overlay flow, you can get the underlay path used for that flow and map the path in the topology view.

## Architecture and Data Collection

Accumulation of the data to map an overlay flow to its underlay path is performed in several steps across Contrail modules.

The following outlines the essential steps:

1. The SNMP collector module polls physical routers.

The SNMP collector module receives the authorizations and configurations of the physical routers from the Contrail config module, and polls all of the physical routers, using SNMP protocol. The collector uploads the data to the Contrail analytics collectors. The SNMP information is stored in the pRouter UVEs (physical router user visible entities).

2. IPFIX and sFlow protocols are used to collect the flow statistics.

The physical router is configured to send flow statistics to the collector, using one of the collection protocols: Internet Protocol Flow Information Export (IPFIX) or sFlow (an industry standard for sampled flow of packet export at Layer 2).

3. The topology module reads the SNMP information.

The Contrail topology module reads SNMP information from the pRouter UVEs from the analytics API, computes the neighbor list, and writes the neighbor information into the pRouter UVEs. This neighbor list is used by the Contrail WebUI to display the physical topology.

4. The Contrail user interface reads and displays the topology and statistics.

The Contrail user interface module reads the topology information from the Contrail analytics and displays the physical topology. It also uses information stored in the analytics to display graphs for link statistics, and to show the map of the overlay flows on the underlay network.

## New Processes/Services for Underlay Overlay Mapping

The **contrail-snmp-collector** and the **contrail-topology** are new daemons that are both added to the **contrail-analytics** node. The **contrail-analytics** package contains these new features and their associated files. The **contrail-status** displays the new services.

**Example:** The following is an example of using **contrail-status** to show the status of the new process and service for underlay overlay mapping.

```
user@host:~# contrail-status

== Contrail Control ==

supervisor-control:      active

contrail-control         active

...
```

```

== Contrail Analytics ==

supervisor-analytics:    active

...

contrail-query-engine    active

contrail-snmp-collector  active

contrail-topology        active

```

**Example: Service Command** The **service** command can be used to start, stop, and restart the new services. See the following example.

```

user@host:~# service contrail-snmp-collector status

contrail-snmp-collector    RUNNING    pid 12179, uptime 1 day, 14:59:11

```

## External Interfaces Configuration for Underlay Overlay Mapping

This section outlines the external interface configurations necessary for successful underlay overlay mapping for Contrail analytics.


### Physical Topology

The typical physical topology includes:

- Servers connected to the ToR switches.
- ToR switches connected to spine switches.
- Spine switches connected to core switches.

The following is an example of how the topology is depicted in the Contrail WebUI analytics.

Figure 89: Analytics Topology



### SNMP Configuration

Configure SNMP on the physical devices so that the **contrail-snmp-collector** can read SNMP data.

The following shows an example SNMP configuration from a Juniper Networks device.

```
set snmp community public authorization read-only
```

## Link Layer Discovery Protocol (LLDP) Configuration

Configure LLDP on the physical device so that the **contrail-snmp-collector** can read the neighbor information of the routers.

The following is an example of LLDP configuration on a Juniper Networks device.

```
set protocols lldp interface all
```

```
set protocols lldp-med interface all
```

## IPFIX and sFlow Configuration

Flow samples are sent to the **contrail-collector** by the physical devices. Because the **contrail-collector** supports the sFlow and IPFIX protocols for receiving flow samples, the physical devices, such as MX Series devices or ToR switches, must be configured to send samples using one of those protocols.

### Example: sFlow Configuration

The following shows a sample sFlow configuration. In the sample, the IP variable *<source ip>* refers to the loopback or IP that can be reachable of the device that acts as an sflow source, and the other IP variable *<collector\_IP\_data>* is the address of the collector device.

```
root@host> show configuration protocols sflow | display set

set protocols sflow polling-interval 0

set protocols sflow sample-rate ingress 10

set protocols sflow source-ip <source ip>4

set protocols sflow collector <collector_IP_data> udp-port 6343

set protocols sflow interfaces ge-0/0/0.0

set protocols sflow interfaces ge-0/0/1.0

set protocols sflow interfaces ge-0/0/2.0

set protocols sflow interfaces ge-0/0/3.0

set protocols sflow interfaces ge-0/0/4.0
```

### Example: IPFIX Configuration

The following is a sample IPFIX configuration from a Juniper Networks device. The IP address variable *<ip\_sflow collector>* represents the sflow collector (control-collector analytics node) and *<source ip>* represents the source (outgoing) interface on the router/switch device used for sending flow data to the collector. This could also be the lo0 address, if it is reachable from the Contrail cluster.

```
root@host> show configuration chassis | display set

set chassis tfeb slot 0 sampling-instance sample-ins1

set chassis network-services all-ethernet
```

```
root@host> show configuration chassis tfeb | display set

set chassis tfeb slot 0 sampling-instance sample-ins1


root@host > show configuration services flow-monitoring | display set

set services flow-monitoring version-ipfix template t1 flow-active-timeout 30

set services flow-monitoring version-ipfix template t1 flow-inactive-timeout 30

set services flow-monitoring version-ipfix template t1 template-refresh-rate packets 10

set services flow-monitoring version-ipfix template t1 ipv4-template


root@host > show configuration interfaces | display set | match sampling

set interfaces ge-1/0/0 unit 0 family inet sampling input

set interfaces ge-1/0/1 unit 0 family inet sampling input


root@host> show configuration forwarding-options sampling | display set

set forwarding-options sampling instance sample-ins1 input rate 1

set forwarding-options sampling instance sample-ins1 family inet output flow-server <ip_sflow collector> port 4739

set forwarding-options sampling instance sample-ins1 family inet output flow-server <ip_sflow collector> version-ipfix template t1

set forwarding-options sampling instance sample-ins1 family inet output inline-jflow source-address <source ip>
```

## Sending pRouter Information to the SNMP Collector in Contrail

Information about the physical devices must be sent to the SNMP collector before the full analytics information can be read and displayed. Typically, the pRouter information is taken from the **contrail-config** file, but the information can also be sent to the SNMP collector by means of a **device.ini** file.

*SNMP collector getting pRouter information from contrail-config file*

The physical routers are added to the **contrail-config** by using the Contrail user interface or by using direct API, by means of provisioning or other scripts. Once the configuration is in the **contrail-config**, the **contrail-snmp-collector** gets the physical router information from **contrail-config**. The SNMP collector uses this list and the other configuration parameters to perform SNMP queries and to populate pRouter UVEs.



Figure 90: Add Physical Router Window

The screenshot shows the Juniper Networks configuration interface. On the left, a sidebar lists navigation options: Infrastructure, Physical Devices, Physical Routers, Interfaces, Networking, Services, and DNS. The main panel displays the 'Physical Routers' configuration page. A modal window titled 'Add Physical Router' is open, showing the following fields and sections:

- Name:** new-router
- Vendor:** (empty field)
- Model:** (empty field)
- Management IP:** 1.1.1.1
- Tunnel Source IP:** (empty field)
- User Credentials:** (empty field)
- Virtual Router:** (empty field)
- BGP Router:** (empty field)
- SNMP Credentials:**
  - Version:** 2 (selected), 3
  - Community:** public

At the bottom right of the modal window are 'Cancel' and 'Save' buttons. The background shows a list of physical routers with checkboxes and names: a7-ex2, a7-ex3, a7-mx80-1, and a7-mx80-2. A status bar at the bottom indicates 'Total: 4 records' and '50 Records'.

## pRouter UVEs

pRouter UVEs are accessed from the REST APIs on your system from **contrail-analytics-api**, using a URL of the form:

**http://<host ip>:8081/analytics/uves/prouters**

The following is sample output from a pRouter REST API:

Figure 91: Sample Output From a pRouter REST API

```
[
  {
    href: "http://10.84.63.130:8081/analytics/uves/prouter/a7-mx80-1?flat",
    name: "a7-mx80-1"
  },
  {
    href: "http://10.84.63.130:8081/analytics/uves/prouter/a7-mx80-2?flat",
    name: "a7-mx80-2"
  },
  {
    href: "http://10.84.63.130:8081/analytics/uves/prouter/a7-ex3?flat",
    name: "a7-ex3"
  },
  {
    href: "http://10.84.63.130:8081/analytics/uves/prouter/a7-ex2?flat",
    name: "a7-ex2"
  }
]
```

Details of a pRouter UVE can be obtained from your system, using a URL of the following form:

**http://<host ip>:8081/analytics/uves/prouter/a7-ex3?flat**

The following is sample output of a pRouter UVE.

Figure 92: Sample Output From a pRouter UVE

```

{
  - PRouterFlowEntry: {
    flow_export_source_ip: "10.84.63.114"
  },
  - PRouterLinkEntry: {
    - link_table: [
      - {
        remote_interface_name: "ge-1/0/1",
        local_interface_name: "ge-0/0/0.0",
        remote_interface_index: 517,
        local_interface_index: 503,
        type: 1,
        remote_system_name: "a7-mx80-1"
      },
      - {
        remote_interface_name: "ge-1/0/1",
        local_interface_name: "ge-0/0/1.0",
        remote_interface_index: 517,
        local_interface_index: 505,
        type: 1,
        remote_system_name: "a7-mx80-2"
      },
      - {
        remote_interface_name: "eth1",
        local_interface_name: "ge-0/0/2.0",
        remote_interface_index: 1,
        local_interface_index: 507,
        type: 2,
        remote_system_name: "a7s35"
      },
      - {
        remote_interface_name: "eth1",
        local_interface_name: "ge-0/0/3.0",
        remote_interface_index: 1,
        local_interface_index: 509,
        type: 2,
        remote_system_name: "a7s36"
      }
    ]
  },
  - PRouterEntry: {
    + ipMib: [...],
    + ifTable: [...],
    + ifXTable: [...],
    + arpTable: [...],
    + lldpTable: {...},
    + ifStats: [...]
  }
}

```

s042435


## Contrail User Interface for Underlay Overlay Analytics

The topology view and related functionality is accessed from the Contrail Web user interface, **Monitor > Physical Topology**.

### Viewing Topology to the Virtual Machine Level

In the Contrail user interface, it is possible to drill down through displayed topology to the virtual machine level. The following diagram shows the virtual machines instantiated on a7s36 vRouter and the full physical topology related to each.

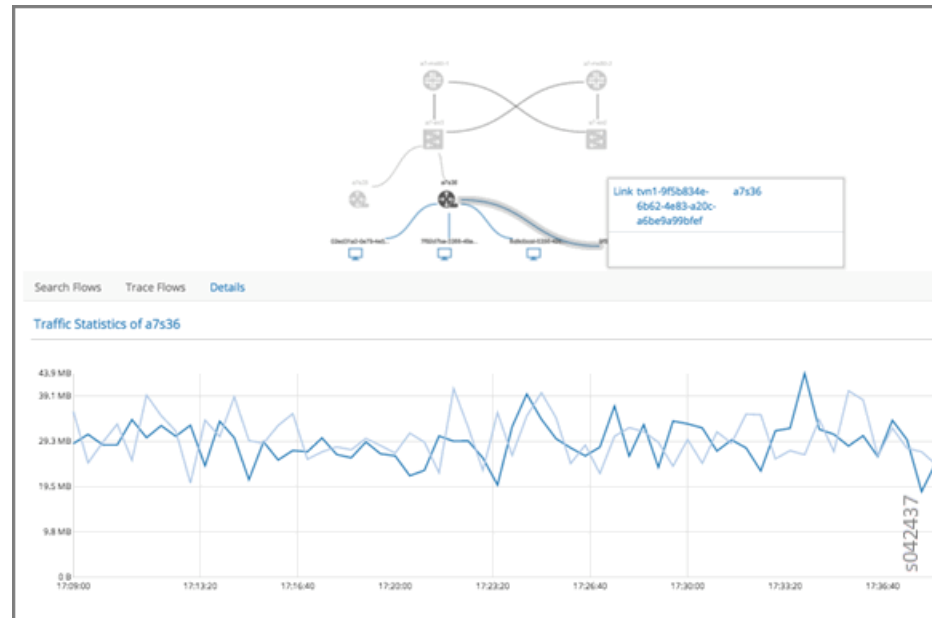
Figure 93: Physical Topology Related to a vRouter



### Viewing the Traffic of any Link

At **Monitor > Physical Topology**, double click any link on the topology to display the traffic statistics graph for that link. The following is an example.

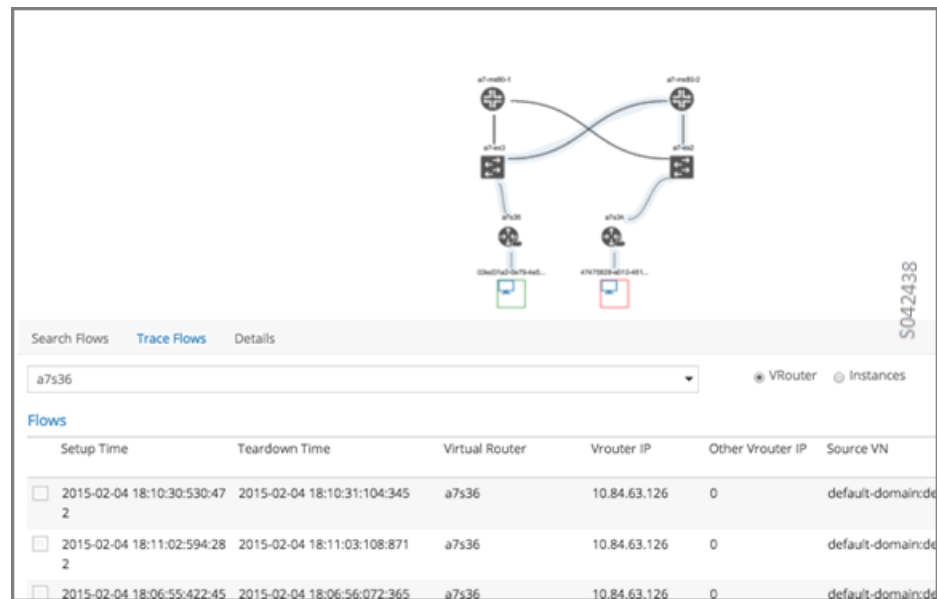
Figure 94: Traffic Statistics Graph



### Trace Flows

Click the **Trace Flows** tab to see a list of active flows. To see the path of a flow, click a flow in the active flows list, then click the **Trace Flow** button. The path taken in the underlay by the selected flow displays. The following is an example.

Figure 95: List of Active Flows



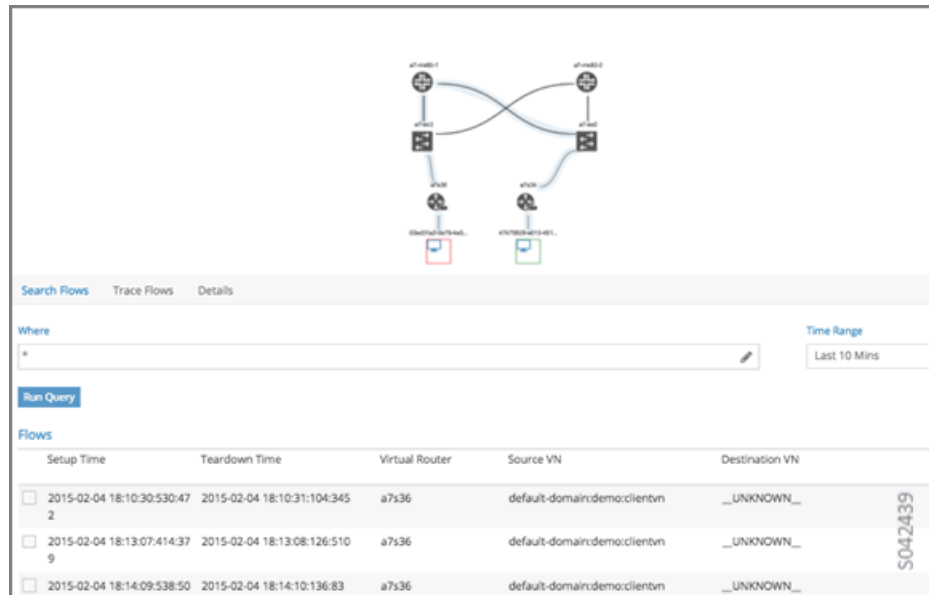
#### Limitations of Trace Flow Feature

Because the Trace Flow feature uses ip traceroute to determine the path between the two vRouters involved in the flow, it has the same limitations as the ip traceroute, including that Layer 2 routers in the path are not listed, and therefore do not appear in the topology.

#### Search Flows and Map Flows

Click the **Search Flows** tab to open a search dialog, then click the **Search** button to list the flows that match the search criteria. You can select a flow from the list and click **Map Flow** to display the underlay path taken by the selected flow in the topology. The following is an example.

Figure 96: Underlay Path



## Overlay to Underlay Flow Map Schemas

The schema to query the underlay mapping information for an overlay flow is obtained from a REST API, which can be accessed on your system using a URL of the following form:

<http://<host ip>:8081/analytics/table/OverlayToUnderlayFlowMap/schema>

**Example: Overlay to Underlay Flow Map Schema**

```
{
  "type": "FLOW",
  "columns": [
    {
      "datatype": "string", "index": true, "name": "o_svn", "select": false,
      "suffixes": ["o_sip"]
    },
    {
      "datatype": "string", "index": false, "name": "o_sip", "select": false,
      "suffixes": null
    },
    {
      "datatype": "string", "index": true, "name": "o_dvn", "select": false,
      "suffixes": ["o_dip"]
    },
    {
      "datatype": "string", "index": false, "name": "o_dip", "select": false,
      "suffixes": null
    },
    {
      "datatype": "int", "index": false, "name": "o_sport", "select": false,
      "suffixes": null
    },
    {
      "datatype": "int", "index": false, "name": "o_dport", "select": false,
      "suffixes": null
    },
    {
      "datatype": "int", "index": true, "name": "o_protocol", "select": false,
      "suffixes": ["o_sport", "o_dport"]
    },
    {
      "datatype": "string", "index": true, "name": "o_vrouter", "select": false,
      "suffixes": null
    }
  ]
}
```

```
{
  "datatype": "string", "index": false, "name": "u_prouter", "select": null,
  "suffixes": null},
{
  "datatype": "int", "index": false, "name": "u_pifindex", "select": null,
  "suffixes": null},
{
  "datatype": "int", "index": false, "name": "u_vlan", "select": null,
  "suffixes": null},
{
  "datatype": "string", "index": false, "name": "u_sip", "select": null,
  "suffixes": null},
{
  "datatype": "string", "index": false, "name": "u_dip", "select": null,
  "suffixes": null},
{
  "datatype": "int", "index": false, "name": "u_sport", "select": null,
  "suffixes": null},
{
  "datatype": "int", "index": false, "name": "u_dport", "select": null,
  "suffixes": null},
{
  "datatype": "int", "index": false, "name": "u_protocol", "select": null,
  "suffixes": null},
{
  "datatype": "string", "index": false, "name": "u_flowtype", "select": null,
  "suffixes": null},
{
  "datatype": "string", "index": false, "name": "u_otherinfo", "select": null,
  "suffixes": null}}}]}
```

The schema for underlay data across pRouters is defined in the Contrail installation at:

<http://<host ip>:8081/analytics/table/StatTable.UFlowData.flow/schema>

#### Example: Flow Data Schema for Underlay

```
{
  "type": "STAT",
  "columns": [
    {
      "datatype": "string", "index": true, "name": "Source", "suffixes": null},
    {
      "datatype": "int", "index": false, "name": "T", "suffixes": null},
    {
      "datatype": "int", "index": false, "name": "CLASS(T)", "suffixes": null},
    {
      "datatype": "int", "index": false, "name": "T=", "suffixes": null},
    {
      "datatype": "int", "index": false, "name": "CLASS(T=)", "suffixes": null},
    {
      "datatype": "uuid", "index": false, "name": "UUID", "suffixes": null},
    {
      "datatype": "int", "index": false, "name": "COUNT(flow)", "suffixes": null},
    {
      "datatype": "string", "index": true, "name": "name", "suffixes":
        ["flow.pifindex"]},
    {
      "datatype": "int", "index": false, "name": "flow.pifindex", "suffixes": null},
    {
      "datatype": "int", "index": false, "name": "SUM(flow.pifindex)", "suffixes":
        null},
```

```

{"datatype": "int", "index": false, "name": "CLASS(flow.pifindex)", "suffixes":
  null},

{"datatype": "int", "index": false, "name": "flow.sport", "suffixes": null},

{"datatype": "int", "index": false, "name": "SUM(flow.sport)", "suffixes":
  null},

{"datatype": "int", "index": false, "name": "CLASS(flow.sport)", "suffixes":
  null},

{"datatype": "int", "index": false, "name": "flow.dport", "suffixes": null},

{"datatype": "int", "index": false, "name": "SUM(flow.dport)", "suffixes":
  null},

{"datatype": "int", "index": false, "name": "CLASS(flow.dport)", "suffixes":
  null},

{"datatype": "int", "index": true, "name": "flow.protocol", "suffixes":
  ["flow.sport", "flow.dport"]},

{"datatype": "int", "index": false, "name": "SUM(flow.protocol)", "suffixes":
  null},

{"datatype": "int", "index": false, "name": "CLASS(flow.protocol)", "suffixes":
  null},

{"datatype": "string", "index": true, "name": "flow.sip", "suffixes": null},

{"datatype": "string", "index": true, "name": "flow.dip", "suffixes": null},

{"datatype": "string", "index": true, "name": "flow.vlan", "suffixes": null},

{"datatype": "string", "index": false, "name": "flow.flowtype", "suffixes":
  null},

{"datatype": "string", "index": false, "name": "flow.otherinfo", "suffixes":
  null}}}

```

#### Example: Typical Query for Flow Map

The following is a typical query. Internally, the **analytics-api** performs a query into the **FlowRecordTable**, then into the **StatTable.UFlowData.flow**, to return list of (**prouter**, **pifindex**) pairs that give the underlay path taken for the given overlay flow.

```

FROM

OverlayToUnderlayFlowMap

SELECT

prouter, pifindex

WHERE

o_svn, o_sip, o_dvn, o_dip, o_sport, o_dport, o_protocol = <overlay flow>

```

## Module Operations for Overlay Underlay Mapping

### SNMP Collector Operation

The Contrail SNMP collector uses a Net-SNMP library to talk to a physical router or any SNMP agent. Upon receiving SNMP packets, the data is translated to the Python dictionary, and corresponding UVE objects are created. The UVE objects are then posted to the SNMP collector.

The SNMP module sleeps for some configurable period, then forks a collector process and waits for the process to complete. The collector process goes through a list of devices to be queried. For each device, it forks a greenlet task (Python coroutine), accumulates SNMP data, writes the summary to a JSON file, and exits. The parent process then reads the JSON file, creates UVEs, sends the UVEs to the collector, then goes to sleep again.

The pRouter UVE sent by the SNMP collector carries only the raw MIB information.

**Example: pRouter  
Entry Carried in  
pRouter UVE**

The definition below shows the **pRouterEntry** carried in the **pRouterUVE**. Additionally, an example **LldpTable** definition is shown.

The following create a virtual table as defined by:

```
http://<host ip>:8081/analytics/table/StatTable.UFlowData.flow/schema

struct LldpTable {
    1: LldpLocalSystemData lldpLocalSystemData
    2: optional list<LldpRemoteSystemsData> lldpRemoteSystemsData
}

struct PRouterEntry {
    1: string name (key="ObjectPRouter")
    2: optional bool deleted
    3: optional LldpTable lldpTable
    4: optional list<ArpTable> arpTable
    5: optional list<IfTable> ifTable
    6: optional list<IfXTable> ifXTable
    7: optional list<IfStats> ifStats (tags="name:.ifIndex")
    8: optional list<IpMib> ipMib
}

uve sandesh PRouterUVE {
    1: PRouterEntry data
}
```



## Topology Module Operation

The topology module reads UVEs posted by the SNMP collector and computes the neighbor table, populating the table with remote system name, local and remote interface names, the remote type (pRouter or vRouter) and local and remote ifindices. The topology module sleeps for a while, reads UVEs, then computes the neighbor table and posts the UVE to the collector.

The pRouter UVE sent by the topology module carries the neighbor list, so the clients can put together all of the pRouter neighbor lists to compute the full topology.

The corresponding pRouter UVE definition is the following.

```
struct LinkEntry {
    1: string remote_system_name
    2: string local_interface_name
    3: string remote_interface_name
    4: RemoteType type
    5: i32 local_interface_index
    6: i32 remote_interface_index
}

struct PRouterLinkEntry {
    1: string name (key="ObjectPRouter")
    2: optional bool deleted
    3: optional list<LinkEntry> link_table
}

uve sandesh PRouterLinkUVE {
    1: PRouterLinkEntry data
}
```

## IPFIX and sFlow Collector Operation

An IPFIX and sFlow collector has been implemented in the Contrail collector. The collector receives the IPFIX and sFlow samples and stores them as statistics samples in the analytics database.

**Example: IPFIX sFlow  
Collector Data**

The following definition shows the data stored for the statistics samples and the indices that can be used to perform queries.

```
struct UFlowSample {  
    1: u64 pifindex  
    2: string sip  
    3: string dip  
    4: u16 sport  
    5: u16 dport  
    6: u16 protocol  
    7: u16 vlan  
    8: string flowtype  
    9: string otherinfo  
}  
  
struct UFlowData {  
    1: string name (key="ObjectPRouterIP")  
    2: optional bool deleted  
    3: optional list<UFlowSample> flow (tags="name:.pifindex, .sip, .dip,  
    .protocol:.sport, .protocol:.dport, .vlan")  
}
```

## Troubleshooting Underlay Overlay Mapping

This section provides a variety of links where you can research errors that may occur with underlay overlay mapping.

**System Logs** Logs for **contrail-snmp-collector** and **contrail-topology** are in the following locations on an installed Contrail system:

`/var/log/contrail/contrail-snmp-collector-stdout.log`

`/var/log/contrail/contrail-topology.log`

**Introspect Utility** Use URLs of the following forms on your Contrail system to access the introspect utilities for SNMP data and for topology data.

- SNMP data introspect

`http://<host ip>:5920/Snh_SandeshUVECacheReq?x=PRouterEntry`

- Topology data introspect

`http://<host ip>:5921/Snh_SandeshUVECacheReq?x=PRouterLinkEntry`

## Script to add pRouter Objects

The usual mechanism for adding pRouter objects to **contrail-config** is through Contrail UI. But you also have the ability to add these objects using the Contrail **vnc-api**. To add one pRouter, save the file with the name **cfg-snmp.py**, and then execute the command as shown:

**python cfg-snmp.py**

**Example: Content for  
cfg-snmp.py**

```
#!/python
from vnc_api import vnc_api
from vnc_api.gen.resource_xsd import SNMPCredentials

vnc = vnc_api.VncApi('admin', 'abcde123', 'admin')
apr = vnc_api.gen.resource_client.PhysicalRouter(name='a7-mx80-1')
apr.set_physical_router_management_ip('ip_address')
apr.set_physical_router_dataplane_ip('ip_address')
apr.set_physical_router_snmp_credentials(SNMPCredentials(version=2,
v2_community='public'))

vnc.physical_router_create(apr)

#$ABC123

apr = vnc_api.gen.resource_client.PhysicalRouter(name='a7-mx80-2')
apr.set_physical_router_management_ip('ip_address')
apr.set_physical_router_dataplane_ip('ip_address')
apr.set_physical_router_snmp_credentials(SNMPCredentials(version=2,
v2_community='public'))
```

```
vnc.physical_router_create(apr)

#$ABC123'

apr = vnc_api.gen.resource_client.PhysicalRouter(name='a7-ex3')

apr.set_physical_router_management_ip('source_ip')

apr.set_physical_router_dataplane_ip('source_ip')

apr.set_physical_router_snmp_credentials(SNMPCredentials(version=2,
v2_community='public'))

vnc.physical_router_create(apr)

#$ABC123'

apr = vnc_api.gen.resource_client.PhysicalRouter(name='a7-ex2')

apr.set_physical_router_management_ip('ip_address')

apr.set_physical_router_dataplane_ip('ip_address')

apr.set_physical_router_snmp_credentials(SNMPCredentials(version=2,
v2_community='public'))

vnc.physical_router_create(apr)

#$ABC123'
```

- Related Documentation**
- [Contrail Analytics Overview on page 245](#)
  - [Contrail Alerts on page 246](#)

# Configuring Contrail Analytics

- [Analytics Scalability on page 267](#)
- [High Availability for Analytics on page 268](#)
- [System Log Receiver in Contrail Analytics on page 269](#)
- [Sending Flow Messages to the Contrail System Log on page 270](#)
- [Ceilometer Support in a Contrail Cloud on page 271](#)

## Analytics Scalability

---

The Contrail monitoring and analytics services (*collector* role) collect and store data generated by various system components and provide the data to the Contrail interface by means of representational state transfer (REST) application program interface (API) queries.

The Contrail components are horizontally scalable to ensure consistent performance as the system grows. Scalability is provided for the generator components (*control* and *compute* roles) and for the REST API users (*webui* role).


This section provides a brief description of the recommended configuration of analytics in Contrail to achieve horizontal scalability.

The following is the recommended locations for the various component roles of the Contrail system for a 5-node configuration.

- Node 1 —config role, web-ui role
- Node 2 —control role, analytics role, database role
- Node 3 —control role, analytics role, database role
- Node 4 —compute role
- Node 5 —compute role

[Figure 97 on page 268](#) illustrates scalable connections for analytics in a 5-node system, with the nodes configured for roles as recommended above. The analytics load is distributed between the two analytics nodes. This configuration can be extended to any number of analytics nodes.

Figure 97: Analytics Scalability



The analytics nodes collect and store data and provide this data through various REST API queries. Scalability is provided for the control nodes, the compute nodes, and the REST API users, with the API output displayed in the Contrail user interface. As the number of control and compute nodes increase in the system, the analytics nodes can also be increased.

## High Availability for Analytics

Contrail supports multiple instances of analytics for high availability and load balancing.

Contrail analytics provides two broad areas of functionality:

- **contrail-collector** —Receives status, logs, and flow information from all Contrail processing elements (for example, generators) and records them.

Every generator is connected to one of the **contrail-collector** instances at any given time. If an instance fails (or is shut down), all the generators that are connected to it are automatically moved to another functioning instance, typically in a few seconds or less. Some messages may be lost during this movement. UVEs are resilient to message loss, so the state shown in a UVE is kept consistent to the state in the generator.

- **contrail-opserver** —Provides an external API to report UVEs and to query logs and flows.

Each analytics component exposes a northbound REST API represented by the **contrail-opserver** service (port 8081) so that the failure of one analytics component or one **contrail-opserver** service should not impact the operation of other instances.

These are the ways to manage connectivity to the **contrail-opserver** endpoints:

- Periodically poll the **contrail-opserver** service on a set of analytics nodes to determine the list of functioning endpoints, then make API requests from one or more of the functioning endpoints.
- Subscribe to the Contrail Discovery Service to get a list of functioning endpoints. If there are any failures, it can take 5-30 minutes for the Contrail Discovery Service to send an update.

The Contrail user interface makes use of the same northbound REST API to present dashboards, and reacts to any **contrail-opserver** high availability event automatically, using the Contrail Discovery Service.

## System Log Receiver in Contrail Analytics

- [Overview on page 269](#)
- [Redirecting System Logs to Contrail Collector on page 269](#)
- [Exporting Logs from Contrail Analytics on page 269](#)

### Overview

The **contrail-collector** process on the Contrail Analytics node can act as a system log receiver.

### Redirecting System Logs to Contrail Collector

You can enable the **contrail-collector** to receive system logs by giving a valid **syslog\_port** as a command line option:

```
--DEFAULT.syslog_port <arg>
```

or by adding **syslog\_port** in the **DEFAULT** section of the configuration file at **/etc/contrail/contrail-collector.conf**.

For nodes to send system logs to the **contrail-collector**, the system log configuration for the node should be set up to direct the system logs to **contrail-collector**.

**Example** Add the following line in **/etc/rsyslog.d/50-default.conf** on an Ubuntu system to redirect the system logs to **contrail-collector**.

```
*.* @<collector_ip>:<collector_syslog_port> :: @ for udp, @@ for tcp
```

The logs can be retrieved by using Contrail tool, either by using the **contrail-logs** utility on the analytics node or by using the Contrail user interface on the system log query page.

### Exporting Logs from Contrail Analytics

You can also export logs stored in Contrail analytics to another system log receiver by using the **contrail-logs** utility.

The **contrail-logs** utility can take these options: **--send-syslog**, **--syslog-server**, **--syslog-port**, to query Contrail analytics, then send the results as system logs to a system log server. This is an on-demand command, one can write a cron job or a job that continuously invokes **contrail-logs** to achieve continuous sending of logs to another system log server.

## Sending Flow Messages to the Contrail System Log

---

The **contrail-vrouter-agent** can be configured to send flow messages and other messages to the system log (syslog). To send flow messages to syslog, configure the following parameters in **/etc/contrail/contrail-vrouter-agent.conf**.

The following parameters are under the section **DEFAULT**:

- **log\_flow=1**—Enables logging of all flow messages.
- **use\_syslog=1**—Enables sending of all messages, including flow messages, to syslog.
- **syslog\_facility=LOG\_LOCAL0**—Enables sending messages from the **contrail-vrouter-agent** to the syslog, using the facility **LOCAL0**. You can configure **LOCAL0** to your required facility.
- **log\_level=SYS\_INFO**—Changes the logging level of **contrail-vrouter-agent** to INFO.

If syslog is enabled, flow messages are *not* sent to Contrail Analytics because the two destinations are mutually exclusive.

Flow log sampling settings apply regardless of the flow log destination specified. If sampling is enabled, the syslog messages will be sampled using the same rules that would apply to Contrail Analytics. If non-sampled flow data is required, sampling must be disabled by means of configuration settings.

Flow events for termination will include both the appropriate tear-down fields and the appropriate setup fields.

The flow messages will be sent to the syslog with a severity of INFO.

The user can configure the remote system log (**rsyslog**) on the compute node to send syslog messages with facility LOCAL0, severity of INFO (and lower), to the remote syslog server. Messages with a higher severity than INFO can be logged to a local file to allow for debugging.

Flow messages appear in the syslog in a format similar to the following log example:

```
May 24 14:40:13 a7s10 contrail-vrouter-agent[29930]: 2016-05-24 Tue 14:40:13:921.098
PDT a7s10 [Thread 139724471654144, Pid 29930]: [SYS_INFO]: FlowLogDataObject:
flowdata= [ [ [ flowuuid = 7ea8bf8f-b827-496e-b93e-7622a0c8eeea direction_ing = 1
sourcevn = default-domain:mock-gen-test:vn8 sourceip = 1.0.0.9 destvn =
default-domain:mock-gen-test:vn58 destip = 1.0.0.59 protocol = 1 sport = -29520 dport =
20315 setup_time = 1464125225556930 bytes = 1035611592 packets = 2024830 diff_bytes
= 27240 diff_packets = 40 ], ] ]
```



**NOTE:** Several individual flow messages might be packed into a single syslog message for improved efficiency.

---



## Ceilometer Support in a Contrail Cloud

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 on page 271](#)
- [Ceilometer Details on page 271](#)
- [Verification of Ceilometer Operation on page 272](#)
- [Contrail Ceilometer Plugin on page 274](#)
- [Ceilometer Installation and Provisioning on page 276](#)

### 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 Fabric with **enable\_ceilometer = True** in the **testbed.py** file.
- Alternately, provisioned using Server Manager with **enable\_ceilometer = True** in the **cluster.json** or **cluster** configuration.



**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	a726f93a-65fa-4cad-828b-54dbfcf4a119	cumulative	B		None
ip.floating.receive.packets	a726f93a-65fa-4cad-828b-54dbfcf4a119	cumulative	packet		None
ip.floating.transmit.bytes	a726f93a-65fa-4cad-828b-54dbfcf4a119	cumulative	B		None
ip.floating.transmit.packets	a726f93a-65fa-4cad-828b-54dbfcf4a119	cumulative	packet		None
network	7fa6796b-756e-4320-9e73-87d4c52ecc83	gauge	network	15c0240142084d16b3127d6f844adb9	
network	9408e287-d3e7-41e2-89f0-5c691c9ca450	gauge	network	15c0240142084d16b3127d6f844adb9	
network	b3b72b98-f61e-4e1f-9a9b-84f4f3ddec0b	gauge	network	15c0240142084d16b3127d6f844adb9	
network	cb829abd-e6a3-42e9-a82f-0742db55d329	gauge	network	15c0240142084d16b3127d6f844adb9	
network.create	7fa6796b-756e-4320-9e73-87d4c52ecc83	delta	network	15c0240142084d16b3127d6f844adb9	
network.create	9408e287-d3e7-41e2-89f0-5c691c9ca450	delta	network	15c0240142084d16b3127d6f844adb9	
network.create	b3b72b98-f61e-4e1f-9a9b-84f4f3ddec0b	delta	network	15c0240142084d16b3127d6f844adb9	
network.create	cb829abd-e6a3-42e9-a82f-0742db55d329	delta	network	15c0240142084d16b3127d6f844adb9	
port	0d401d96-c2bf-4672-abf2-880eecf25ceb	gauge	port	01edcedd989f43b3a2d6121d424b254d	
port	211b94a4-581d-45d0-8710-c6c69df15709	gauge	port	01edcedd989f43b3a2d6121d424b254d	
port	2287ce25-4eef-4212-b77f-3cf590943d36	gauge	port	01edcedd989f43b3a2d6121d424b254d	
port.create	f62f3732-222e-4c40-8783-5bcbcf1fd6a1c	delta	port	01edcedd989f43b3a2d6121d424b254d	
port.create	f8c89218-3cad-48e2-8bd8-46c1bc33e752	delta	port	01edcedd989f43b3a2d6121d424b254d	

```

| port.update | delta | port |
43ed422d-b073-489f-877f-515a3cc0b8c4 | 15c0240142084d16b3127d6f844adb9 |
ded208991de34fe4bb7dd725097f1c7e |
| subnet | gauge | subnet |
09105ed1-1654-4b5f-8c12-f0f2666fa304 | 15c0240142084d16b3127d6f844adb9 |
ded208991de34fe4bb7dd725097f1c7e |
| subnet | gauge | subnet |
4bf00aac-407c-4266-a048-6ff52721ad82 | 15c0240142084d16b3127d6f844adb9 |
ded208991de34fe4bb7dd725097f1c7e |
| subnet.create | delta | subnet |
09105ed1-1654-4b5f-8c12-f0f2666fa304 | 15c0240142084d16b3127d6f844adb9 |
ded208991de34fe4bb7dd725097f1c7e |
| subnet.create | delta | subnet |
4bf00aac-407c-4266-a048-6ff52721ad82 | 15c0240142084d16b3127d6f844adb9 |
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	User ID
Project 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		
451c93eb-e728-4ba1-8665-6e7c7a8b49e2	None			None
ip.floating.transmit.packets	cumulative	packet		
9cf76844-8f09-4518-a09e-e2b8832bf894	None			None

In the meter -list output, the Resource ID refers to the floating IP.

The following example shows the output from the **ceilometer resource-show -r 451c93eb-e728-4ba1-8665-6e7c7a8b49e2** command:

Property	Value
metadata	{u'router_id': u'None', u'status': u'ACTIVE', u'tenant_id': u'ceed483222f9453ab1d7bcdd353971bc', u'floating_network_id': u'6d0cca50-4be4-4b49-856a-6848133eb970', u'fixed_ip_address': u'2.2.2.4', u'floating_ip_address': u'3.3.3.4', u'port_id': ad98-4e56-ae65-ab7c62a67355', u'id': u'451c93eb-e728-4ba1-8665-6e7c7a8b49e2', u'device_id': u'00953f62-df11-4b05-97ca-30c3f6735ffd'}
project_id	None
resource_id	451c93eb-e728-4ba1-8665-6e7c7a8b49e2
source	openstack

```
| user_id      | None
|
+-----+
```

The following example shows the output from the **ceilometer statistics** command and the **ceilometer sample-list** command for the **ip.floating.receive.packets** meter:

```
+-----+-----+-----+-----+-----+-----+-----+-----+
| 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 |
+-----+-----+-----+-----+-----+-----+-----+
```

```
+-----+-----+-----+-----+-----+-----+-----+-----+
| Resource ID | Name | Type |
| Volume | Unit | Timestamp |
+-----+-----+-----+-----+-----+-----+-----+
| 9cf76844-8f09-4518-a09e-e2b8832bf894 | ip.floating.receive.packets |
cumulative | 208.0 | packet | 2015-02-18T21:48:30.469000 |
| 451c93eb-e728-4ba1-8665-6e7c7a8b49e2 | ip.floating.receive.packets |
cumulative | 325.0 | packet | 2015-02-18T21:48:28.354000 |
| 9cf76844-8f09-4518-a09e-e2b8832bf894 | ip.floating.receive.packets |
cumulative | 0.0 | packet | 2015-02-18T21:38:30.350000 |
```

## 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 automatically as part of the OpenStack controller node and the compute agent service is installed as part of the compute node.

The following fabric tasks are added to facilitate the installation and provisioning:

**fab install\_ceilometer**—Installs the Ceilometer packages on the OpenStack controller node.

**fab install\_ceilometer\_compute**—Installs the Ceilometer packages on the compute node.

**fab setup\_ceilometer**—Provisions the Ceilometer controller services on the OpenStack controller node.

**fab setup\_ceilometer\_compute**—Provisions the Contrail Ceilometer plugin package on the OpenStack controller node.

**fab install\_contrail\_ceilometer\_plugin**—Installs the Contrail Ceilometer plugin package on the OpenStack controller node.

**fab setup\_contrail\_ceilometer\_plugin**—Provisions the Contrail Ceilometer plugin package on the OpenStack controller node.



**NOTE:** The fabric tasks are automatically called as part of the **fab install\_openstack** and **fab setup\_openstack** commands for the OpenStack controller node, and as part of the **fab install\_vrouter**, **fab setup\_vrouter** commands for the compute node





## CHAPTER 14

# Using Contrail Analytics to Monitor and Troubleshoot the Network

- [Monitoring the System on page 280](#)
- [Debugging Processes Using the Contrail Introspect Feature on page 282](#)
- [Monitor > Infrastructure > Dashboard on page 286](#)
- [Monitor > Infrastructure > Control Nodes on page 288](#)
- [Monitor > Infrastructure > Virtual Routers on page 295](#)
- [Monitor > Infrastructure > Analytics Nodes on page 306](#)
- [Monitor > Infrastructure > Config Nodes on page 311](#)
- [Monitor > Networking on page 314](#)
- [Query > Flows on page 322](#)
- [Query > Logs on page 329](#)
- [Example: Debugging Connectivity Using Monitoring for Troubleshooting on page 334](#)


## Monitoring the System

The **Monitor** icon on the Contrail Controller provides numerous options so you can view and analyze usage and other activity associated with all nodes of the system, through the use of reports, charts, and detailed lists of configurations and system activities.

Monitor pages support monitoring of infrastructure components—control nodes, virtual routers, analytics nodes, and config nodes. Additionally, users can monitor networking and debug components.

Use the menu options available from the **Monitor** icon to configure and view the statistics you need for better understanding of the activities in your system. See [Figure 98 on page 280](#)

Figure 98: Monitor Menu



See [Table 32 on page 280](#) for descriptions of the items available under each of the menu options from the **Monitor** icon.

Table 32: Monitor Menu Options

Option	Description
Infrastructure > Dashboard	Shows “at-a-glance” status view of the infrastructure components, including the numbers of virtual routers, control nodes, analytics nodes, and config nodes currently operational, and a bubble chart of virtual routers showing the CPU and memory utilization, log messages, system information, and alerts. See <a href="#">“Monitor &gt; Infrastructure &gt; Dashboard” on page 286</a> .

Table 32: Monitor Menu Options (*continued*)

Option	Description
<b>Infrastructure &gt; Control Nodes</b>	<p>View a summary for all control nodes in the system, and for each control node, view:</p> <ul style="list-style-type: none"> <li>Graphical reports of memory usage and average CPU load.</li> <li>Console information for a specified time period.</li> <li>A list of all peers with details about type, ASN, and the like.</li> <li>A list of all routes, including next hop, source, local preference, and the like.</li> </ul> <p>See <a href="#">"Monitor &gt; Infrastructure &gt; Control Nodes"</a> on page 288.</p>
<b>Infrastructure &gt; Virtual Routers</b>	<p>View a summary of all vRouters in the system, and for each vRouter, view:</p> <ul style="list-style-type: none"> <li>Graphical reports of memory usage and average CPU load.</li> <li>Console information for a specified time period.</li> <li>A list of all interfaces with details such as label, status, associated network, IP address, and the like.</li> <li>A list of all associated networks with their ACLs and VRFs.</li> <li>A list of all active flows with source and destination details, size, and time.</li> </ul> <p>See <a href="#">"Monitor &gt; Infrastructure &gt; Virtual Routers"</a> on page 295.</p>
<b>Infrastructure &gt; Analytics Nodes</b>	<p>View activity for the analytics nodes, including memory and CPU usage, analytics host names, IP address, status, and more. See <a href="#">"Monitor &gt; Infrastructure &gt; Analytics Nodes"</a> on page 306.</p>
<b>Infrastructure &gt; Config Nodes</b>	<p>View activity for the config nodes, including memory and CPU usage, config host names, IP address, status, and more. See <a href="#">"Monitor &gt; Infrastructure &gt; Config Nodes"</a> on page 311.</p>
<b>Networking &gt; Networks</b>	<p>For all virtual networks for all projects in the system, view graphical traffic statistics, including:</p> <ul style="list-style-type: none"> <li>Total traffic in and out.</li> <li>Inter VN traffic in and out.</li> <li>The most active ports, peers, and flows for a specified duration.</li> <li>All traffic ingress and egress from connected networks, including their attached policies.</li> </ul> <p>See <a href="#">"Monitor &gt; Networking"</a> on page 314.</p>
<b>Networking &gt; Dashboard</b>	<p>For all virtual networks for all projects in the system, view graphical traffic statistics, including:</p> <ul style="list-style-type: none"> <li>Total traffic in and out.</li> <li>Inter VN traffic in and out.</li> </ul> <p>You can view the statistics in varying levels of granularity, for example, for a whole project, or for a single network. See <a href="#">"Monitor &gt; Networking"</a> on page 314.</p>
<b>Networking &gt; Projects</b>	<p>View essential information about projects in the system including name, associated networks, and traffic in and out.</p>

Table 32: Monitor Menu Options (*continued*)

Option	Description
<b>Networking &gt; Networks</b>	View essential information about networks in the system including name and traffic in and out.
<b>Networking &gt; Instances</b>	View essential information about instances in the system including name, associated networks, interfaces, vRouters, and traffic in and out.
<b>Debug &gt; Packet Capture</b>	<ul style="list-style-type: none"> <li>• Add and manage packet analyzers.</li> <li>• Attach packet captures and configure their details.</li> <li>• View a list of all packet analyzers in the system and the details of their configurations, including source and destination networks, ports, and IP addresses.</li> </ul>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Monitor &gt; Infrastructure &gt; Dashboard on page 286</a></li> <li>• <a href="#">Monitor &gt; Infrastructure &gt; Control Nodes on page 288</a></li> <li>• <a href="#">Monitor &gt; Infrastructure &gt; Virtual Routers on page 295</a></li> <li>• <a href="#">Monitor &gt; Networking on page 314</a></li> <li>• <a href="#">Query &gt; Logs on page 329</a></li> <li>• <a href="#">Query &gt; Flows on page 322</a></li> </ul>

## Debugging Processes Using the Contrail Introspect Feature

This topic describes how to use the Sandesh infrastructure and the Contrail Introspect feature to debug processes.

Introspect is a mechanism for taking a program object and querying information about it.

Sandesh is the name of a unified infrastructure in the Contrail Virtual Networking solution.

Sandesh is a way for the Contrail daemons to provide a request-response mechanism. Requests and responses are defined in Sandesh format and the Sandesh compiler generates code to process the requests and send responses.

Sandesh also provides a way to use a Web browser to send Sandesh requests to a Contrail daemon and get the Sandesh responses. This feature is used to debug processes by looking into the operational status of the daemons.

Each Contrail daemon starts an HTTP server, with the following page types:

- The main index.html listing all Sandesh modules and the links to them.
- Sandesh module pages that present HTML forms for each Sandesh request.

- XML-based dynamically-generated pages that display Sandesh responses.
- An automatically generated page that shows all code needed for rendering and all HTTP server-client interactions.


You can display the HTTP introspect of a Contrail daemon directly by accessing the following Introspect ports:

- `<controller-ip>:8083`. This port displays the *contrail-control* introspect port.
- `<compute-ip>:8085`. This port displays the *contrail-vrouter-agent* introspect port.

Another way to launch the Introspect page is by browsing to a particular node page using the Contrail Web user interface.

Figure 99 on page 283 shows the *contrail-control* infrastructure page. Notice the Introspect link at the bottom of the Control Nodes Details tab window.

**Figure 99: Control Nodes Details Tab Window**



The following are the Sandesh modules for the Contrail control process (*contrail-control*) Introspect port.

- `bgp_peer.xml`
- `control_node.xml`
- `cpuinfo.xml`

- discovery\_client\_stats.xml
- ifmap\_log.xml
- ifmap\_server\_show.xml
- rtarget\_group.xml
- sandesh\_trace.xml
- sandesh\_uve.xml
- service\_chaining.xml
- static\_route.xml
- task.xml
- xmpp\_server.xml

Figure 100 on page 284 shows the Controller Introspect window.

Figure 100: Controller Introspect Window

The screenshot shows the 'Controller Introspect' window in the Contrail interface. The title bar includes 'Contrail' and navigation buttons: 'Collapse', 'Expand', 'Wrap', and 'NoWrap'. The main content area has a header 'Controller Introspect'. Below it, there is a text input field containing 'AgentXmppConnectionStatusReq'. To the right of this field is a 'Send' button. The bottom right corner of the window displays the IP address '10.2.3.88'.

Figure 101 on page 284 shows an example of the BGP Peer (bgp\_peer.xml) Introspect page.

Figure 101: BGP Peer Introspect Page

The screenshot shows the 'Bgp\_peer Introspect' page in the Contrail interface. The title bar includes 'Contrail' and navigation buttons: 'Collapse', 'Expand', 'Wrap', and 'NoWrap'. The main content area has a header 'Bgp\_peer Introspect'. On the left side, there is a sidebar with a list of introspectable components: 'BgpNeighborReq', 'ShowBgpNeighborSummaryReq', 'ClearBgpNeighborReq', 'ShowRouteReq', 'ShowRouteSummaryReq', 'ShowRoutingInstanceReq', 'ShowRoutingInstanceSummaryReq', 'ShowMulticastManagerReq', 'ShowBgpInstanceConfigReq', 'ShowBgpNeighborConfigReq', and 'ShowBgpServerReq'. The main area contains three sections: 1. 'BgpNeighborReq' with input fields for 'neighbor(string)' and 'domain(string)', and a 'Send' button. 2. 'ShowBgpNeighborSummaryReq' with a 'Send' button. 3. 'ClearBgpNeighborReq' with an input field for 'name(string)' and a 'Send' button. The bottom right corner of the window displays the IP address '10.2.3.88'.

Figure 102 on page 285 shows an example of the BGP Neighbor Summary Introspect page.

Figure 102: BGP Neighbor Summary Introspect Page

Contrail									
ShowBgpNeighborSummaryResp									
neighbors									
peer	deleted	deleted_at	peer_address	peer_id	peer_asn	encoding	peer_type	state	local_address
b6s23	false	-	192.168.68.1	192.168.68.1	64512	BGP	internal	Established	192.168.68.2
b6s25	false	-	192.168.68.3	192.168.68.3	64512	BGP	internal	Established	192.168.68.2
mx1	false	-	192.168.100.1	192.168.100.1	64512	BGP	internal	Established	192.168.68.2
mx2	false	-	192.168.100.2	192.168.100.2	64512	BGP	internal	Established	192.168.68.2
b6s28	false	-	192.168.68.6	-	0	XMPP	internal	Established	192.168.68.2
b6s18	false	-	192.168.69.5	-	0	XMPP	internal	Established	192.168.68.2
b6s13	false	-	192.168.69.8	-	0	XMPP	internal	Established	192.168.68.2
b6s7	false	-	192.168.69.11	-	0	XMPP	internal	Established	192.168.68.2
b6s33	false	-	192.168.68.11	-	0	XMPP	internal	Established	192.168.68.2
b6s9	false	-	192.168.69.10	-	0	XMPP	internal	Established	192.168.68.2
b6s26	false	-	192.168.68.4	-	0	XMPP	internal	Established	192.168.68.2

The following are the Sandesh modules for the Contrail vRouter agent (**contrail-vrouter-agent**) Introspect port.

- agent.xml
- agent\_stats\_interval.xml
- cfg.xml
- controller.xml
- cpuinfo.xml
- diag.xml
- discovery\_client\_stats.xml
- flow\_stats\_interval.xml
- ifmap\_agent.xml
- kstate.xml
- multicast.xml
- pkt.xml
- port\_ipc.xml
- sandesh\_trace.xml
- sandesh\_uve.xml
- services.xml
- stats\_interval.xml
- task.xml
- xmpp\_server.xml

Figure 103 on page 286 shows an example of the Agent (agent.xml) Introspect page.

Figure 103: Agent Introspect Page

Contrail									
AgentXmppConnectionStatus									
peer									
controller_ip	state	cfg_controller	mcast_controller	last_state	last_event	last_state_at	flap_count	flap_time	rx
192.168.68.3	Established	Yes	No	OpenSent	xmsm::EvXmppKeepAlive	2015-Jul-21 01:20:57.616019	2	2015-Jul-21 01:20:57.555077	rx
192.168.68.2	Established	No	Yes	OpenSent	xmsm::EvXmppKeepAlive	2015-Jul-21 01:20:59.599875	2	2015-Jul-21 01:20:59.548692	rx

## Monitor > Infrastructure > Dashboard


Use **Monitor > Infrastructure > Dashboard** to get an “at-a-glance” view of the system infrastructure components, including the numbers of virtual routers, control nodes, analytics nodes, and config nodes currently operational, a bubble chart of virtual routers showing the CPU and memory utilization, log messages, system information, and alerts.

- [Monitor Dashboard on page 286](#)
- [Monitor Individual Details from the Dashboard on page 287](#)
- [Using Bubble Charts on page 287](#)
- [Color-Coding of Bubble Charts on page 288](#)

## Monitor Dashboard

Click **Monitor > Infrastructure > Dashboard** on the left to view the **Dashboard**. See [Figure 104 on page 286](#).

Figure 104: Monitor > Infrastructure > Dashboard






## Monitor Individual Details from the Dashboard

Across the top of the **Dashboard** screen are summary boxes representing the components of the system that are shown in the statistics. See [Figure 105 on page 287](#). Any of the control nodes, virtual routers, analytics nodes, and config nodes can be monitored individually and in detail from the **Dashboard** by clicking an associated box, and drilling down for more detail.

**Figure 105: Dashboard Summary Boxes**



Detailed information about monitoring each of the areas represented by the boxes is provided in the links in [Table 33 on page 287](#).


**Table 33: Dashboard Summary Boxes**

Box	For More Information
vRouters	"Monitor > Infrastructure > Virtual Routers" on page 295
Control Nodes	"Monitor > Infrastructure > Control Nodes" on page 288
Analytics Nodes	"Monitor > Infrastructure > Analytics Nodes" on page 306
Config Nodes	"Monitor > Infrastructure > Config Nodes" on page 311

## Using Bubble Charts

Bubble charts show the CPU and memory utilization of components contributing to the current analytics display, including vRouters, control nodes, config nodes, and the like. You can hover over any bubble to get summary information about the component it represents; see [Figure 106 on page 288](#). You can click through the summary information to get more details about the component.

Figure 106: Bubble Summary Information



### Color-Coding of Bubble Charts

Bubble charts use the following color-coding scheme:

#### Control Nodes

- Blue—working as configured.
- Red—error, at least one configured peer is down.

#### vRouters

- Blue—working, but no instance is launched.
- Green—working with at least one instance launched.
- Red—error, there is a problem with connectivity or a vRouter is in a failed state.

#### Related Documentation

- [Monitor > Infrastructure > Virtual Routers on page 295](#)
- [Monitor > Infrastructure > Control Nodes on page 288](#)
- [Monitor > Infrastructure > Analytics Nodes on page 306](#)
- [Monitor > Infrastructure > Config Nodes on page 311](#)

## Monitor > Infrastructure > Control Nodes

Use **Monitor > Infrastructure > Control Nodes** to gain insight into usage statistics for control nodes.

- [Monitor Control Nodes Summary on page 289](#)
- [Monitor Individual Control Node Details on page 289](#)
- [Monitor Individual Control Node Console on page 291](#)
- [Monitor Individual Control Node Peers on page 293](#)
- [Monitor Individual Control Node Routes on page 294](#)

## Monitor Control Nodes Summary

Select **Monitor > Infrastructure > Control Nodes** to see a graphical chart of average memory usage versus average CPU percentage usage for all control nodes in the system. Also on this screen is a list of all control nodes in the system. See [Figure 107 on page 289](#). See [Table 34 on page 289](#) for descriptions of the fields on this screen.

Figure 107: Control Nodes Summary



Table 34: Control Nodes Summary Fields


Field	Description
Host name	The name of the control node.
IP Address	The IP address of the control node.
Version	The software version number that is installed on the control node.
Status	The current operational status of the control node — Up or Down.
CPU (%)	The CPU percentage currently in use by the selected control node.
Memory	The memory in MB currently in use and the total memory available for this control node.
Total Peers	The total number of peers for this control node.
Established in Sync Peers	The total number of peers in sync for this control node.
Established in Sync vRouters	The total number of vRouters in sync for this control node.

## Monitor Individual Control Node Details

Click the name of any control nodes listed under the **Control Nodes** title to view an array of graphical reports of usage and numerous details about that node. There are several

tabs available to help you probe into more details about the selected control node. The first tab is the **Details** tab; see [Figure 108 on page 290](#).

**Figure 108: Individual Control Node—Details Tab**



The Details tab provides a summary of the status and activity on the selected node, and presents graphical displays of CPU and memory usage. See [Table 35 on page 290](#) for descriptions of the fields on this tab.

**Table 35: Individual Control Node—Details Tab Fields**

Field	Description
Hostname	The host name defined for this control node.
IP Address	The IP address of the selected node.
Status	The operational status of the control node.
Control Node Manager	The operational status of the control node manager.
Config Node	The IP address of the configuration node associated with this control node.
Analytics Node	The IP address of the node from which analytics (monitor) information is derived.
Analytics Messages	The total number of analytics messages in and out from this node.
Peers	The total number of peers established for this control node and how many are in sync and of what type.

Table 35: Individual Control Node—Details Tab Fields (*continued*)

Field	Description
CPU	The average percent of CPU load incurred by this control node.
Memory	The average memory usage incurred by this control node.
Last Log	The date and time of the last log message issued about this control node.
Control Node CPU/Memory Utilization	A graphic display x, y chart of the average CPU load and memory usage incurred by this control node over time.

## Monitor Individual Control Node Console

Click the **Console** tab for an individual control node to display system logging information for a defined time period, with the last 5 minutes of information as the default display. See [Figure 109 on page 291](#).

Figure 109: Individual Control Node—Console Tab

Monitor > Infrastructure > Control Nodes > b3s9

Search Sitemap

Details Console Peers Routes

### Console Logs

Time Range: Custom

From Time: Oct 08, 2013 02:26:33 PM

To Time: Oct 08, 2013 02:31:33 PM

Log Category: All

Log Type: any

Log Level: SYS\_DEBUG

Limit: Limit 10 mess

Auto Refresh: ☒

Display Logs Reset

Time	Category	Log Type	Log
2013-10-08 14:31:30:351:353	BGP	BgpStateMachineSessionMessageLog	Bgp Peer 10.84.23.252 : P fsm::EvConnectTimerExp
2013-10-08 14:31:27:971:482	BGP	BgpStateMachineSessionMessageLog	Bgp Peer 10.84.23.253 : P state Connect
2013-10-08 14:31:24:970:157	BGP	BgpStateMachineSessionMessageLog	Bgp Peer 10.84.23.253 : P fsm::EvConnectTimerExp
2013-10-08 14:30:58:220:866	BGP	BgpStateMachineSessionMessageLog	Bgp Peer 10.84.23.252 : P state Connect

See [Table 36 on page 291](#) for descriptions of the fields on the **Console** tab screen.

Table 36: Control Node: Console Tab Fields

Field	Description
Time Range	Select a timeframe for which to review logging information as sent to the console. There are 11 options, ranging from the <b>Last 5 mins</b> through to the <b>Last 24 hrs</b> . The default display is for the <b>Last 5 mins</b> .

Table 36: Control Node: Console Tab Fields (*continued*)

Field	Description
<b>Log Category</b>	Select a log category to display:  All _default_ XMPP IFMap TCP
<b>Log Type</b>	Select a log type to display.
<b>Log Level</b>	Select a log severity level to display:  SYS_EMERG SYS_ALERT SYS_CRIT SYS_ERR SYS_WARN SYS_NOTICE SYS_INFO SYS_DEBUG
<b>Search</b>	Enter any text string to search and display logs containing that string.
<b>Limit</b>	Select from a list an amount to limit the number of messages displayed:  No Limit Limit 10 messages Limit 50 messages Limit 100 messages Limit 200 messages Limit 500 messages
<b>Auto Refresh</b>	Click the check box to automatically refresh the display if more messages occur.
<b>Display Logs</b>	Click this button to refresh the display if you change the display criteria.
<b>Reset</b>	Click this button to clear any selected display criteria and reset all criteria to their default settings.
<b>Time</b>	This column lists the time received for each log message displayed.
<b>Category</b>	This column lists the log category for each log message displayed.
<b>Log Type</b>	This column lists the log type for each log message displayed.
<b>Log</b>	This column lists the log message for each log displayed.

## Monitor Individual Control Node Peers

The **Peers** tab displays the peers for an individual control node and their peering state. Click the expansion arrow next to the address of any peer to reveal more details. See [Figure 110 on page 293](#).

Figure 110: Individual Control Node—Peers Tab

Monitor > Infrastructure > Control Nodes > b3s9

Search Sitemap

Details Console **Peers** Routes

**Peers**

Peer	Peer Type	Peer ASN	Status	Last flap	Messages (Recv/Sent)
▶ 10.84.23.252	BGP	64512	Active, -	-	0/ 0
▶ 10.84.23.8	BGP	64512	Established, in sync	-	3754/ 3758
▶ 10.84.23.253	BGP	64512	Connect, -	-	0/ 0
▶ 10.84.21.4	XMPP	-	Established, in sync	-	2751/ 5189
▶ 10.84.21.5	XMPP	-	Established, in sync	-	2753/ 5802
▶ 10.84.21.6	XMPP	-	Established, in sync	-	2752/ 4264
▲ 10.84.21.34	XMPP	-	Established, in sync	-	2753/ 5659

Details:

```

- {
  name: "b3s9:10.84.21.34",
  value: - {
    XmppPeerInfoData: - {
      state_info: - {
        last_state: "Active",
        state: "Established",
        last_state_at: 1381198447915913
      },
      peer_stats_info: - {

```

s041579

See [Table 37 on page 293](#) for descriptions of the fields on the **Peers** tab screen.

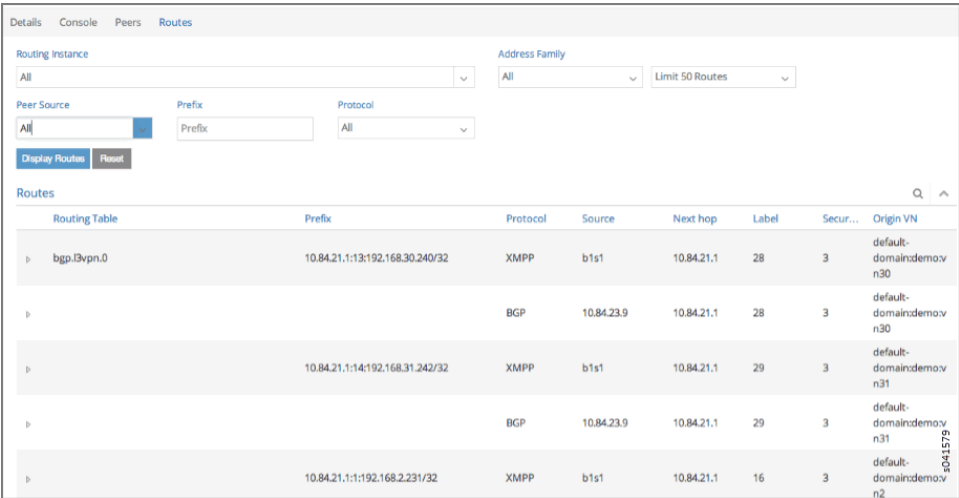
Table 37: Control Node: Peers Tab Fields

Field	Description
Peer	The hostname of the peer.
Peer Type	The type of peer.
Peer ASN	The autonomous system number of the peer.
Status	The current status of the peer.
Last flap	The last flap detected for this peer.
Messages (Recv/Sent)	The number of messages sent and received from this peer.

Monitor Individual Control Node Routes

The **Routes** tab displays active routes for this control node and lets you query the results. Use horizontal and vertical scroll bars to view more results. Click the expansion icon next to a routing table name to reveal more details about the selected route. See [Figure 111 on page 294](#).

Figure 111: Individual Control Node—Routes Tab



See [Table 38 on page 294](#) for descriptions of the fields on the **Routes** tab screen.

Table 38: Control Node: Routes Tab Fields

Field	Description
Routing Instance	You can select a single routing instance from a list of all instances for which to display the active routes.
Address Family	Select an address family for which to display the active routes:  All (default) l3vpn inet inetmcast
(Limit Field)	Select to limit the display of active routes:  Limit 10 Routes Limit 50 Routes Limit 100 Routes Limit 200 Routes
Peer Source	Select from a list of available peers the peer for which to display the active routes, or select All.



Table 38: Control Node: Routes Tab Fields (*continued*)

Field	Description
<b>Prefix</b>	Enter a route prefix to limit the display of active routes to only those with the designated prefix.
<b>Protocol</b>	Select a protocol for which to display the active routes:  All (default) XMPP BGP ServiceChain Static
<b>Display Routes</b>	Click this button to refresh the display of routes after selecting different display criteria.
<b>Reset</b>	Click this button to clear any selected criteria and return the display to default values.
<i>Column</i>	<i>Description</i>
<b>Routing Table</b>	The name of the routing table that stores this route.
<b>Prefix</b>	The route prefix for each active route displayed.
<b>Protocol</b>	The protocol used by the route.
<b>Source</b>	The host source for each active route displayed.
<b>Next hop</b>	The IP address of the next hop for each active route displayed.
<b>Label</b>	The label for each active route displayed.
<b>Security</b>	The security value for each active route displayed.
<b>Origin VN</b>	The virtual network from which the route originates.
<b>AS Path</b>	The AS path for each active route displayed.

## Monitor > Infrastructure > Virtual Routers


- [Monitor vRouters Summary on page 296](#)
- [Monitor Individual vRouters Tabs on page 297](#)
- [Monitor Individual vRouter Details Tab on page 297](#)
- [Monitor Individual vRouters Interfaces Tab on page 298](#)
- [Configuring Interface Monitoring and Mirroring on page 299](#)
- [Monitor Individual vRouters Networks Tab on page 300](#)
- [Monitor Individual vRouters ACL Tab on page 301](#)

- [Monitor Individual vRouters Flows Tab on page 302](#)
- [Monitor Individual vRouters Routes Tab on page 303](#)
- [Monitor Individual vRouter Console Tab on page 304](#)

Monitor vRouters Summary

Click **Monitor > Infrastructure > Virtual Routers** to view the **vRouters** summary screen.  
See [Figure 112 on page 296](#).

Figure 112: vRouters Summary



See [Table 39 on page 296](#) for descriptions of the fields on the **vRouters Summary** screen.

Table 39: vRouters Summary Fields

Field	Description
Host name	The name of the vRouter. Click the name of any vRouter to reveal more details.
IP Address	The IP address of the vRouter.
Version	The version of software installed on the system.
Status	The current operational status of the vRouter — Up or Down.
CPU (%)	The CPU percentage currently in use by the selected vRouter.
Memory (MB)	The memory currently in use and the total memory available for this vRouter.
Networks	The total number of networks for this vRouter.

Table 39: vRouters Summary Fields (*continued*)

Field	Description
<b>Instances</b>	The total number of instances for this vRouter.
<b>Interfaces</b>	The total number of interfaces for this vRouter.

## Monitor Individual vRouters Tabs

Click the name of any vRouter to view details about performance and activities for that vRouter. Each individual vRouters screen has the following tabs.

- **Details**—similar display of information as on individual control nodes **Details** tab. See [Figure 113 on page 297](#).
- **Console**—similar display of information as on individual control nodes **Console** tab. See [Figure 121 on page 305](#).
- **Interfaces**—details about associated interfaces. See [Figure 114 on page 299](#).
- **Networks**—details about associated networks. See [Figure 117 on page 301](#).
- **ACL**—details about access control lists. See [Figure 118 on page 302](#).
- **Flows**—details about associated traffic flows. See [Figure 119 on page 303](#).
- **Routes**—details about associated routes. See [Figure 120 on page 304](#).

## Monitor Individual vRouter Details Tab

The **Details** tab provides a summary of the status and activity on the selected node, and presents graphical displays of CPU and memory usage; see [Figure 113 on page 297](#). See [Table 40 on page 298](#) for descriptions of the fields on this tab.

Figure 113: Individual vRouters—Details Tab

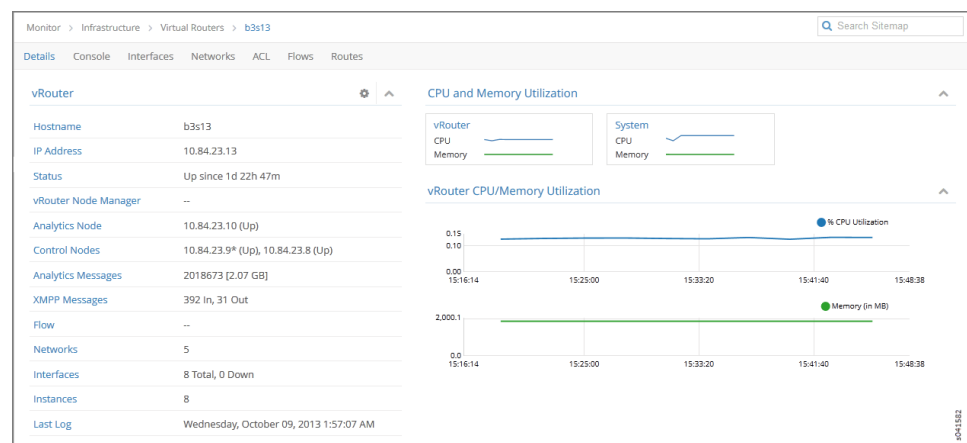


Table 40: vRouters Details Tab Fields

Field	Description
Hostname	The hostname of the vRouter.
IP Address	The IP address of the selected vRouter.
Status	The operational status of the vRouter.
vRouter Node Manager	The operational status of the vRouter node manager.
Analytics Node	The IP address of the node from which analytics (monitor) information is derived.
Control Nodes	The IP address of the configuration node associated with this vRouter.
Analytics Messages	The total number of analytics messages in and out from this node.
XMPP Messages	The total number of XMPP messages that have gone in and out of this vRouter.
Flow	The number of active flows and the total flows for this vRouter.
Networks	The number of networks associated with this vRouter.
Interfaces	The number of interfaces associated with this vRouter.
Instances	The number of instances associated with this vRouter.
Last Log	The date and time of the last log message issued about this vRouter.
vRouter CPU/Memory Utilization	Graphs (x, y) displaying CPU and memory utilization averages over time for this vRouter, in comparison to system utilization averages.

### Monitor Individual vRouters Interfaces Tab

The **Interfaces** tab displays details about the interfaces associated with an individual vRouter. Click the expansion arrow next to any interface name to reveal more details. Use horizontal and vertical scroll bars to access all portions of the screen. See [Figure 114 on page 299](#). See [Table 41 on page 299](#) for descriptions of the fields on the **Interfaces** tab screen.

Figure 114: Individual vRouters—Interfaces Tab

Monitor > Infrastructure > Virtual Routers > b1s36

Search Sitemap

Details Console **Interfaces** Networks ACL Flows Routes

**Interfaces**

Name	Label	Status	Network	IP Address	Floating IP	Instance
tap25e5cee3-07	18	Up	default-domain:demo:vn30	192.168.30.247	None	005132fd-0d83-4db7-88c8-bd49d68e9480
tap4d91aab1-f1	25	Up	default-domain:demo:vn26	192.168.26.247	None	65d6c6e9-7a82-43d8-a706-f74d81715920
tap5a8cd9dd-5b	27	Up	default-domain:demo:vn23	192.168.23.249	None	a159c518-4fb6-402a-ae0d-eb5b4457b551
tap603a5e0b-8b	16	Up	default-domain:demo:vn19	192.168.19.247	None	fe622580-b0cf-4c6d-89e5-d2065e7e87e4
tap68ad232c-76	19	Up	default-domain:demo:vn28	192.168.28.247	None	91089d89-76b5-46c2-abc9-b9693bcb37ac

Details :

```
{
  index: "6",
  name: "tap68ad232c-76",
  uuid: "68ad232c-76d1-4fe2-a200-42182497545e",
  vrf_name: "default-domain:demo:vn28:vn28",
  active: "Active",
  dhcp_service: "Enable",
}
```

Table 41: vRouters: Interfaces Tab Fields

Field	Description
<b>Name</b>	The name of the interface.
<b>Label</b>	The label for the interface.
<b>Status</b>	The current status of the interface.
<b>Network</b>	The network associated with the interface.
<b>IP Address</b>	The IP address of the interface.
<b>Floating IP</b>	Displays any floating IP addresses associated with the interface.
<b>Instance</b>	The name of any instance associated with the interface.

## Configuring Interface Monitoring and Mirroring


Contrail supports user monitoring of traffic on any guest virtual machine interface when using the Juniper Contrail user interface.

When interface monitoring (packet capture) is selected, a default analyzer is created and all traffic from the selected interface is mirrored and sent to the default analyzer. If a mirroring instance is already launched, the traffic will be redirected to the selected instance. The interface traffic is only mirrored during the time that the monitor packet capture interface is in use. When the capture screen is closed, interface mirroring stops.

To configure interface mirroring:

1. Select **Monitor > Infrastructure > Virtual Routers**, then select the vRouter that has the interface to mirror.
2. In the list of attributes for the vRouter, select **Interfaces**; see [Figure 115 on page 300](#).

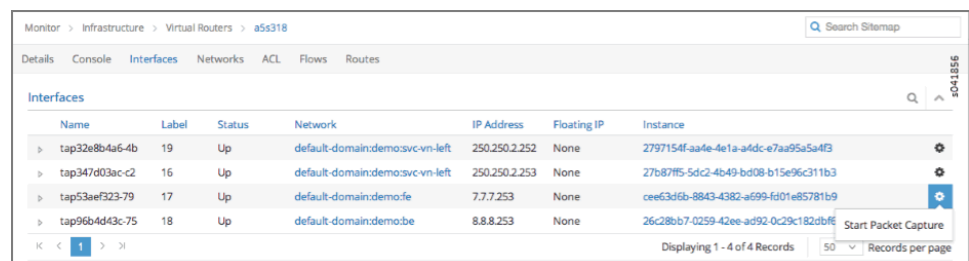
**Figure 115: Individual vRouter**



A list of interfaces for that vRouter appears.

3. For the interface to mirror, click the Action icon in the last column and select the option **Packet Capture**; see [Figure 116 on page 300](#).

**Figure 116: Interfaces**



The mirror packet capture starts and displays at this screen.

The mirror packet capture stops when you exit this screen.

## Monitor Individual vRouters Networks Tab

The **Networks** tab displays details about the networks associated with an individual vRouter. Click the expansion arrow at the name of any network to reveal more details. See [Figure 117 on page 301](#). See [Table 42 on page 301](#) for descriptions of the fields on the **Networks** tab screen.

Figure 117: Individual vRouters—Networks Tab

Name	ACLs	VRF
default-domain:demo:vn24	a372751f-6497-41e9-b409-fa4ab5ce6b7f	default-domain:demo:vn24:vn24
default-domain:demo:vn22	195af177-0a28-49a1-9cf0-2ceac22af5a1	default-domain:demo:vn22:vn22
default-domain:demo:vn30	362cce6e-2894-42d6-ba03-3ee98cac8809	default-domain:demo:vn30:vn30
default-domain:demo:vn21	5918a068-1cd5-4993-9cff-386a807940ca	default-domain:demo:vn21:vn21
default-domain:demo:vn28	dd87c461-97c0-4d47-bff0-89040e7d6ab0	default-domain:demo:vn28:vn28
default-domain:demo:vn19	f0465432-6fc0-4fb3-967c-392100617408	default-domain:demo:vn19:vn19
default-domain:demo:vn2	1c46e7e0-f799-4bc6-ae09-e4654c263aa6	default-domain:demo:vn2:vn2

Details:

```

- {
  name: "default-domain:demo:vn2",
  uuid: "63d08f7a-b342-4892-9171-edab9f4c397f",
  acl_uuid: "1c46e7e0-f799-4bc6-ae09-e4654c263aa6",
  mirror_acl_uuid: - {},
  mirror_cfg_acl_uuid: - {},
  vrf_name: "default-domain:demo:vn2:vn2",
  ipam_data: - {
    list: - {

```

Table 42: vRouters: Networks Tab Fields

Field	Description
Name	The name of each network associated with this vRouter.
ACLs	The name of the access control list associated with the listed network.
VRF	The identifier of the VRF associated with the listed network.
Action	Click the icon to select the action: Edit, Delete

### Monitor Individual vRouters ACL Tab

The **ACL** tab displays details about the access control lists (ACLs) associated with an individual vRouter. Click the expansion arrow next to the UUID of any ACL to reveal more details. See [Figure 118 on page 302](#). See [Table 43 on page 302](#) for descriptions of the fields on the **ACL** tab screen.

Figure 118: Individual vRouters—ACL Tab

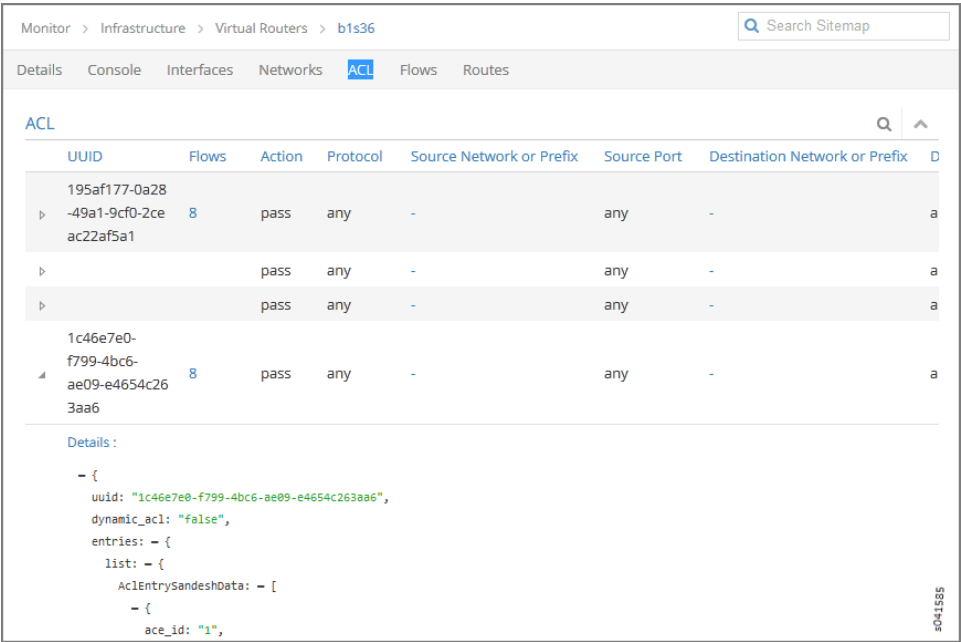


Table 43: vRouters: ACL Tab Fields

Field	Description
UUID	The universal unique identifier (UUID) associated with the listed ACL.
Flows	The flows associated with the listed ACL.
Action	The traffic action defined by the listed ACL.
Protocol	The protocol associated with the listed ACL.
Source Network or Prefix	The name or prefix of the source network associated with the listed ACL.
Source Port	The source port associated with the listed ACL.
Destination Network or Prefix	The name or prefix of the destination network associated with the listed ACL.
Destination Port	The destination port associated with the listed ACL.
ACE Id	The ACE ID associated with the listed ACL.

Monitor Individual vRouters Flows Tab

The **Flows** tab displays details about the flows associated with an individual vRouter. Click the expansion arrow next to any ACL/SG UUID to reveal more details. Use the horizontal and vertical scroll bars to access all portions of the screen. See



Figure 119 on page 303. See Table 44 on page 303 for descriptions of the fields on the **Flows** tab screen.

Figure 119: Individual vRouters—Flows Tab




Table 44: vRouters: Flows Tab Fields

Field	Description
ACL UUID	The default is to show <b>All</b> flows, however, you can select from a drop down list any single flow to view its details.
ACL / SG UUID	The universal unique identifier (UUID) associated with the listed ACL or SG.
Protocol	The protocol associated with the listed flow.
Src Network	The name of the source network associated with the listed flow.
Src IP	The source IP address associated with the listed flow.
Src Port	The source port of the listed flow.
Dest Network	The name of the destination network associated with the listed flow.
Dest IP	The destination IP address associated with the listed flow.
Dest Port	The destination port associated with the listed flow.
Bytes/Pkts	The number of bytes and packets associated with the listed flow.
Setup Time	The setup time associated with the listed flow.

Monitor Individual vRouters Routes Tab

The **Routes** tab displays details about unicast and multicast routes in specific VRFs for an individual vRouter. Click the expansion arrow next to the route prefix to reveal more details. See Figure 120 on page 304. See Table 45 on page 304 for descriptions of the fields on the **Routes** tab screen.

Figure 120: Individual vRouters—Routes Tab




Table 45: vRouters: Routes Tab Fields

Field	Description
VRF	Select from a drop down list the virtual routing and forwarding (VRF) to view.
Show Routes	Select to show the route type: <b>Unicast</b> or <b>Multicast</b> .
Prefix	The IP address prefix of a route.
Next hop	The next hop method for this route.
Next hop details	The next hop details for this route.

Monitor Individual vRouter Console Tab

Click the **Console** tab for an individual vRouter to display system logging information for a defined time period, with the last 5 minutes of information as the default display. See [Figure 121 on page 305](#). See [Table 46 on page 305](#) for descriptions of the fields on the **Console** tab screen.

Figure 121: Individual vRouter—Console Tab

Monitor > Infrastructure > Virtual Routers > b1s36

Details Console Interfaces Networks ACL Flows Routes

### Console Logs

Time Range: Custom  
 From Time: Oct 02, 2013 05:00:39 AM  
 To Time: Oct 02, 2013 05:05:39 AM  
 Log Category: All  
 Log Type: any  
 Log Level: SYS\_INFO  
 Limit: Limit 10 messages  
 Auto Refresh: ☒

Display Logs Reset

Time	Category	Log Type	Log
2013-10-02 05:05:39:572:199	Agent	AgentRouteLog	Added route 192.168.31.222/32 in VRF default-domain:demo:vn31:vn31 10.84.23.9
2013-10-02 05:05:34:761:107	Agent	AgentRouteLog	Added route 192.168.31.224/32 in VRF default-domain:demo:vn31:vn31 10.84.23.9
2013-10-02 05:05:34:731:318	Agent	AgentRouteLog	Added route 192.168.31.223/32 in VRF default-domain:demo:vn31:vn31 10.84.23.9
2013-10-02 05:05:32:283:326	Agent	AgentRouteLog	Added route 192.168.31.225/32 in VRF default-domain:demo:vn31:vn31 10.84.23.8
2013-10-02 05:05:31:282:424	Agent	AgentRouteLog	Added route 192.168.31.227/32 in VRF default-domain:demo:vn31:vn31 10.84.23.8
2013-10-02 05:05:29:319:521	Agent	AgentRouteLog	Added route 192.168.31.229/32 in VRF default-domain:demo:vn31:vn31 10.84.23.9

Table 46: Control Node: Console Tab Fields

Field	Description
<b>Time Range</b>	Select a timeframe for which to review logging information as sent to the console. There are several options, ranging from <b>Last 5 mins</b> through to the <b>Last 24 hrs</b> , plus a <b>Custom</b> time range.
<b>From Time</b>	If you select <b>Custom</b> in <b>Time Range</b> , enter the start time.
<b>To Time</b>	If you select <b>Custom</b> in <b>Time Range</b> , enter the end time.
<b>Log Category</b>	Select a log category to display: <ul style="list-style-type: none"> <li>• All</li> <li>• _default_</li> <li>• XMPP</li> <li>• IFMap</li> <li>• TCP</li> </ul>
<b>Log Type</b>	Select a log type to display.
<b>Log Level</b>	Select a log severity level to display: <ul style="list-style-type: none"> <li>• SYS_EMERG</li> <li>• SYS_ALERT</li> <li>• SYS_CRIT</li> <li>• SYS_ERR</li> <li>• SYS_WARN</li> <li>• SYS_NOTICE</li> <li>• SYS_INFO</li> <li>• SYS_DEBUG</li> </ul>

Table 46: Control Node: Console Tab Fields (*continued*)

Field	Description
<b>Limit</b>	Select from a list an amount to limit the number of messages displayed: <ul style="list-style-type: none"> <li>• No Limit</li> <li>• Limit 10 messages</li> <li>• Limit 50 messages</li> <li>• Limit 100 messages</li> <li>• Limit 200 messages</li> <li>• Limit 500 messages</li> </ul>
<b>Auto Refresh</b>	Click the check box to automatically refresh the display if more messages occur.
<b>Display Logs</b>	Click this button to refresh the display if you change the display criteria.
<b>Reset</b>	Click this button to clear any selected display criteria and reset all criteria to their default settings.
<i>Columns</i>	
<b>Time</b>	This column lists the time received for each log message displayed.
<b>Category</b>	This column lists the log category for each log message displayed.
<b>Log Type</b>	This column lists the log type for each log message displayed.
<b>Log</b>	This column lists the log message for each log displayed.

## Monitor > Infrastructure > Analytics Nodes

Select **Monitor > Infrastructure > Analytics Nodes** to view the console logs, generators, and query expansion (QE) queries of the analytics nodes.

- [Monitor Analytics Nodes on page 306](#)
- [Monitor Analytics Individual Node Details Tab on page 307](#)
- [Monitor Analytics Individual Node Generators Tab on page 308](#)
- [Monitor Analytics Individual Node QE Queries Tab on page 309](#)
- [Monitor Analytics Individual Node Console Tab on page 310](#)

## Monitor Analytics Nodes

Select **Monitor > Infrastructure > Analytics Nodes** to view a summary of activities for the analytics nodes; see [Figure 122 on page 307](#). See [Table 47 on page 307](#) for descriptions of the fields on the analytics summary.

Figure 122: Analytics Nodes Summary



Table 47: Fields on Analytics Nodes Summary

Field	Description
Host name	The name of this node.
IP address	The IP address of this node.
Version	The version of software installed on the system.
Status	The current operational status of the node — Up or Down — and the length of time it is in that state.
CPU (%)	The average CPU percentage usage for this node.
Memory	The average memory usage for this node.
Generators	The total number of generators for this node.

Monitor Analytics Individual Node Details Tab

Click the name of any analytics node displayed on the analytics summary to view the **Details** tab for that node. See [Figure 123 on page 308](#).

See [Table 48 on page 308](#) for descriptions of the fields on this screen.

Figure 123: Monitor Analytics Individual Node Details Tab




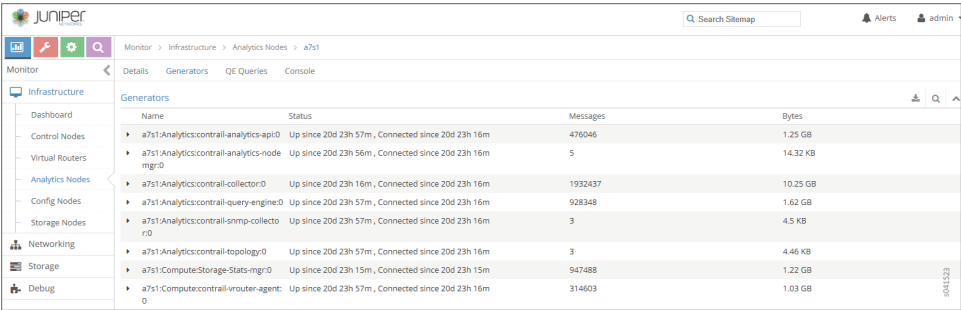
Table 48: Monitor Analytics Individual Node Details Tab Fields

Field	Description
Hostname	The name of this node.
IP Address	The IP address of this node.
Version	The installed version of the software.
Overall Node Status	The current operational status of the node — Up or Down — and the length of time in this state.
Processes	The current status of each analytics process, including Collector, Query Engine, and OpServer.
CPU (%)	The average CPU percentage usage for this node.
Memory	The average memory usage of this node.
Messages	The total number of messages for this node.
Generators	The total number of generators associated with this node.
Last Log	The date and time of the last log message issued about this node.

Monitor Analytics Individual Node Generators Tab

The **Generators** tab displays information about the generators for an individual analytics node; see [Figure 124 on page 309](#). Click the expansion arrow next to any generator name to reveal more details. See [Table 49 on page 309](#) for descriptions of the fields on the **Peers** tab screen.

Figure 124: Individual Analytics Node—Generators Tab



Name	Status	Messages	Bytes
a7s1:Analytics:contrail-analytics-api0	Up since 20d 23h 57m, Connected since 20d 23h 16m	476046	1.25 GB
a7s1:Analytics:contrail-analytics-node-mgr0	Up since 20d 23h 56m, Connected since 20d 23h 16m	5	14.32 KB
a7s1:Analytics:contrail-collector0	Up since 20d 23h 16m, Connected since 20d 23h 16m	1932437	10.25 GB
a7s1:Analytics:contrail-query-engine0	Up since 20d 23h 57m, Connected since 20d 23h 16m	928348	1.62 GB
a7s1:Analytics:contrail-snmp-collector0	Up since 20d 23h 57m, Connected since 20d 23h 16m	3	4.5 KB
a7s1:Analytics:contrail-topology0	Up since 20d 23h 57m, Connected since 20d 23h 16m	3	4.46 KB
a7s1:Compute:Storage-Stats-mgr0	Up since 20d 23h 15m, Connected since 20d 23h 15m	947488	1.22 GB
a7s1:Compute:contrail-vrouter-agent0	Up since 20d 23h 57m, Connected since 20d 23h 16m	314603	1.03 GB

Table 49: Monitor Analytics Individual Node Generators Tab Fields


Field	Description
Name	The host name of the generator.
Status	The current status of the peer— Up or Down — and the length of time in that state.
Messages	The number of messages sent and received from this peer.
Bytes	The total message size in bytes.

Monitor Analytics Individual Node QE Queries Tab

The **QE Queries** tab displays the number of query expansion (QE) messages that are in the queue for this analytics node. See [Figure 125 on page 309](#).

See [Table 50 on page 309](#) for descriptions of the fields on the **QE Queries** tab screen.

Figure 125: Individual Analytics Node—QE QueriesTab



Enqueue Time	Query	Progress
No QE Queries to display		

Table 50: Analytics Node QE Queries Tab Fields

Field	Description
Enqueue Time	The length of time this message has been in the queue waiting to be delivered.
Query	The query message.
Progress (%)	The percentage progress for the message delivery.

## Monitor Analytics Individual Node Console Tab

Click the **Console** tab for an individual analytics node to display system logging information for a defined time period. See [Figure 126 on page 310](#). See [Table 51 on page 310](#) for descriptions of the fields on the **Console** tab screen.

**Figure 126: Analytics Individual Node—Console Tab**

**Table 51: Monitor Analytics Individual Node Console Tab Fields**

Field	Description
<b>Time Range</b>	Select a timeframe for which to review logging information as sent to the console. There are 11 options, ranging from the <b>Last 5 mins</b> through to the <b>Last 24 hrs</b> . The default display is for the <b>Last 5 mins</b> .
<b>Log Category</b>	Select a log category to display: <ul style="list-style-type: none"> <li>All</li> <li>_default_</li> <li>XMPP</li> <li>IFMap</li> <li>TCP</li> </ul>
<b>Log Type</b>	Select a log type to display.
<b>Log Level</b>	Select a log severity level to display: <ul style="list-style-type: none"> <li>SYS_EMERG</li> <li>SYS_ALERT</li> <li>SYS_CRIT</li> <li>SYS_ERR</li> <li>SYS_WARN</li> <li>SYS_NOTICE</li> <li>SYS_INFO</li> <li>SYS_DEBUG</li> </ul>
<b>Keywords</b>	Enter any text string to search for and display logs containing that string.



Table 51: Monitor Analytics Individual Node Console Tab Fields (*continued*)

Field	Description
(Limit field)	Select the number of messages to display:  No Limit Limit 10 messages Limit 50 messages Limit 100 messages Limit 200 messages Limit 500 messages
<b>Auto Refresh</b>	Click the check box to automatically refresh the display if more messages occur.
<b>Display Logs</b>	Click this button to refresh the display if you change the display criteria.
<b>Reset</b>	Click this button to clear any selected display criteria and reset all criteria to their default settings.
<b>Time</b>	This column lists the time received for each log message displayed.
<b>Category</b>	This column lists the log category for each log message displayed.
<b>Log Type</b>	This column lists the log type for each log message displayed.
<b>Log</b>	This column lists the log message for each log displayed.

## Monitor > Infrastructure > Config Nodes

Select **Monitor > Infrastructure > Config Nodes** to view the information about the system config nodes.

- [Monitor Config Nodes on page 311](#)
- [Monitor Individual Config Node Details on page 312](#)
- [Monitor Individual Config Node Console on page 313](#)

## Monitor Config Nodes

Select **Monitor > Infrastructure > Config Nodes** to view a summary of activities for the analytics nodes. See [Figure 127 on page 312](#).

Figure 127: Config Nodes Summary




Table 52 on page 312 describes the fields in the Config Nodes summary.

Table 52: Config Nodes Summary Fields

Field	Description
Host name	The name of this node.
IP address	The IP address of this node.
Version	The version of software installed on the system.
Status	The current operational status of the node — Up or Down — and the length of time it is in that state.
CPU (%)	The average CPU percentage usage for this node.
Memory	The average memory usage for this node.

## Monitor Individual Config Node Details

Click the name of any config node displayed on the config nodes summary to view the **Details** tab for that node; see [Figure 128 on page 312](#).

Figure 128: Individual Config Nodes— Details Tab




Table 53 on page 313 describes the fields on the Details screen.


Table 53: Individual Config Nodes— Details Tab Fields

Field	Description
Hostname	The name of the config node.
IP Address	The IP address of this node.
Version	The installed version of the software.
Overall Node Status	The current operational status of the node — Up or Down — and the length of time it is in this state.
Processes	The current operational status of the processes associated with the config node, including AI Server, Schema Transformer, Service Monitor, Discovery, and Ifmap.
Analytics Node	The analytics node associated with this node.
CPU (%)	The average CPU percentage usage for this node.
Memory	The average memory usage by this node.

## Monitor Individual Config Node Console

Click the **Console** tab for an individual config node to display system logging information for a defined time period. See [Figure 129 on page 313](#).

Figure 129: Individual Config Node—Console Tab



See [Table 54 on page 313](#) for descriptions of the fields on the **Console** tab screen.

Table 54: Individual Config Node-Console Tab Fields

Field	Description
Time Range	Select a timeframe for which to review logging information as sent to the console. Use the drop down calendar in the fields From Time and To Time to select the date and times to include in the time range for viewing.
Log Category	Select from the drop down menu a log category to display. The option to view All is also available.
Log Type	Select a log type to display.

Table 54: Individual Config Node-Console Tab Fields (*continued*)

Field	Description
Log Level	Select a log severity level to display:
Limit	Select from a list an amount to limit the number of messages displayed:  All Limit 10 messages Limit 50 messages Limit 100 messages Limit 200 messages Limit 500 messages
Keywords	Enter any key words by which to filter the log messages displayed.
Auto Refresh	Click the check box to automatically refresh the display if more messages occur.
Display Logs	Click this button to refresh the display if you change the display criteria.
Reset	Click this button to clear any selected display criteria and reset all criteria to their default settings.

## Monitor > Networking


The **Monitor -> Networking** pages give an overview of the networking traffic statistics and health of domains, projects within domains, virtual networks within projects, and virtual machines within virtual networks.

- [Monitor > Networking Menu Options on page 314](#)
- [Monitor -> Networking -> Dashboard on page 315](#)
- [Monitor > Networking > Projects on page 316](#)
- [Monitor Projects Detail on page 317](#)
- [Monitor > Networking > Networks on page 319](#)

### Monitor > Networking Menu Options

[Figure 130 on page 315](#) shows the menu options available under **Monitor > Networking**.

Figure 130: Monitor Networking Menu Options



### Monitor -> Networking -> Dashboard

Select **Monitor -> Networking -> Dashboard** to gain insight into usage statistics for domains, virtual networks, projects, and virtual machines. When you select this option, the Traffic Statistics for Domain window is displayed as shown in [Figure 131 on page 315](#).

Figure 131: Traffic Statistics for Domain Window



Table 55 on page 316 describes the fields in the Traffic Statistics for Domain window.

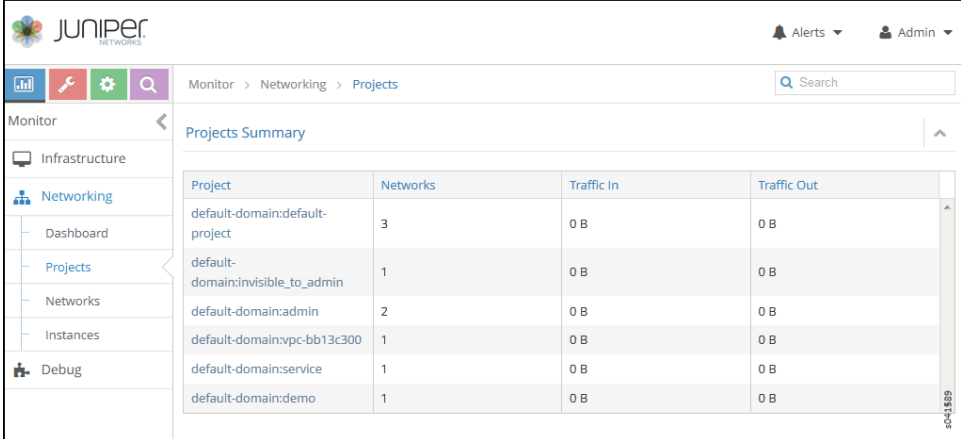
**Table 55: Projects Summary Fields**

Field	Description
<b>Total Traffic In</b>	The volume of traffic into this domain
<b>Total Traffic Out</b>	The volume of traffic out of this domain.
<b>Inter VN Traffic In</b>	The volume of inter-virtual network traffic into this domain.
<b>Inter VN Traffic Out</b>	The volume of inter-virtual network traffic out of this domain.
<b>Projects</b>	This chart displays the networks and interfaces for projects with the most throughput over the past 30 minutes. Click <b>Projects</b> then select <b>Monitor &gt; Networking &gt; Projects</b> , to display more detailed statistics.
<b>Networks</b>	This chart displays the networks for projects with the most throughput over the past 30 minutes. Click <b>Networks</b> then select <b>Monitor &gt; Networking &gt; Networks</b> , to display more detailed statistics.

## Monitor > Networking > Projects

Select **Monitor > Networking > Projects** to see information about projects in the system. See Figure 132 on page 316.

**Figure 132: Monitor > Networking > Projects**



Project	Networks	Traffic In	Traffic Out
default-domain:default-project	3	0 B	0 B
default-domain:invisible_to_admin	1	0 B	0 B
default-domain:admin	2	0 B	0 B
default-domain:vpc-bb13c300	1	0 B	0 B
default-domain:service	1	0 B	0 B
default-domain:demo	1	0 B	0 B

See Table 56 on page 316 for descriptions of the fields on this screen.

**Table 56: Projects Summary Fields**

Field	Description
<b>Projects</b>	The name of the project. You can click the name to access details about connectivity for this project.
<b>Networks</b>	The volume of inter-virtual network traffic out of this domain.


Table 56: Projects Summary Fields (*continued*)

Field	Description
Traffic In	The volume of traffic into this domain.
Traffic Out	The volume of traffic out of this domain.

## Monitor Projects Detail

You can click any of the projects listed on the Projects Summary to get details about connectivity, source and destination port distribution, and instances. When you click an individual project, the Summary tab for Connectivity Details is displayed as shown in [Figure 133 on page 317](#). Hover over any of the connections to get more details.


Figure 133: Monitor Projects Connectivity Details



In the Connectivity Details window you can click the links between the virtual networks to view the traffic statistics between the virtual networks.

The Traffic Statistics information is also available when you select **Monitor > Networking > Networks** as shown in [Figure 134 on page 317](#).

Figure 134: Traffic Statistics Between Networks



In the Connectivity Details window you can click the Instances tab to get a summary of details for each of the instances in this project.

Figure 135: Projects Instances Summary

Instance	Virtual Network	Interfaces	vRouter	IP Address	Floating IP	Traffic (In/Out)
out	default-domain:admin:right	1	hp1	2.2.2.252		129.87 KB / 119.83 KB
NAT1_1	default-domain:admin:right	1	hp1	2.2.2.253 250.250.1.253 (1 more)		3.69 MB / 1.15 MB
in	default-domain:admin:left	1	hp1	1.1.1.252		132.75 KB / 122.02 KB

See Table 3 for a description of the fields on this screen.


Table 57: Projects Instances Summary Fields

Field	Description
Instance	The name of the instance. Click the name then select <b>Monitor &gt; Networking &gt; Instances</b> to display details about the traffic statistics for this instance.
Virtual Network	The virtual network associated with this instance.
Interfaces	The number of interfaces associated with this instance.
vRouter	The name of the vRouter associated with this instance.
IP Address	Any IP addresses associated with this instance.
Floating IP	Any floating IP addresses associated with this instance.
Traffic (In/Out)	The volume of traffic in KB or MB that is passing in and out of this instance.

Select **Monitor > Networking > Instances** to display instance traffic statistics as shown in [Figure 136 on page 319](#).



Figure 136: Instance Traffic Statistics



Monitor > Networking > Networks

Select **Monitor > Networking > Networks** to view a summary of the virtual networks in your system. See [Figure 137 on page 319](#).

Figure 137: Network Summary

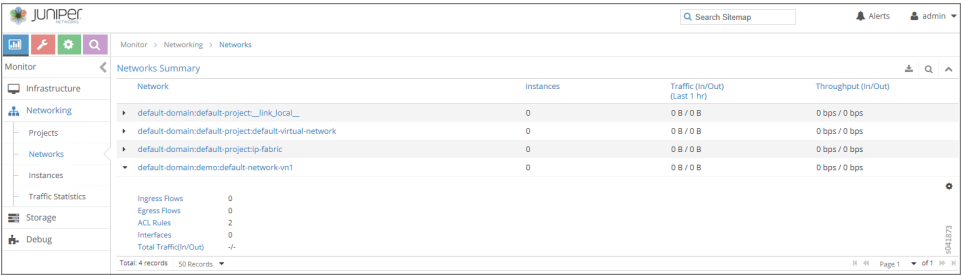


Table 58: Network Summary Fields

Field	Description
Network	The domain and network name of the virtual network. Click the arrow next to the name to display more information about the network, including the number of ingress and egress flows, the number of ACL rules, the number of interfaces, and the total traffic in and out.
Instances	The number of instances launched in this network.
Traffic (In/Out)	The volume of inter-virtual network traffic in and out of this network.
Throughput (In/Out)	The throughput of inter-virtual network traffic in and out of this network.

At **Monitor > Networking > Networks** you can click on the name of any of the listed networks to get details about the network connectivity, traffic statistics, port distribution, instances, and other details, by clicking the tabs across the top of the page.

Figure 138 on page 320 shows the **Summary** tab for an individual network, which displays connectivity details and traffic statistics for the selected network.

**Figure 138: Individual Network Connectivity Details—Summary Tab**




Figure 139 on page 320 shows the **Port Map** tab for an individual network, which displays the relative distribution of traffic for this network by protocol, by port.

**Figure 139: Individual Network— Port Map Tab**




Figure 140 on page 321 shows the **Port Distribution** tab for an individual network, which displays the relative distribution of traffic in and out by source port and destination port.

Figure 140: Individual Network— Port Distribution Tab




Figure 141 on page 321 shows the **Instances** tab for an individual network, which displays details for each instance associated with this network, including the number of interfaces, the associated vRouter, the instance IP address, and the volume of traffic in and out.

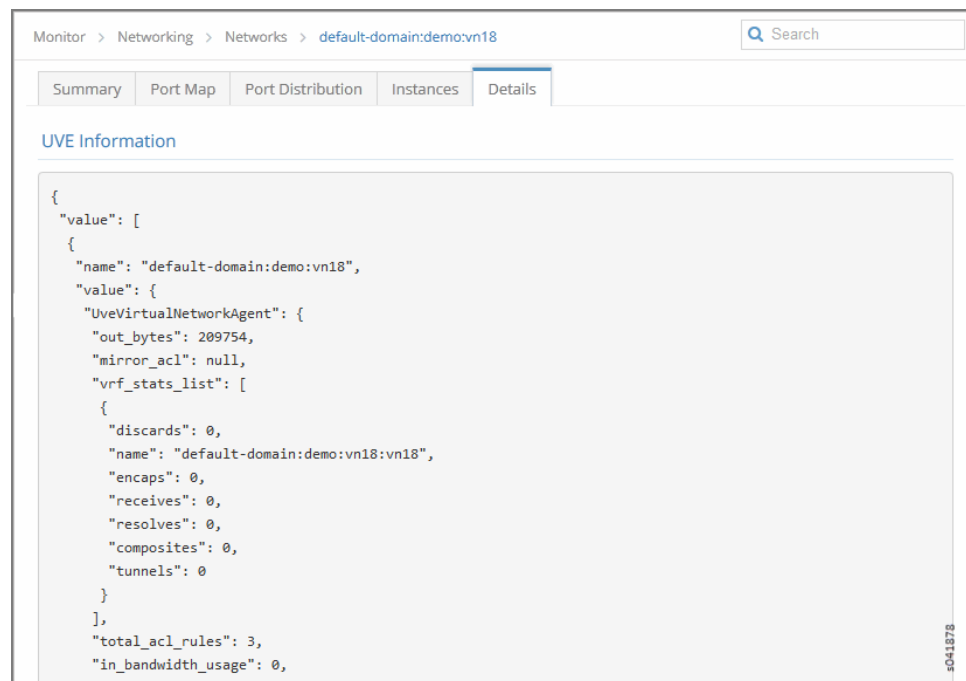
Additionally, you can click the arrow near the instance name to reveal even more details about the instance—the interfaces and their addresses, UUID, CPU (usage), and memory used of the total amount available.

Figure 141: Individual Network Instances Tab

Monitor > Networking > Networks > default-domain:demo:vn18						
Search						
Summary Port Map Port Distribution <b>Instances</b> Details						
Instances Summary						
	Instance	Interfaces	vRouter	IP Address	Floating IP	Traffic (In/Out)
▶	vn18_vm-b342ca93-9acd-4275-acb8-df7b5843884c	1	b1s29	192.168.18.225		1.13 KB / 712.00 B
▲	vn18_vm-22a42bf6-fccc-4db3-b5ac-80082bbefbef	1	b1s42	192.168.18.236		1.13 KB / 712.00 B
	Interfaces IP Address: 192.168.18.236 Label: 17 Mac Address: 02:e9:94:e7:0e:56 Network: default-domain:demo:vn18 Traffic (In/Out): 1.13 KB/712.00 B UUID 22a42bf6-fccc-4db3-b5ac-80082bbefbef CPU 0.01 Memory (Used/Total) 1.23 GB / 15.63 GB					
▶	vn18_vm-f676567a-826f-4e9d-9a81-b4649b7fcde2	1	b1s15	192.168.18.235		1.13 KB / 712.00 B

Figure 142 on page 322 shows the **Details** tab for an individual network, which displays the code used to define this network --the User Virtual Environment (UVE) code.

Figure 142: Individual Network Details Tab



## Query > Flows

Select **Query > Flows** to perform rich and complex SQL-like queries on flows in the Contrail Controller. You can use the query results for such things as gaining insight into the operation of applications in a virtual network, performing historical analysis of flow issues, and pinpointing problem areas with flows.

- [Query > Flows > Flow Series on page 322](#)
- [Example: Query Flow Series on page 325](#)
- [Query > Flow Records on page 326](#)
- [Query > Flows > Query Queue on page 328](#)

### Query > Flows > Flow Series

Select **Query > Flows > Flow Series** to create queries of the flow series table. The results are in the form of time series data for flow series. See [Figure 143 on page 323](#)

Figure 14-3: Query Flow Series Window

The query fields available on the screen for the **Flow Series** tab are described in [Table 59 on page 323](#). Enter query data into the fields to create a SQL-like query to display and analyze flows.

Table 59: Query Flow Series Fields

Field	Description
<b>Time Range</b>	<p>Select a range of time to display the flow series:</p> <ul style="list-style-type: none"> <li>• Last 10 Mins</li> <li>• Last 30 Mins</li> <li>• Last 1 Hr</li> <li>• Last 6 Hrs</li> <li>• Last 12 Hrs</li> <li>• Custom</li> </ul> <p>Click <b>Custom</b> to enter a specific custom time range in two fields: <b>From Time</b> and <b>To Time</b>.</p>
<b>Select</b>	Click the edit button (pencil icon) to open a <b>Select</b> window ( <a href="#">Figure 14-4 on page 324</a> ), where you can click one or more boxes to select the fields to display from the flow series, such as <b>Source VN</b> , <b>Dest VN</b> , <b>Bytes</b> , <b>Packets</b> , and more.
<b>Where</b>	Click the edit button (pencil icon) to open a query-writing window, where you can specify query values for variables such as <b>sourcevn</b> , <b>sourceip</b> , <b>destvn</b> , <b>destip</b> , <b>protocol</b> , <b>sport</b> , <b>dport</b> .
<b>Direction</b>	Select the desired flow direction: <b>INGRESS</b> or <b>EGRESS</b> .
<b>Filter</b>	Click the edit button (pencil icon) to open a <b>Filter</b> window ( <a href="#">Figure 14-5 on page 325</a> ), where you can select filter items to sort by, the sort order, and limits to the number of results returned.
<b>Run Query</b>	Click <b>Run Query</b> to retrieve the flows that match the query you created. The flows are listed on the lower portion of the screen in a box with columns identifying the selected fields for each flow.
(graph buttons)	When <b>Time Granularity</b> is selected, you have the option to view results in graph or flowchart form. Graph buttons appear on the screen above the <b>Export</b> button. Click a graph button to transform the tabular results into a graphical chart display.

Table 59: Query Flow Series Fields (*continued*)

Field	Description
<b>Export</b>	The Export button is displayed after you click <b>Run Query</b> . This allows you to export the list of flows to a text <b>.csv</b> file.

The **Select** window allows you to select one or more attributes of a flow series by clicking the check box for each attribute desired, see [Figure 144 on page 324](#). The upper section of the **Select** window includes field names, and the lower portion lets you select units. Select **Time Granularity** and then select **SUM(Bytes)** or **SUM(Packets)** to aggregate bytes and packets in intervals.

Figure 144: Flow Series Select

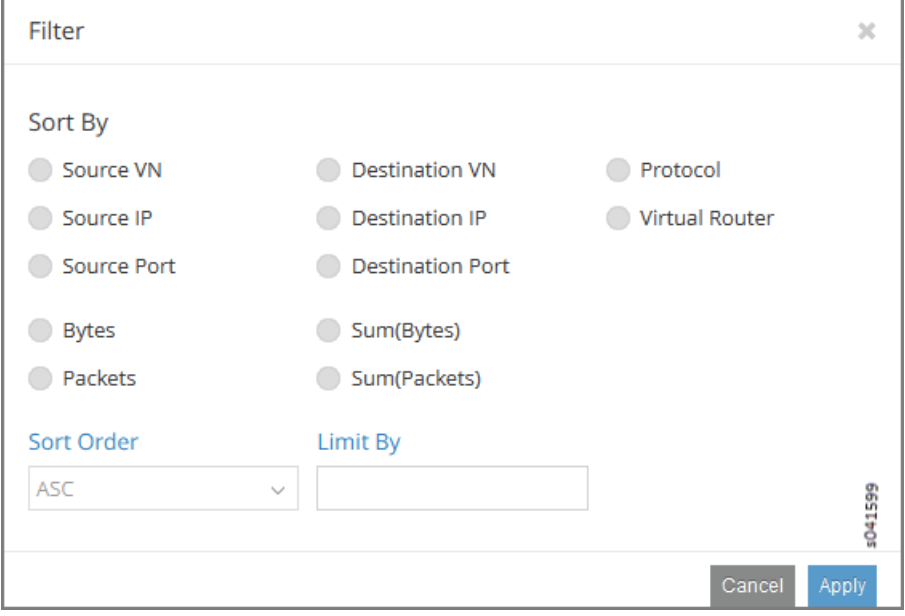
The screenshot shows a 'Select' dialog box with a close button (X) in the top right corner. The dialog contains a list of attributes with checkboxes next to them, arranged in three columns:

- Column 1: ☐ Source VN, ☐ Source IP, ☐ Source Port, ☐ Bytes, ☐ Packets
- Column 2: ☐ Destination VN, ☐ Destination IP, ☐ Destination Port, ☐ SUM(Bytes), ☐ SUM(Packets)
- Column 3: ☐ Time Granularity, ☐ Protocol, ☐ Virtual Router

At the bottom right of the dialog, there are two buttons: 'Cancel' and 'Apply'. On the far right edge of the dialog, there is a vertical label 's041600'.

Use the **Filter** window to refine the display of query results for flows, by defining an attribute by which to sort the results, the sort order of the results, and any limit needed to restrict the number of results. See [Figure 145 on page 325](#).

Figure 145: Flow Series Filter



The 'Filter' dialog box contains the following elements:

- Sort By:** A grid of radio buttons for selecting the sort criteria: Source VN, Destination VN, Protocol, Source IP, Destination IP, Virtual Router, Source Port, Destination Port, Bytes, Sum(Bytes), Packets, and Sum(Packets).
- Sort Order:** A dropdown menu currently set to 'ASC'.
- Limit By:** An empty text input field.
- Buttons:** 'Cancel' and 'Apply' buttons at the bottom right.
- Watermark:** A vertical watermark '© 2016 Juniper Networks, Inc.' is visible on the right side.

### Example: Query Flow Series

The following is an example flow series query that returns the time series of the summation traffic in bytes for all combinations of source VN and destination VN for the last 10 minutes, with the bytes aggregated in 10 second intervals. See [Figure 146 on page 325](#).

Figure 146: Example: Query Flow Series



The 'Query Flow Series' interface includes the following fields and controls:

- Time Range:** A dropdown menu set to 'Last 10 Mins'.
- Select:** A text input field containing the query 'sourcevn, destvn, time-granularity, sum(bytes)'.
- Where:** An empty text input field.
- Filter:** An empty text input field.
- Time Granularity:** A dropdown menu set to '10' with a unit of 'secs'.
- Direction:** A dropdown menu set to 'INGRESS'.
- Buttons:** A 'Run Query' button at the bottom left.
- Watermark:** A vertical watermark '© 2016 Juniper Networks, Inc.' is visible on the right side.

The query returns tabular time series data, see [Figure 147 on page 326](#), for the following combinations of Source VN and Dest VN:


1. Flow Class 1: Source VN = default-domain:demo:front-end, Dest VN= \_\_UNKNOWN\_\_
2. Flow Class 2: Source VN = default-domain:demo:front-end, Dest VN=default-domain:demo:back-end

Figure 147: Query Flow Series Tabular Results

Query Flow Series				
Query Results				
Time	Source VN	Dest. VN	Direction	SUM(Bytes)
2013-08-05 18:59:30:0	default-domain:demo:vn0	default-domain:demo:vn16	INGRESS	421,128
2013-08-05 18:59:40:0	default-domain:demo:vn0	default-domain:demo:vn16	INGRESS	227,000
2013-08-05 18:59:50:0	default-domain:demo:vn0	default-domain:demo:vn16	INGRESS	216,816
2013-08-05 19:00:00:0	default-domain:demo:vn0	default-domain:demo:vn16	INGRESS	387,036
2013-08-05 18:59:30:0	default-domain:demo:vn1	default-domain:demo:vn17	INGRESS	52,944
2013-08-05 18:59:40:0	default-domain:demo:vn1	default-domain:demo:vn17	INGRESS	52,692
2013-08-05 18:59:50:0	default-domain:demo:vn1	default-domain:demo:vn17	INGRESS	58,040
2013-08-05 19:00:00:0	default-domain:demo:vn1	default-domain:demo:vn17	INGRESS	42,480
2013-08-05 18:59:30:0	default-domain:demo:vn16	default-domain:demo:vn0	INGRESS	17,832
2013-08-05 18:59:40:0	default-domain:demo:vn16	default-domain:demo:vn0	INGRESS	27,320
2013-08-05 18:59:50:0	default-domain:demo:vn16	default-domain:demo:vn0	INGRESS	20,792
2013-08-05 19:00:00:0	default-domain:demo:vn16	default-domain:demo:vn0	INGRESS	10,404

Because **Time Granularity** is selected, the results can also be displayed as graphical charts. Click the graph button on the right side of the tabular results. The results are displayed in a graphical flow chart. See [Figure 148 on page 326](#).

Figure 148: Query Flow Series Graphical Results



## Query > Flow Records

Select **Query > Flow Records** to create queries of individual flow records for detailed debugging of connectivity issues between applications and virtual machines. Queries at this level return records of the active flows within a given time period.



Figure 149: Flow Records

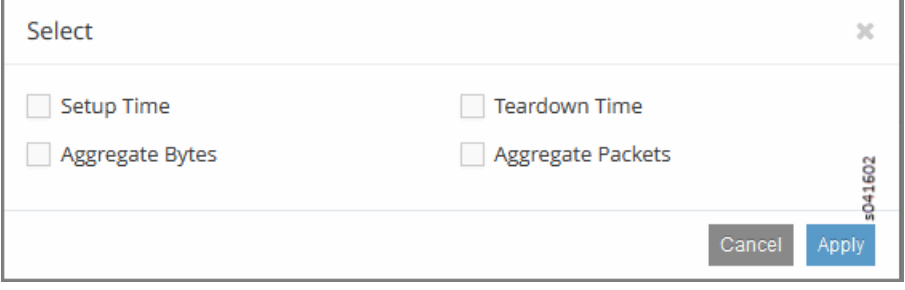
The query fields available on the screen for the **Flow Records** tab are described in [Table 60 on page 327](#). Enter query data into the fields to create an SQL-like query to display and analyze flows.

Table 60: Query Flow Records Fields

Field	Description
<b>Time Range</b>	<p>Select a range of time for the flow records:</p> <ul style="list-style-type: none"> <li>• Last 10 Mins</li> <li>• Last 30 Mins</li> <li>• Last 1 Hr</li> <li>• Last 6 Hrs</li> <li>• Last 12 Hrs</li> <li>• Custom</li> </ul> <p>Click <b>Custom</b> to enter a specified custom time range in two fields: <b>From Time</b> and <b>To Time</b>.</p>
<b>Select</b>	Click the edit button (pencil icon) to open a <b>Select</b> window ( <a href="#">Figure 150 on page 328</a> ), where you can click one or more boxes to select attributes to display for the flow records, including <b>Setup Time</b> , <b>Teardown Time</b> , <b>Aggregate Bytes</b> , and <b>Aggregate Packets</b> .
<b>Where</b>	Click the edit button (pencil icon) to open a query-writing window where you can specify query values for <b>sourcevn</b> , <b>sourceip</b> , <b>destvn</b> , <b>destip</b> , <b>protocol</b> , <b>sport</b> , <b>dport</b> . .
<b>Direction</b>	Select the desired flow direction: <b>INGRESS</b> or <b>EGRESS</b> .
<b>Run Query</b>	Click <b>Run Query</b> to retrieve the flow records that match the query you created. The records are listed on the lower portion of the screen in a box with columns identifying the fields for each flow.
<b>Export</b>	The <b>Export</b> button is displayed after you click <b>Run Query</b> , allowing you to export the list of flows to a text <b>.csv</b> file.

The **Select** window allows you to select one or more attributes to display for the flow records selected, see [Figure 150 on page 328](#).

Figure 150: Flow Records Select Window



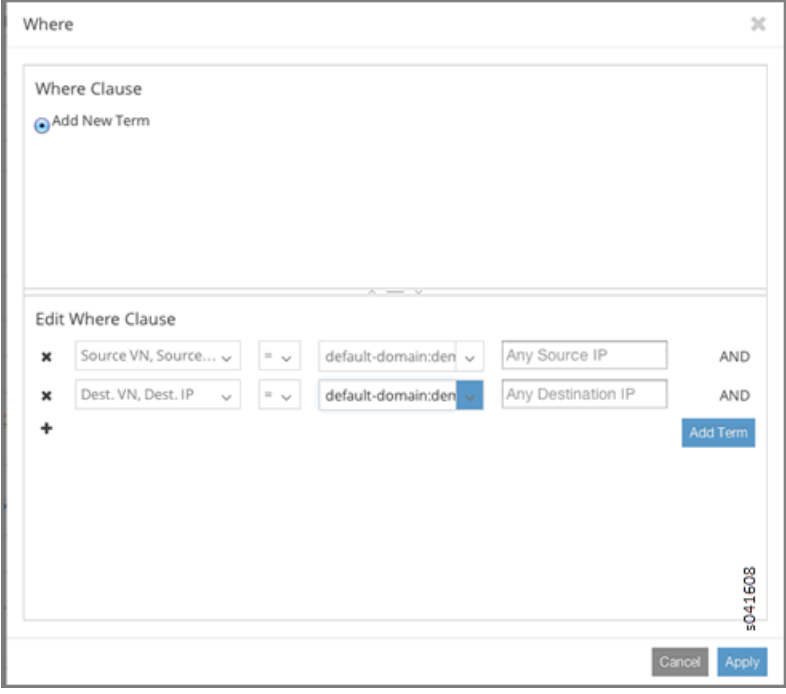
The 'Select' window contains four checkboxes: 'Setup Time', 'Teardown Time', 'Aggregate Bytes', and 'Aggregate Packets'. At the bottom right, there are 'Cancel' and 'Apply' buttons. A vertical label 's041602' is positioned on the right side of the window.

You can restrict the query to a particular source VN and destination VN combination using the **Where** section.

The **Where Clause** supports logical AND and logical OR operations, and is modeled as a logical OR of multiple AND terms. For example: ( (term1 AND term2 AND term3..) OR (term4 AND term5) OR...).

Each term is a single variable expression such as **Source VN = VN1**.

Figure 151: Where Clause Window



The 'Where' window shows a 'Where Clause' section with a radio button for 'Add New Term'. Below it is an 'Edit Where Clause' section. This section contains two rows of conditions, each with a dropdown menu, an equals sign, another dropdown menu, and a text input field. The first row is 'Source VN, Source...' = 'default-domain:den' 'Any Source IP' AND. The second row is 'Dest. VN, Dest. IP' = 'default-domain:den' 'Any Destination IP' AND. There is an 'Add Term' button at the bottom right of the 'Edit Where Clause' section. At the bottom of the window are 'Cancel' and 'Apply' buttons. A vertical label 's041608' is positioned on the right side of the window.

## Query > Flows > Query Queue

Select **Query > Flows > Query Queue** to display queries that are in the queue waiting to be performed on the data. See [Figure 152 on page 329](#).

Figure 152: Flows Query Queue

Date	Query	Progress	Records	Status	Time Taken
2013-10-09 18:07:06	{ "table": "FlowSeriesTable", "start_time": 1381267020000000, "end_time": 1381277820000000, "select_fields": { "flow_class_id", "direction_ing", "sum(bytes)", "T=60", "dir": 1 } }	100%	180	completed	150 secs
2013-10-09 17:55:48	{ "table": "FlowSeriesTable", "start_time": 1381267020000000, "end_time": 1381277820000000, "select_fields": { "flow_class_id", "direction_ing", "sum(bytes)", "T=60", "dir": 1 } }	100%	180	completed	145 secs
2013-10-09 17:29:39	{ "table": "FlowSeriesTable", "start_time": 1381267020000000, "end_time": 1381277820000000, "select_fields": { "flow_class_id", "direction_ing", "sum(bytes)", "T=60", "dir": 1 } }	100%	180	completed	170 secs
2013-10-09 16:57:10	{ "table": "FlowSeriesTable", "start_time": 1381267020000000, "end_time": 1381277820000000, "select_fields": { "flow_class_id", "direction_ing", "sum(bytes)", "T=60", "dir": 1 } }	100%	180	completed	270 secs
2013-10-09 16:39:48	{ "table": "FlowSeriesTable", "start_time": 1381360140000000, "end_time": 1381361940000000, "select_fields": { "flow_class_id", "direction_ing", "T=60", "sum(bytes)", "dir": 1 } }	100%	30	completed	60 secs
2013-10-09 11:07:29	{ "table": "FlowSeriesTable", "start_time": 1381338420000000, "end_time": 1381342020000000, "select_fields": { "flow_class_id", "direction_ing", "sum(bytes)", "T=60", "dir": 1 } }	100%	7	completed	15 secs

The query fields available on the screen for the **Flow Records** tab are described in [Table 61 on page 329](#). Enter query data into the fields to create an SQL-like query to display and analyze flows.

Table 61: Query Flow Records Fields

Field	Description
<b>Date</b>	The date and time the query was started.
<b>Query</b>	A display of the parameters set for the query.
<b>Progress</b>	The percentage completion of the query to date.
<b>Records</b>	The number of records matching the query to date.
<b>Status</b>	The status of the query, such as <b>completed</b> .
<b>Time Taken</b>	The amount of time in seconds it has taken the query to return the matching records.
(Action icon)	Click the <b>Action</b> icon and select <b>View Results</b> to view a list of the records that match the query, or click <b>Delete</b> to remove the query from the queue.

## Query > Logs


The **Query > Logs** option allows you to access the system log and object log activity of any Contrail Controller component from one central location.

- [Query > Logs Menu Options on page 329](#)
- [Query > Logs > System Logs on page 330](#)
- [Sample Query for System Logs on page 331](#)
- [Query > Logs > Object Logs on page 332](#)

## Query > Logs Menu Options

Click **Query > Logs** to access the **Query Logs** menu, where you can select **System Logs** to view system log activity, **Object Logs** to view object logs activity, and **Query Queue** to create custom queries of log activity; see [Figure 153 on page 330](#).

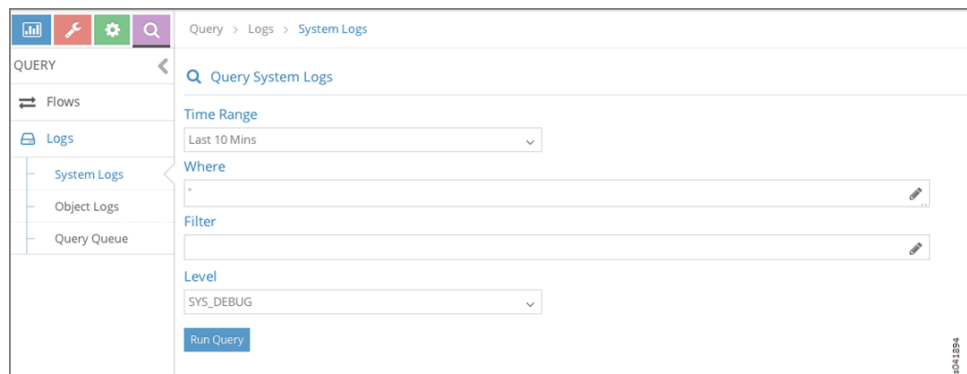
Figure 153: Query &gt; Logs



### Query > Logs > System Logs

Click **Query > Logs > System Logs** to access the **Query System Logs** menu, where you can view system logs according to criteria that you determine. See [Figure 154 on page 330](#).

Figure 154: Query &gt; Logs &gt; System Logs



The query fields available on the **Query System Logs** screen are described in [Table 62 on page 330](#).

Table 62: Query System Logs Fields

Field	Description
<b>Time Range</b>	<p>Select a range of time for which to see the system logs:</p> <ul style="list-style-type: none"> <li>• Last 10 Mins</li> <li>• Last 30 Mins</li> <li>• Last 1 Hr</li> <li>• Last 6 Hrs</li> <li>• Last 12 Hrs</li> <li>• Custom</li> </ul> <p>If you click Custom, enter a desired time range in two new fields: <b>From Time</b> and <b>To Time</b>.</p>

Table 62: Query System Logs Fields (*continued*)

Field	Description
<b>Where</b>	Click the edit button (pencil icon) to open a query-writing window, where you can specify query values for variables such as Source, Module, MessageType, and the like, in order to retrieve specific information.
<b>Level</b>	<p>Select the message severity level to view:</p> <ul style="list-style-type: none"> <li>• SYS_NOTICE</li> <li>• SYS_EMERG</li> <li>• SYS_ALERT</li> <li>• SYS_CRIT</li> <li>• SYS_ERR</li> <li>• SYS_WARN</li> <li>• SYS_INFO</li> <li>• SYS_DEBUG</li> </ul>
<b>Run Query</b>	Click this button to retrieve the system logs that match the query. The logs are listed in a box with columns showing the <b>Time</b> , <b>Source</b> , <b>Module Id</b> , <b>Category</b> , <b>Log Type</b> , and <b>Log</b> message.
<b>Export</b>	This button appears after you click <b>Run Query</b> , allowing you to export the list of system messages to a text/csv file.

### Sample Query for System Logs

This section shows a sample system logs query designed to show all **System Logs** from **ModuleId = VRouterAgent** on **Source = b1s16** and filtered by **Level = SYS\_DEBUG**.

1. At the **Query System Logs** screen, click in the **Where** field to access the **Where** query screen and enter information defining the location to query in the **Edit Where Clause** section and click **OK**; see [Figure 155 on page 332](#).

Figure 155: Edit Where Clause

2. The information you defined at the Where screen displays on the **Query System Logs**. Enter any more defining information needed; see [Figure 156 on page 332](#). When finished, click **Run Query** to display the results.

Figure 156: Sample Query System Logs

## Query > Logs > Object Logs

Object logs allow you to search for logs associated with a particular object, for example, all logs for a specified virtual network. Object logs record information related to modifications made to objects, including creation, deletion, and other modifications; see [Figure 157 on page 333](#).

Figure 157: Query &gt; Logs &gt; Object Logs

The query fields available on the **Object Logs** screen are described in [Table 63 on page 333](#).

Table 63: Object Logs Query Fields

Field	Description
<b>Time Range</b>	<p>Select a range of time for which to see the logs:</p> <ul style="list-style-type: none"> <li>• Last 10 Mins</li> <li>• Last 30 Mins</li> <li>• Last 1 Hr</li> <li>• Last 6 Hrs</li> <li>• Last 12 Hrs</li> <li>• Custom</li> </ul> <p>If you click Custom, enter a desired time range in two new fields: <b>From Time</b> and <b>To Time</b>.</p>
<b>Object Type</b>	<p>Select the object type for which to show logs:</p> <ul style="list-style-type: none"> <li>• Virtual Network</li> <li>• Virtual Machine</li> <li>• Virtual Router</li> <li>• BGP Peer</li> <li>• Routing Instance</li> <li>• XMPP Connection</li> </ul>
<b>Object Id</b>	Select from a list of available identifiers the name of the object you wish to use.
<b>Select</b>	<p>Click the edit button (pencil icon) to open a window where you can select searchable types by clicking a checkbox:</p> <ul style="list-style-type: none"> <li>• ObjectLog</li> <li>• SystemLog</li> </ul>

Table 63: Object Logs Query Fields (*continued*)

Field	Description
Where	Click the edit button (pencil icon) to open the query-writing window, where you can specify query values for variables such as <b>Source</b> , <b>ModuleId</b> , and <b>MessageType</b> , in order to retrieve information as specific as you wish.
Run Query	Click this button to retrieve the system logs that match the query. The logs are listed in a box with columns showing the <b>Time</b> , <b>Source</b> , <b>Module Id</b> , <b>Category</b> , <b>Log Type</b> , and <b>Log</b> message.
Export	This button appears after you click <b>Run Query</b> , allowing you to export the list of system messages to a text/csv file.

## Example: Debugging Connectivity Using Monitoring for Troubleshooting

- [Using Monitoring to Debug Connectivity on page 334](#)

### Using Monitoring to Debug Connectivity

This example shows how you can use monitoring to debug connectivity in your Contrail system. You can use the demo setup in Contrail to use these steps on your own.


1. Navigate to **Monitor -> Networking -> Networks -> default-domain:demo:vn0**, Instance **ed6abd16-250e-4ec5-a382-5cbc458fb0ca** with IP address **192.168.0.252** in the virtual network **vn0**; see [Figure 158 on page 334](#)

Figure 158: Navigate to Instance



2. Click the instance to view **Traffic Statistics for Instance**. see [Figure 159 on page 334](#).

Figure 159: Traffic Statistics for Instance




3. Instance **d26c0b31-c795-400e-b8be-4d3e6de77dcf** with IP address **192.168.0.253** in the virtual network **vn16**. see [Figure 160 on page 335](#) and [Figure 161 on page 335](#).



Figure 160: Navigate to Instance

Monitor > Networking > Networks > default-domain:demo:vn16		
Summary Port Map Port Distribution Instances Details		
Instance	Traffic In	Traffic Out
d26c0b31-c795-400e-b8be-4d3e6de77dcf	2.18 MB	2.13 MB
23045415-b679-4d7a-8f5d-96c162d628be	2.11 MB	2.16 MB

Figure 161: Traffic Statistics for Instance



4. From **Monitor->Infrastructure->Virtual Routers->a3s18->Interfaces**, we can see that **Instance ed6abd16-250e-4ec5-a382-5cbc458fb0ca** is hosted on **Virtual Router a3s18**; see [Figure 162 on page 335](#).

Figure 162: Navigate to a3s18 Interfaces

Monitor > Infrastructure > Virtual Routers > a3s18						
Details Console Interfaces Networks ACL Flows Routes						
Name	Label	Status	Network	IP Address	Floating IP	Instance
tap1d4e0121-4c	16	Up	default-domain:demo:vn0	192.168.0.252	None	ed6abd16-250e-4ec5-a382-5cbc458fb0ca
tap3d9dc3e1-97	18	Up	default-domain:demo:vn16	192.168.16.252	None	23045415-b679-4d7a-8f5d-96c162d628be
tap93b3b463-74	19	Up	default-domain:demo:vn17	192.168.17.252	None	99311eda-261e-47d8-b4a7-8d13d5d7d98f
tapc7d943c-6b	17	Up	default-domain:demo:vn1	192.168.1.252	None	20244e9f-a4ed-4a32-8039-15c75323572e

5. From **Monitor->Infrastructure->Virtual Routers->a3s19->Interfaces**, we can see that **Instance d26c0b31-c795-400e-b8be-4d3e6de77dcf** is hosted on **Virtual Router a3s19**; see [Figure 163 on page 335](#).

Figure 163: Navigate to a3s19 Interfaces

Monitor > Infrastructure > Virtual Routers > a3s19						
Details Console Interfaces Networks ACL Flows Routes						
Name	Label	Status	Network	IP Address	Floating IP	Instance
tap29583b2f-c2	19	Up	default-domain:demo:vn16	192.168.16.253	None	d26c0b31-c795-400e-b8be-4d3e6de77dcf
tapb257d21d-d3	18	Up	default-domain:demo:vn1	192.168.1.253	None	eeb0c321-7536-46e7-a454-cd1f13ac895
tapc3e3d87-66	17	Up	default-domain:demo:vn17	192.168.17.253	None	b2425f5-67e-4060-9478-81a4a8275541
tapc5ea87e3-55	16	Up	default-domain:demo:vn0	192.168.0.253	None	682b7414-c8ba-45ee-91bc-9c22cd86c9d

6. **Virtual Routers a3s18 and a3s19** have the **ACL** entries to allow connectivity between **default-domain:demo:vn0** and **default-domain:demo:vn16** networks; see [Figure 164 on page 336](#) and [Figure 165 on page 336](#).

Figure 164: ACL Connectivity a3s18

Monitor > Infrastructure > Virtual Routers > a3s18									
Details Console Interfaces Networks ACL Flows Routes									
UUID	Flows	Action	Protocol	Source Network or Prefix	Source Port	Destination Network or Prefix	Destination Port	Source Policy Rule	ACE Id
a7249326-3f50-477a-ad...	16	pass	any	default-domain:demo:vn0	any	default-domain:demo:vn16	any		1
		pass	any	default-domain:demo:vn16	any	default-domain:demo:vn0	any		2
		pass	any	default-domain:demo:vn0	any	default-domain:demo:vn0	any		3
b32143a3-0e80-4ae2-9c...	16	pass	any	default-domain:demo:vn1	any	default-domain:demo:vn17	any		1
		pass	any	default-domain:demo:vn17	any	default-domain:demo:vn1	any		2
		pass	any	default-domain:demo:vn1	any	default-domain:demo:vn1	any		3
b8c9810-e9fc-41b8-aa7...	16	pass	any	default-domain:demo:vn0	any	default-domain:demo:vn16	any		1
		pass	any	default-domain:demo:vn16	any	default-domain:demo:vn0	any		2
		pass	any	default-domain:demo:vn16	any	default-domain:demo:vn16	any		3
d1b47291-7a21-469e-8d...	16	pass	any	default-domain:demo:vn1	any	default-domain:demo:vn17	any		1
		pass	any	default-domain:demo:vn17	any	default-domain:demo:vn1	any		2
		pass	any	default-domain:demo:vn17	any	default-domain:demo:vn17	any		3

Figure 165: ACL Connectivity a3s19

Monitor > Infrastructure > Virtual Routers > a3s19									
Details Console Interfaces Networks ACL Flows Routes									
UUID	Flows	Action	Protocol	Source Network or Prefix	Source Port	Destination Network or Prefix	Destination Port	Source Policy Rule	ACE Id
a7249326-3f50-477a-ad...	16	pass	any	default-domain:demo:vn0	any	default-domain:demo:vn16	any		1
		pass	any	default-domain:demo:vn16	any	default-domain:demo:vn0	any		2
		pass	any	default-domain:demo:vn0	any	default-domain:demo:vn0	any		3
b32143a3-0e80-4ae2-9c...	16	pass	any	default-domain:demo:vn1	any	default-domain:demo:vn17	any		1
		pass	any	default-domain:demo:vn17	any	default-domain:demo:vn1	any		2
		pass	any	default-domain:demo:vn1	any	default-domain:demo:vn1	any		3
b8c9810-e9fc-41b8-aa7...	16	pass	any	default-domain:demo:vn0	any	default-domain:demo:vn16	any		1
		pass	any	default-domain:demo:vn16	any	default-domain:demo:vn0	any		2
		pass	any	default-domain:demo:vn16	any	default-domain:demo:vn16	any		3
d1b47291-7a21-469e-8d...	16	pass	any	default-domain:demo:vn1	any	default-domain:demo:vn17	any		1
		pass	any	default-domain:demo:vn17	any	default-domain:demo:vn1	any		2
		pass	any	default-domain:demo:vn17	any	default-domain:demo:vn17	any		3

- Next, verify the routes on the control node for routing instances **default-domain:demo:vn0:vn0** and **default-domain:demo:vn16:vn16**; see [Figure 166 on page 336](#) and [Figure 167 on page 337](#).

Figure 166: Routes default-domain:demo:vn0:vn0

Monitor > Infrastructure > Control Nodes > a3s15							
Details Console Peers Routes							
Routing Instance		default-domain:demo:vn0	Address Family		All	Limit 50 Routes	
Peer Source		All	Prefix		Prefix	Display Routes Reset	
Prefix	Address Family	Protocol	Source	Next hop	Label	Local Preference	AS Path
192.168.0.252/32	inet	XMPP	a3s18	10.84.17.4	16	100	-
192.168.0.253/32	inet	BGP	10.84.17.3	10.84.17.4	16	100	AS_PATH: 0
192.168.0.253/32	inet	XMPP	a3s19	10.84.17.5	16	100	-
192.168.0.253/32	inet	BGP	10.84.17.3	10.84.17.5	16	100	AS_PATH: 0
192.168.16.252/32	inet	XMPP	a3s18	10.84.17.4	17	100	-
192.168.16.253/32	inet	BGP	10.84.17.3	10.84.17.4	17	100	AS_PATH: 0
192.168.16.253/32	inet	XMPP	a3s19	10.84.17.5	17	100	-
192.168.16.253/32	inet	BGP	10.84.17.3	10.84.17.5	17	100	AS_PATH: 0
10.84.17.4:1:192.168.0.255,0.0.0.0	inetmcast	XMPP	a3s18	10.84.17.4	0	100	-
10.84.17.4:1:255.255.255.0.0.0.0	inetmcast	XMPP	a3s18	10.84.17.4	0	100	-
10.84.17.5:1:192.168.0.255,0.0.0.0	inetmcast	XMPP	a3s19	10.84.17.5	0	100	-
10.84.17.5:1:255.255.255.0.0.0.0	inetmcast	XMPP	a3s19	10.84.17.5	0	100	-

Figure 167: Routes default-domain:demo:vn16:vn16

Prefix	Address Family	Protocol	Source	Next hop	Label	Local Preference	AS Path
192.168.0.252/32	inet	XMPP	a3s18	10.84.17.4	16	100	-
192.168.0.252/32	inet	BGP	10.84.17.3	10.84.17.4	16	100	AS_PATH: 0
192.168.0.253/32	inet	XMPP	a3s19	10.84.17.5	16	100	-
192.168.16.252/32	inet	BGP	10.84.17.3	10.84.17.5	16	100	AS_PATH: 0
192.168.16.252/32	inet	XMPP	a3s18	10.84.17.4	17	100	-
192.168.16.253/32	inet	BGP	10.84.17.3	10.84.17.4	17	100	AS_PATH: 0
192.168.16.253/32	inet	XMPP	a3s19	10.84.17.5	17	100	-
192.168.16.253/32	inet	BGP	10.84.17.3	10.84.17.5	17	100	AS_PATH: 0
10.84.17.4:2:192.168.16.255.0.0.0.0	inetmcast	XMPP	a3s18	10.84.17.4	0	100	-
10.84.17.4:2:255.255.255.255.0.0.0.0	inetmcast	XMPP	a3s18	10.84.17.4	0	100	-
10.84.17.5:2:192.168.16.255.0.0.0.0	inetmcast	XMPP	a3s19	10.84.17.5	0	100	-
10.84.17.5:2:255.255.255.255.0.0.0.0	inetmcast	XMPP	a3s19	10.84.17.5	0	100	-

8. We can see that VRF **default-domain:demo:vn0:vn0** on Virtual Router **a3s18** has the appropriate route and next hop to reach VRF **default-domain:demo:front-end** on Virtual Router **a3s19**; see [Figure 168 on page 337](#).

Figure 168: Verify Route and Next Hop a3s18

Prefix	Next ho...	Next hop details
169.254.169.254 / 32	receive	Source: MData Dest VN: default-domain:default-project:___link_local___
192.168.0.252 / 32	interface	Interface: tap1dae0121-4c Dest VN: default-domain:demo:vn0
	interface	Interface: tap1dae0121-4c Dest VN: default-domain:demo:vn0
	interface	Interface: tap1dae0121-4c Dest VN: default-domain:demo:vn0
192.168.0.253 / 32	tunnel	Dest IP: 10.84.17.5 Dest VN: default-domain:demo:vn0 Label: 16
	tunnel	Dest IP: 10.84.17.5 Dest VN: default-domain:demo:vn0 Label: 16
192.168.0.254 / 32	interface	Interface: pkt0 Dest VN: default-domain:demo:vn0
192.168.16.252 / 32	interface	Interface: tap249de2e1-97 Dest VN: default-domain:demo:vn16
	interface	Interface: tap249de2e1-97 Dest VN: default-domain:demo:vn16
192.168.16.253 / 32	tunnel	Dest IP: 10.84.17.5 Dest VN: default-domain:demo:vn16 Label: 19

9. We can see that VRF **default-domain:demo:vn16:vn16** on Virtual Router **a3s19** has the appropriate route and next hop to reach VRF **default-domain:demo:vn0:vn0** on Virtual Router **a3s18**; see [Figure 169 on page 338](#).

Figure 169: Verify Route and Next Hop a3s19

Monitor > Infrastructure > Virtual Routers > a3s19		
Details	Console	Interfaces
Networks	ACL	Flows
Routes		
VRF default-domain:demo:vn16:vn16 Show Routes Unicast Multicast		
Prefix	Next ho...	Next hop details
169.254.169.254 / 32	receive	Source: MData Dest VN: default-domain:default-project:link_local__
192.168.0.252 / 32	tunnel	Dest IP: 10.84.17.4 Dest VN: default-domain:demo:vn0 Label: 16
	tunnel	Dest IP: 10.84.17.4 Dest VN: default-domain:demo:vn0 Label: 16
192.168.0.253 / 32	interface	Interface: tape5ea97e3-55 Dest VN: default-domain:demo:vn0
	interface	Interface: tape5ea97e3-55 Dest VN: default-domain:demo:vn0
192.168.16.252 / 32	tunnel	Dest IP: 10.84.17.4 Dest VN: default-domain:demo:vn16 Label: 18
	tunnel	Dest IP: 10.84.17.4 Dest VN: default-domain:demo:vn16 Label: 18
192.168.16.253 / 32	interface	Interface: tap29585b2f-c2 Dest VN: default-domain:demo:vn16
	interface	Interface: tap29585b2f-c2 Dest VN: default-domain:demo:vn16
	interface	Interface: tap29585b2f-c2 Dest VN: default-domain:demo:vn16
192.168.16.254 / 32	interface	Interface: pkt0 Dest VN: default-domain:demo:vn16

10. Finally, flows between instances (IPs 192.168.0.252 and 192.168.16.253) can be verified on Virtual Routers a3s18 and a3s19; see Figure 170 on page 338 and Figure 171 on page 338.

Figure 170: Flows for a3s18

Monitor > Infrastructure > Virtual Routers > a3s18								
Details	Console	Interfaces	Networks	ACL	Flows	Routes		
Active Flows: 64								
Protocol	Source Network	Source IP	Source Port	Destination Network	Destination IP	Destination Port	Bytes/Pkts	Setup Time
TCP	vn0	192.168.0.252	43434	vn16	192.168.16.253	9100	1884088/5417	21:00:22.131180 2013-Aug-06
TCP	vn16	192.168.16.253	9100	vn0	192.168.0.252	43434	1909668/5891	21:00:22.131193 2013-Aug-06
TCP	vn16	192.168.16.253	9101	vn0	192.168.0.252	53369	1903900/5805	21:00:22.206222 2013-Aug-06
TCP	vn0	192.168.0.252	53369	vn16	192.168.16.253	9101	1890088/5302	21:00:22.206207 2013-Aug-06
UDP	vn0	192.168.0.252	39522	vn16	192.168.16.252	9200	0/0	21:00:22.382861 2013-Aug-06
UDP	vn0	192.168.0.252	44794	vn16	192.168.16.253	9201	1707392/3144	21:00:24.104277 2013-Aug-06
UDP	vn16	192.168.16.253	9201	vn0	192.168.0.252	44794	1735788/3107	21:00:24.104293 2013-Aug-06
UDP	vn0	192.168.0.252	40561	vn16	192.168.16.253	9200	1693476/3067	21:00:22.037777 2013-Aug-06
UDP	vn16	192.168.16.253	9200	vn0	192.168.0.252	40561	1643324/3061	21:00:22.037887 2013-Aug-06
UDP	vn0	192.168.0.252	39522	vn16	192.168.16.252	9200	1676616/3074	21:00:22.306703 2013-Aug-06
TCP	vn0	192.168.0.252	34236	vn16	192.168.16.252	9100	1891368/5686	21:00:22.395695 2013-Aug-06
TCP	vn0	192.168.0.252	34236	vn16	192.168.16.252	9100	0/0	21:00:22.400371 2013-Aug-06

Figure 171: Flows for a3s19

Monitor > Infrastructure > Virtual Routers > a3s19								
Details	Console	Interfaces	Networks	ACL	Flows	Routes		
Active Flows: 64								
Protocol	Source Network	Source IP	Source Port	Destination Network	Destination IP	Destination Port	Bytes/Pkts	Setup Time
UDP	vn0	192.168.0.252	44794	vn16	192.168.16.253	9201	1069380/1975	21:00:24.111374 2013-Aug-06
UDP	vn16	192.168.16.253	9201	vn0	192.168.0.252	44794	1100604/1963	21:00:24.111380 2013-Aug-06
UDP	vn0	192.168.0.252	40561	vn16	192.168.16.253	9200	1046756/1877	21:00:22.047747 2013-Aug-06
UDP	vn16	192.168.16.253	9200	vn0	192.168.0.252	47270	1061900/1921	21:00:25.379941 2013-Aug-06
UDP	vn16	192.168.16.253	9200	vn0	192.168.0.252	40561	1010568/1914	21:00:22.047756 2013-Aug-06
TCP	vn16	192.168.16.253	9100	vn0	192.168.0.253	53314	1217772/3649	21:00:23.440564 2013-Aug-06
TCP	vn0	192.168.0.252	43434	vn16	192.168.16.253	9100	1196336/3400	21:00:22.137665 2013-Aug-06
TCP	vn16	192.168.16.253	9100	vn0	192.168.0.252	43434	1239616/3724	21:00:22.137679 2013-Aug-06
UDP	vn16	192.168.16.253	9200	vn0	192.168.0.253	47270	0/0	21:00:25.347968 2013-Aug-06
TCP	vn16	192.168.16.253	9100	vn0	192.168.0.253	53314	0/0	21:00:23.440090 2013-Aug-06
UDP	vn16	192.168.16.253	9201	vn0	192.168.0.253	53930	1088652/1953	21:00:25.440166 2013-Aug-06
TCP	vn16	192.168.16.253	9101	vn0	192.168.0.253	34551	0/0	21:00:23.514246 2013-Aug-06
TCP	vn16	192.168.16.253	9101	vn0	192.168.0.253	34551	1304723/3504	21:00:23.514651 2013-Aug-06

## PART 5

# Index

- [Index on page 341](#)



# Index

## Symbols

#, comments in configuration statements.....	xxiii
( ), in syntax descriptions.....	xxiii
< >, in syntax descriptions.....	xxiii
[ ], in configuration statements.....	xxiii
{ }, in configuration statements.....	xxiii
(pipe), in syntax descriptions.....	xxiii

## A

alerts.....	246
ASN	
global.....	16

## B

BGP peers.....	29
braces, in configuration statements.....	xxiii
brackets	
angle, in syntax descriptions.....	xxiii
square, in configuration statements.....	xxiii

## C

Cobbler.....	46
comments, in configuration statements.....	xxiii
compute nodes	
vRouter.....	4
XMPP agent.....	4
configure custom	
hostname.....	14
IP address.....	14
LAN port.....	14
nameserver.....	14
Contrail ISO.....	15
Contrail packages.....	15
contrail-logs.....	21
contrail-server-manager restart.....	48
control node	
configuring.....	29
control nodes.....	4
conventions	
text and syntax.....	xxii
curly braces, in configuration statements.....	xxiii

customer support.....	xxiv
contacting JTAC.....	xxiv

## D

dashboard.....	280
DHCP.....	46, 170
discovery servicr.....	178
DKMS.....	112
DNS.....	169
configuring.....	172
DHCP.....	170
IPAM.....	170
record types.....	171
scripts.....	177
<i>See also</i> Domain Name System	
documentation	
comments on.....	xxiii
Domain Name System.....	169
<i>See also</i> DNS	

## E

EX 4200.....	163
existing OpenStack.....	19

## F

font conventions.....	xxii
-----------------------	------

## H

hardware requirements.....	9
health check.....	223
heat template.....	201
hostname	
configure custom.....	14

## I

image	
creating.....	133
infrastructure.....	280
install.....	15
instance	
virtual machine.....	136
IP address	
configure custom.....	14
IP Address Management.....	169
<i>See also</i> IPAM	
IP address pool	
allocating.....	140
creating.....	139
floating.....	139, 140

IPAM.....169, 170  
     See also IP Address Management

## L

LAN port  
     configure custom.....14

## M

manuals  
     comments on.....xxiii  
 monitor.....280  
 multi-tier example.....157  
 MX 80.....163

## N

nameserver  
     configure custom.....14  
 network  
     create.....125  
     delete.....120, 127  
     Juniper.....120  
     OpenStack.....127  
 network policy  
     associating to a network.....123, 131  
     creating.....121, 129  
     Juniper.....121, 123  
     OpenStack.....129, 131

## O

object log.....21  
 OpenStack.....3

## P

parentheses, in syntax descriptions.....xxiii  
 policy  
     associating to a network.....123, 131  
     creating.....121, 129  
     Juniper.....121, 123  
     OpenStack.....129, 131  
 projects  
     creating.....116

## R

roles  
     cfgm.....11  
     collector.....11  
     compute.....11  
     control.....11  
     webui.....11

rpm.....15

## S

security groups  
     associating to an instance.....142  
 service chain  
     creating.....227, 235, 239  
     example.....227, 235, 239  
 service instance  
     commands.....227, 235, 239  
 service policy  
     commands.....227, 235, 239  
 service template  
     commands.....227, 235, 239  
 support, technical See technical support  
 syntax conventions.....xxii  
 syslog.....21

## T

technical support  
     contacting JTAC.....xxiv  
 testbed definitions.....16, 19  
 testbed.py.....16  
     existing OpenStack.....19  
 trace messages.....21

## U

user tags.....47

## V

virtual machine  
     instance.....136  
     launching.....136  
 virtual network  
     creating.....118, 125  
     Juniper.....118  
     OpenStack.....125  
 VSRX.....227, 235, 239