

Contrail™

Contrail Networking and Security User Guide

Published
2020-12-01

Release
5.1

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

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

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

Contrail™ Contrail Networking and Security User Guide

5.1

Copyright © 2020 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About the Documentation | viii

Documentation and Release Notes | ix

Documentation Conventions | ix

Documentation Feedback | xi

Requesting Technical Support | xii

Self-Help Online Tools and Resources | xii

Creating a Service Request with JTAC | xiii

1

Contrail Security

Security Policy Features | 15

Overview of Existing Network Policy and Security Groups in Contrail | 15

Security Policy Enhancements | 16

Using Tags and Configuration Objects to Enhance Security Policy | 16

Predefined Tags | 17

Tagging Objects | 17

Policy Application | 17

Configuration Objects | 17

Configuration Object Tag Object | 18

Address-Group Configuration Object | 19

Service-Group Configuration Object | 20

Application-policy-set Configuration Object | 20

Policy-management Configuration Object | 20

Firewall-policy Configuration Object | 21

Firewall-rule Configuration Object | 21

Using the Contrail Web User Interface to Manage Security Policies | 22

Adding Security Policies | 23

Managing Policy Tags | 26

Viewing Global Policies | 27

Visualizing Traffic Groups | 28

Security Policy Features in OpenStack | 32

Overview of Existing Network Policy and Security Groups in OpenStack | 33

Security Policy Enhancements | 34

Configuration Objects | 35

Configure Neutron as FWaaS | 35

Policy Generation | 36

Configuring Policy Generation | 37

2

Configuring Virtual Networks

Creating a Virtual Network with Juniper Networks Contrail | 46

Creating a Floating IP Address Pool | 49

Support for IPv6 Networks in Contrail | 51

Overview: IPv6 Networks in Contrail | 52

Creating IPv6 Virtual Networks in Contrail | 53

Address Assignments | 53

Adding IPv6 Peers | 54

Configuring EVPN and VXLAN | 55

Configuring the VXLAN Identifier Mode | 57

Configuring Forwarding | 60

Configuring the VXLAN Identifier | 62

Configuring Encapsulation Methods | 63

Support for EVPN Route Type 5 | 66

Support for EVPN Type 6 Selective Multicast Ethernet Tag Route | 67

Configuring EVPN Type-6 SMET Routes | 67

Support for L3VPN Inter AS Option C | 68

Configuring Inter AS Option C | 69

3

Deploying a Multi-Tier Web Application Using Contrail

Example: Deploying a Multi-Tier Web Application | 73

Multi-Tier Web Application Overview | 73

Example: Setting Up Virtual Networks for a Simple Tiered Web Application | 74

Verifying the Multi-Tier Web Application | 76

Sample Addressing Scheme for Simple Tiered Web Application | 77

Sample Physical Topology for Simple Tiered Web Application | 78

Sample Physical Topology Addressing | 79

Sample Network Configuration for Devices for Simple Tiered Web Application | 80

4

Configuring Services

Configuring DNS Servers | 88

DNS Overview | 88

Defining Multiple Virtual Domain Name Servers in Contrail | 89

IPAM and Virtual DNS | 89

DNS Record Types | 90

Configuring DNS on the User Interface | 91

Configuring DNS Using Scripts | 97

Distributed Service Resource Allocation with Containerized Contrail | 98

Replacement of Centralized Discovery Service | 99

New Distributed Resource Allocation Manager | 99

Deprecation of IF-MAP | 100

Changes in Configuration Files | 100

Support for Broadcast and Multicast | 107

Subnet Broadcast | 108

All-Broadcast/Limited-Broadcast and Link-Local Multicast | 108

Host Broadcast | 109

Configuring Service Chaining

Service Chaining | 111

- Service Chaining Basics | 111

- Service Chaining Configuration Elements | 112

Service Chaining MX Series Configuration | 115

ECMP Load Balancing in the Service Chain | 116

Service Chain Version 2 with Port Tuple | 118

- Overview of Port Tuple | 118

- Service Chain Version 2 Sample Workflow | 119

 - Service Chain with Equal-Cost Multipath in Active-Active Mode | 120

 - Service Chain Active-Standby Mode with Allowed Address Pair | 120

 - Service Chain with Static Route Table | 121

- Service Chain with Health Check | 121

Service Chain Route Reorigination | 122

- Overview: Service Chaining in Contrail | 122

- Route Aggregation | 124

 - Schema for Route Aggregation | 125

 - Configuring and Troubleshooting Route Aggregation | 127

- Routing Policy | 132

 - Applying Routing Policy | 133

 - Routing Policy Configuration | 135

 - Configuring and Troubleshooting Routing Policy | 136

 - Using a VNC Script to Create Routing Policy | 138

 - Verify Routing Policy in API Server | 140

 - Verify Routing Policy in the Control Node | 140

 - Verify Routing Policy Configuration in the Control Node | 141

 - Verify Routing Policy Configuration on the Routing Instance | 141

- Control for Route Reorigination | 142

 - Configuring and Troubleshooting Reorigination Control | 143

Example: Creating an In-Network Service Chain by Using Contrail Command | 144

Example: Creating an In-Network-NAT Service Chain | 153

Example: Creating a Transparent Service Chain by Using Contrail Command | 161

Using Static Routes with Services | 171

Static Routes for Service Instances | 171

Configuring Static Routes on a Service Instance | 172

Configuring Static Routes on Service Instance Interfaces | 173

Configuring Static Routes as Host Routes | 175

Configuring Metadata Service | 176

Optimizing Contrail

Source Network Address Translation (SNAT) | 179

Overview | 179

Neutron APIs for Routers | 180

Network Namespace | 181

Using the Web UI to Configure Routers with SNAT | 181

Using the Web UI to Configure Distributed SNAT | 182

About the Documentation

IN THIS SECTION

- [Documentation and Release Notes | ix](#)
- [Documentation Conventions | ix](#)
- [Documentation Feedback | xi](#)
- [Requesting Technical Support | xii](#)

Use this guide to understand how you can create and orchestrate highly secure virtual networks using Contrail Networking and Security. Contrail Networking provides dynamic end-to-end networking policy and control for any cloud, any workload, and any deployment, from a single user interface. And unlike the conventional firewall configuration of IP addresses and port ranges, Contrail Security implements intent-based policy framework that uses policy constructs such as tags, labels, application policy set, address groups, and service groups.

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

Table 1: Contrail Guides

Guide Name	Description
Contrail Installation and Upgrade Guide	Provides step-by-step instructions to install and bring up Contrail and its various components.
Contrail Deployment Guide	Provides information about the next steps to be taken after a successful installation of Contrail.
Contrail Fabric Lifecycle Management and BMS Guide	Provides information about Contrail underlay management and data center automation.
Contrail Networking and Security User Guide	Provides information about creating and orchestrating highly secure virtual networks.
Contrail Service Provider Focused Features Guide	Provides information about the features that are used by service providers.
Contrail Analytics and Troubleshooting Guide	Provides information about AppFormix and Contrail analytics.

Documentation and Release Notes

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

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

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

Documentation Conventions

Table 2 on page ix defines notice icons used in this guide.

Table 2: Notice Icons






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

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

Table 3: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only

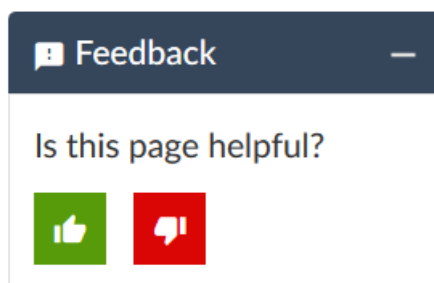
Table 3: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	<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">● In the Logical Interfaces box, select All Interfaces.● To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

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



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

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

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

1

CHAPTER

Contrail Security

Security Policy Features | **15**

Security Policy Features in OpenStack | **32**

Policy Generation | **36**

Configuring Policy Generation | **37**

Security Policy Features

IN THIS SECTION

- [Overview of Existing Network Policy and Security Groups in Contrail | 15](#)
- [Security Policy Enhancements | 16](#)
- [Using Tags and Configuration Objects to Enhance Security Policy | 16](#)
- [Configuration Objects | 17](#)
- [Using the Contrail Web User Interface to Manage Security Policies | 22](#)

Overview of Existing Network Policy and Security Groups in Contrail

Contrail virtual networks are isolated by default. Workloads in a virtual network cannot communicate with workloads in other virtual networks, by default. A Contrail network policy may be used to connect two virtual networks. In addition, Contrail network policy also provides security between two virtual networks by allowing or denying specified traffic.

In modern cloud environments, workloads are moving from one server to another, one rack to another and so on. Therefore, users must rely less on using IP addresses or other network coordinates to identify the endpoints to be protected. Instead users must leverage application attributes to author policies, so that the policies don't need to be updated on account of workload mobility.

You might want to segregate traffic based on the different categories of data origination, such as:

- Protecting the application itself
- Segregating traffic for specific component tiers within the application
- Segregating traffic based on the deployment environment for the application instance
- Segregating traffic based on the specific geographic location where the application is deployed

There are many other possible scenarios where traffic needs to be segregated.

Additionally, you might need to group workloads based on combinations of tags. These intents are hard to express with existing network policy constructs or Security Group constructs. Besides, existing policy constructs leveraging the network coordinates, must continually be rewritten or updated each time workloads move and their associated network coordinates change.

Security Policy Enhancements

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

Starting with Contrail Release 4.1, limitations of the current network policy and security group constructs are addressed by supporting decoupling of routing from security policies, multidimension segmentation, and policy portability. This release also enhances user visibility and analytics functions for security.

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

- Routing and policy decoupling—introducing new firewall policy objects, which decouples policy from routing.
- Multidimension segmentation—segment traffic and add security features, based on multiple dimensions of entities, such as application, tier, deployment, site, usergroup, and so on.
- Policy portability—security policies can be ported to different environments, such as ‘from development to production’, ‘from pci-complaint to production’, ‘to bare metal environment’ and ‘to container environment’.
- Visibility and analytics

Using Tags and Configuration Objects to Enhance Security Policy

Starting with Contrail Release 4.1, tags and configuration objects are used to create new firewall policy objects that decouple routing and network policies, enabling multidimension segmentation and policy portability.

Multidimension traffic segmentation helps you segment traffic based on dimensions such as application, tier, deployment, site, and usergroup.

You can also port security policies to different environments. Portability of policies are enabled by providing match conditions for tags. Match tags must be added to the policy rule to match tag values of source and destination workloads without mentioning tag values. For example, in order for the **‘allow protocol tcp source application-tier=web destination application-tier=application match application and site’** rule to take effect, the application and site values must match.

Predefined Tags

You can choose predefined tags based on the environment and deployment requirements.

Predefined tags include:

- application
- application-tier
- deployment
- site
- label (a special tag that allows the user to label objects)

Custom Tags

You can also define custom tags for a Kubernetes environment. You can define tags in the UI or upload configurations in JSON format.

Example Tag Usage

application = HRApp application-tier = Web site = USA

Tagging Objects

A user can tag the objects project, VN, VM, and VMI with tags and values to map their security requirements. Tags follow the hierarchy of project, VN, VM and VMI and are inherited in that order. This gives an option for the user to provide default settings for any tags at any level. Policies can specify their security in terms of tagged endpoints, in addition to expressing in terms of ip prefix, network, and address groups endpoints.

Policy Application

Policy application is a new object, implemented by means of the application tag. The user can create a list of policies per application to be applied during the flow acceptance evaluation. Introducing global scoped policies and project scoped policies. There are global scoped policies, which can be applied globally for all projects, and project scoped policies, which are applied to specific projects.

Configuration Objects

The following are the configuration objects for the new security features.

- firewall-policy
- firewall-rule
- policy-management

- application-policy
- service-group
- address-group
- tag
- global-application-policy

Configuration Object Tag Object

Each configuration object tag object contains:

- tag: one of the defined tag types, stored as string and a 32-bit ID.
- tag type: Contains the type string and ID (the first 16 bits of the tag) and references to the tag resource type

Each value entered by the user creates a unique ID that is set in the tag_id field. The system can have up to 64 million tag values. On average, each tag can have up to 2k values, but there are no restrictions per tag.

Tags and labels can be attached to any object, for example, project, VN, VM, VMI, and policy, and these objects have a tag reference list to support multiple tags.

RBAC controls the users allowed to modify or remove attached tags. Some tags (typically facts) are attached by the system by default or by means of introspection.

Tag APIs

Tag APIs are used to give RBAC per tag in any object (VMI, VM, Project).

- REST: HTTP POST to /set_tag_<tag_type>/<obj_uuid>
- Python: set_tag_<tag_type> (object_type, object_uuid, tag_value)

Configuration also supports the following APIs:

- tag query
- tags (policy)
- tags (application tag)
- object query
- tags (object)
- tags (type, value)

Label

Label is special tag type, used to assign labels for objects. All of the tag constructs are valid, except that tag type is 'label'. One difference from other tags is that an object can have any number of labels. All other tag types are restricted to one tag per object.

Local and Global Tags

Tags can be defined globally or locally under a project; tag objects are children of either config-root or a project. An object can be tagged with a tag in its project or with a globally-scoped tag.

Analytics

When given a tag query with a SQL where clause and select clause, analytics should give out objects. The query can also contain labels, and the labels can have different operators.

Example:

User might want to know: a list of VMIs where **'site == USA and deployment == Production'**

list of VMIs where **'site == USA and deployment == Production has '**

Given tag SQL where clause and select clause, analytics should give out flows.

Control Node

The control node passes the tags, along with route updates, to agents and other control nodes.

Agent

Agent gets attached tags along with configuration objects. Agent also gets route updates containing tags associated with IP route. This process is similar to getting security group IDs along with the route update.

Address-Group Configuration Object

There are multiple ways to add IP address to address-group.

- Manually add IP prefixes to the address-group by means of configuration.
- Label a work load with the address-group's specified label. All ports that are labelled with the same label are considered to be part of that address-group.
- Use introspect workloads, based on certain criteria, to add ip-address to address-group.

Configuration

The address-group object refers to a label object, description, and list of IP prefixes. The label - object is created using the tag APIs.

Agent

Agent gets address-group and label objects referenced in policy configuration. Agent uses this address group for matching policy rules.

Analytics

When given address group label, analytics gets all the objects associated with it. Given address group label, get all the flows associated with it.

Service-Group Configuration Object

Configuration

The service-group contains a list of ports and protocols. The open stack service-group has a list of service objects; the service object contains attributes: id, name, service group id, protocol, source_port, destination_port, icmp_code, icmp_type, timeout, tenant id.

Agent

Agent gets service-group object as it is referred to in a policy rule. Agent uses this service group during policy evaluation.

Application-policy-set Configuration Object

The application-policy-set configuration object can refer to a tag of type application, network-policy objects, and firewall-policy objects. This object can be local (project) or globally scoped.

When an application tag is attached to an application-policy-set object, the policies referred by that object are automatically applied to the ports that have the same application tag.

Any firewall-policies referred by the application-policy-set objects are ordered using sequence numbers. If the same application tag is attached to multiple application-policy-sets, all those sets will apply, but order among those sets is undefined.

One application-policy-set (called default-policy-application-set) is special in that policies referred by it are applied to all interfaces by default, after applying policies referred to other application-policy-sets.

Upon seeing the application tag for any object, the associated policies are sent to agent. Agent will use this information to find out the list of policies to be applied and their sequence during flow evaluation. User can attach application tag to allowed objects (Project, VN, VM or VMI).

Policy-management Configuration Object

Policy-management is a global container object for all policy-related configuration.

Policy-management object contains

- network-policies (NPs)
- firewall-policies (FWPs)
- application-policy-sets

- global-policy objects
- global-policy-apply objects
- NPs - List of contrail networking policy objects
- FWPs - List of new firewall policy objects
- Application-policies - List of Application-policy objects
- Global-policies - List of new firewall policy objects, that are defined for global access
- Global-policy-apply - List of global policies in a sequence, and these policies applied during flow evaluation.
- Network Policies (NP) references are available, as they are today.

Firewall-policy Configuration Object

Firewall-policy is a new policy object that contains a list of firewall-rule-objects and audited flag.

Firewall-policy can be project or global scoped depending on usage. Includes an audited Boolean flag to indicate that the owner of the policy indicated that the policy is audited. Default is False, and will have to explicitly be set to True after review. Generates a log event for audited with timestamp and user details.

Firewall-rule Configuration Object

Firewall-rule is a new rule object, which contains the following fields. The syntax is to give information about their layout inside the rule.

- <sequence number>
There is a string object sequence number on the link from firewall-policy to firewall-policy-rule objects.
The sequence number decides the order in which the rules are applied.
- [< id >]
uuid
- [name < name >]
Unique name selected by user
- [description < description >]
- public
- {permit | deny}
- [protocol {< protocol-name > | any } destination-port { < port range > | any } [source-port { < port range > | any }]] | service-group < name >
- endpoint-1 { [ip < prefix >] | [virtual-network < vnname >] | [address-group < group name >] | [tags T1 == V1 && T2 == V2 ... && Tn == Vn && label == label name...] | any }
- { -> | <- | <-> }

Specifies connection direction. All the rules are connection oriented and this option gives the direction of the connection.

- endpoint-2 { [ip < prefix >] | [virtual-network < vnname >] | [address-group < group name >] | [tags T1 == V1 && T2 == V2 ... && Tn == Vn && label == label name...] | any }

Tags at endpoints support an expression of tags. We support only '=' and '&&' operators. User can specify labels also as part the expression. Configuration object contains list of tag names (or global:tag-name in case of global tags) for endpoints.

- [match_tags {T1 Tn} | none }

List of tag types or none. User can specify either match with list of tags or none. Match with list of tags mean, source and destination tag values should match for the rule to take effect.

- [log | mirror | alert | activate | drop | reject | sdrop]

complex actions

- { enable | disable }

A boolean flag to indicate the rule is enabled or disabled. Facilitates selectively turn off the rules, without remove the rule from the policy. Default is True.

- filter

Compilation of Rules

Whenever the API server receives a request to create/update a firewall policy rule object, it analyzes the object data to make sure that all virtual-networks, address-group, tag objects exist. If any of them do not exist, the request will be rejected. In addition, it will actually create a reference to those objects mentioned in the two endpoints. This achieves two purposes. First, we don't allow users to name non-existent objects in the rule and second, the user is not allowed to delete those objects without first removing them from all rules that are referring to them.

Using the Contrail Web User Interface to Manage Security Policies

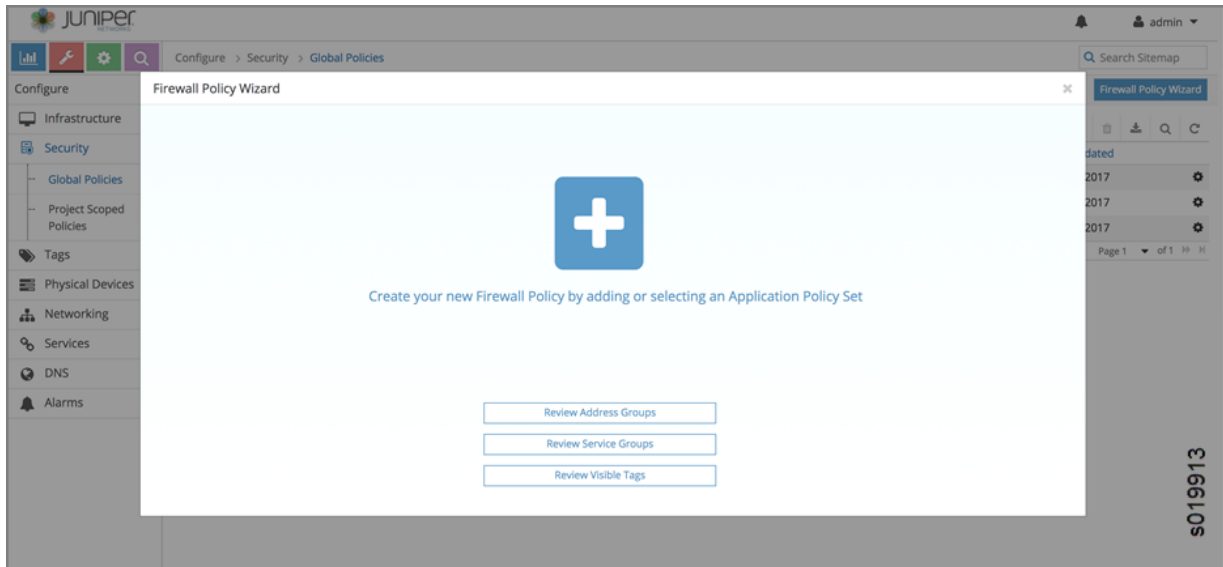
IN THIS SECTION

- [Adding Security Policies | 23](#)
- [Managing Policy Tags | 26](#)
- [Viewing Global Policies | 27](#)
- [Visualizing Traffic Groups | 28](#)

Adding Security Policies

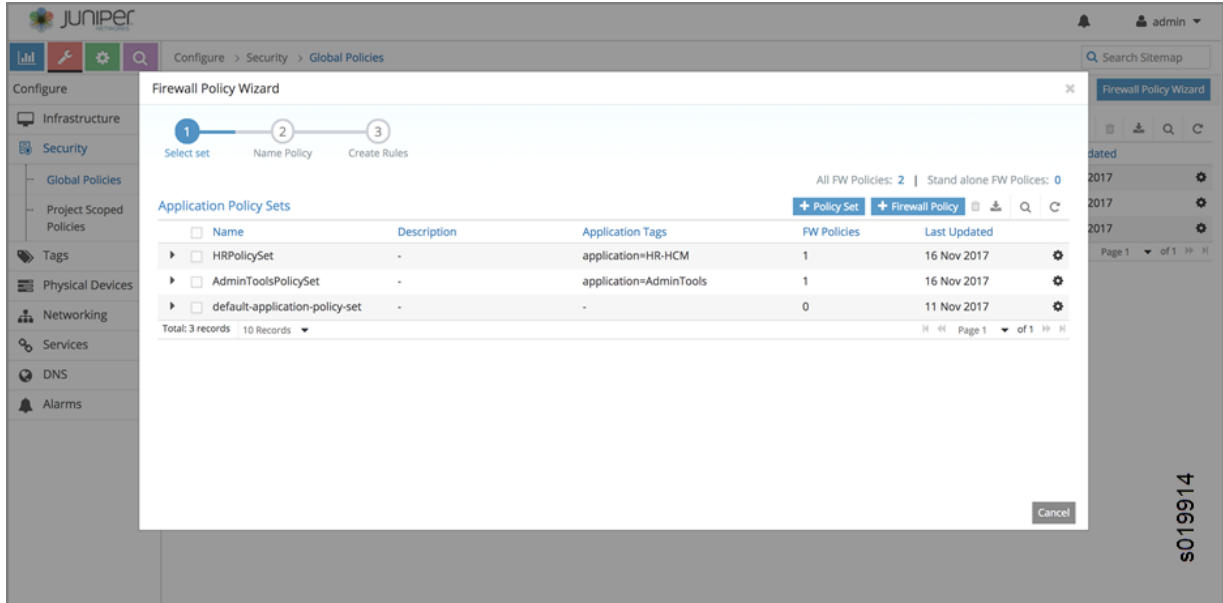
1. To add a security policy, go to **Configure > Security > Global Policies**. Near the upper right, click the button **Firewall Policy Wizard**. The **Firewall Policy Wizard** appears, where you can create your new firewall policy by adding or selecting an application policy set. See [Figure 1 on page 23](#).

Figure 1: Firewall Policy Wizard




2. Click the large + on the Firewall Policy Wizard screen to view the **Application Policy Sets** window. The existing application policy sets are displayed. See [Figure 2 on page 24](#).

Figure 2: Application Policy Sets



- To create a new firewall policy, click the application policy set in the list to which the new firewall policy will belong. The **Edit Application Policy Sets** window appears, displaying a field for the description of the selected policy set and listing firewall policies associated with the set. See [Figure 3 on page 24](#), where the **HRPolicySet** has been selected.

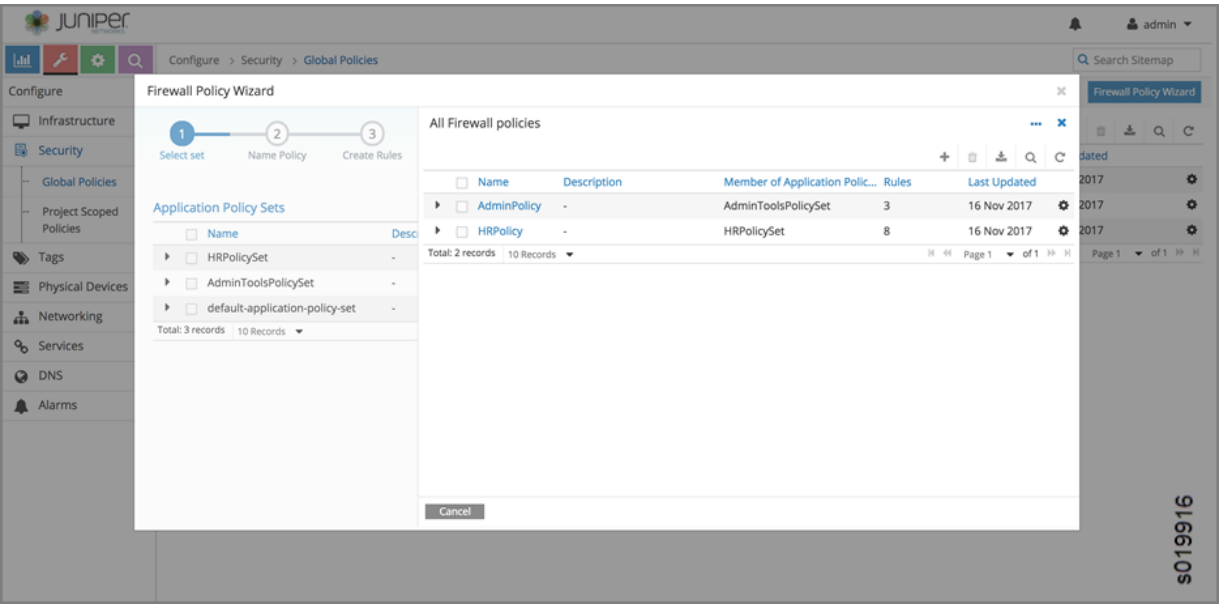
Figure 3: Edit Application Policy Sets



- To view all firewall policies, click the Application Policy Sets link in the left side.

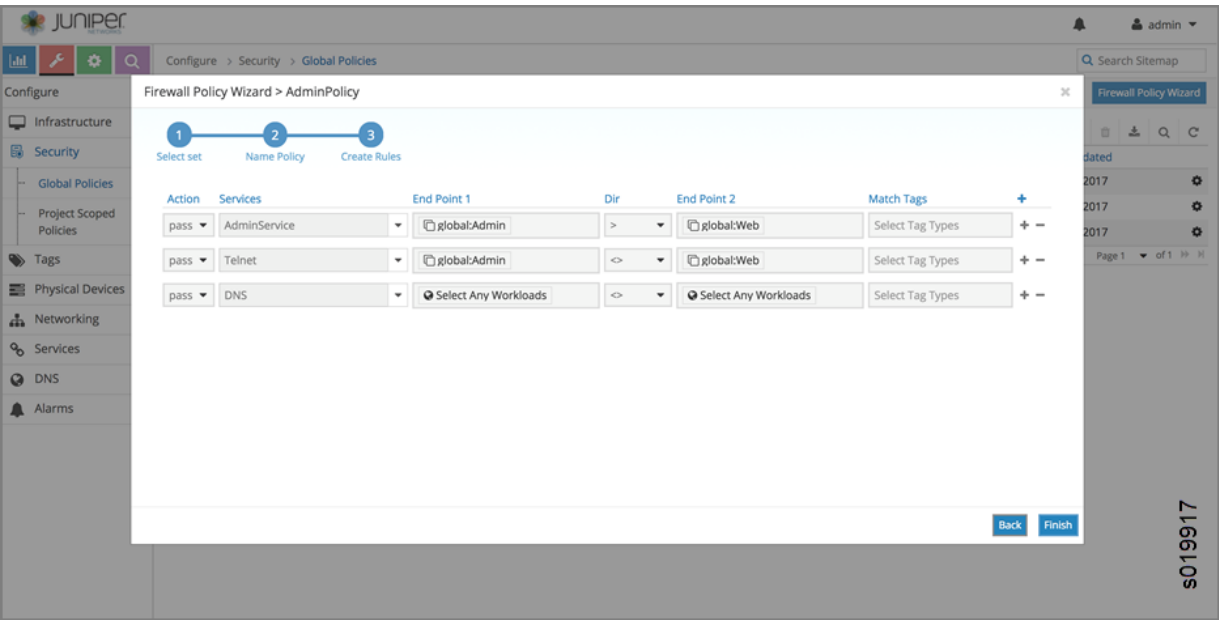
See [Figure 4 on page 25](#).

Figure 4: All Firewall Policies



5. Select any listed firewall policy to view or edit the rules associated with that policy. See [Figure 5 on page 25](#), where all the rules for the **AdminPolicy** are listed. Use the dropdown menus in each field to add or change policy rules, and use the +, - icons to the right of each rule to add or delete the rule.

Figure 5: Firewall Policy Rules



Managing Policy Tags

You can use the Contrail web user interface to create and manage the tags used to provide granularity to security policies. You can have global tags, applicable to the entire system, or project tags, defined for specific uses in specific projects.

1. To manage policy tags, go to **Configure > Tags > Global Tags**. The **Tags** window appears, listing all of the tags in use in the system, with the associated virtual networks, ports, and projects for each tag. Tags are defined first by type, such as application, deployment, site, tier, and so on. See [Figure 6 on page 26](#).

Figure 6: Tags

Name	Associated Virtual Networks	Associated Ports	Associated Projects
application=AdminTools	-	ip-fabric (-)	-
application=HR-HCM	-	-	HRProd HRDev
deployment=Development	-	-	HRDev
deployment=Production	-	-	HRProd
site=France	-	-	HRProd
site=USA	-	-	HRDev
tier=Admin	-	ip-fabric (-)	-
tier=App	appvn (HRProd) appvn (HRDev)	-	-
tier=DB	dbvn (HRDev) dbvn (HRProd)	-	-
tier=Web	webvn (HRProd) webvn (HRDev)	-	-

Total: 10 records 10 Records

Page 1 of 1

2. You can click through any listed tag to see the rules to which the tag is applied. See [Figure 7 on page 27](#), which shows the application tags that are applied to the current application sets. You can also reach this page from **Configure > Security > Global Policies**.

Figure 7: View Application Tags

Name	Description	Application Tags	PW Policies	Last Updated
HRPolySet	-	application=HR-HCM	1	16 Nov 2017
AdminToolsPolicySet	-	application=AdminTools	1	16 Nov 2017
default-application-policy-set	-	-	0	11 Nov 2017

Total: 3 records 50 Records

Viewing Global Policies

From **Configure > Security > Global Policies**, in addition to viewing the policies included in application policy sets, you can also view all firewall policies, all service groups policies, and all address groups policies.

1. To view and manage the global firewall policies, from **Configure > Security > Global Policies**, click the **Firewall Policies** tab to view the details for system firewall policies, see [Figure 8 on page 27](#)

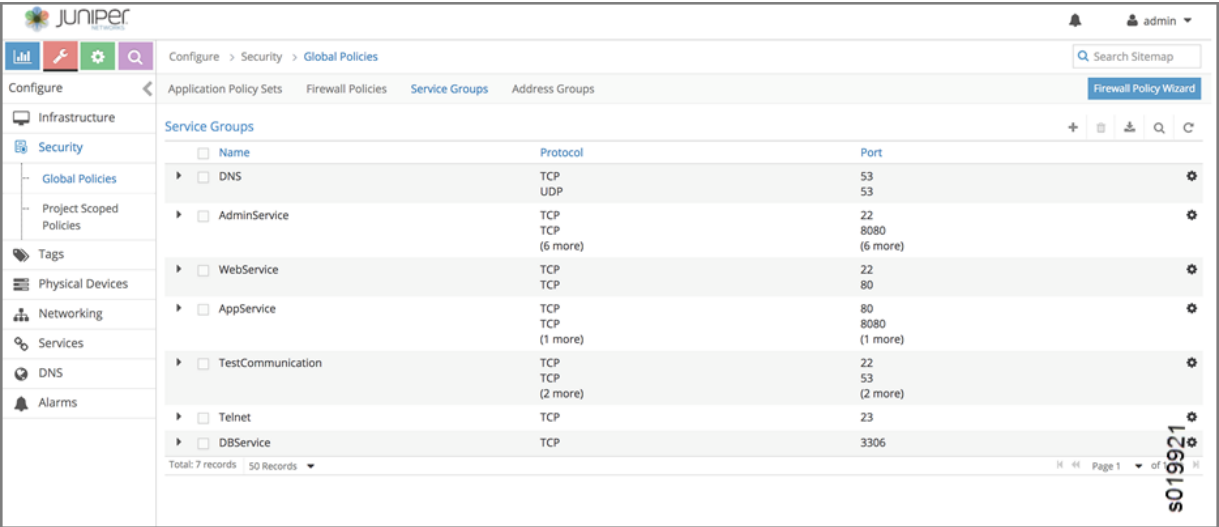
Figure 8: Firewall Policies

Name	Description	Member of Application Policy Sets	Rules	Last Updated
AdminPolicy	-	AdminToolsPolicySet	3	16 Nov 2017
HRPoly	-	HRPolicySet	8	16 Nov 2017

Total: 2 records 50 Records

2. To view and manage the service groups policies, from **Configure > Security > Global Policies**, click the **Service Groups** tab to view the details for system policies for service groups, see [Figure 9 on page 28](#).

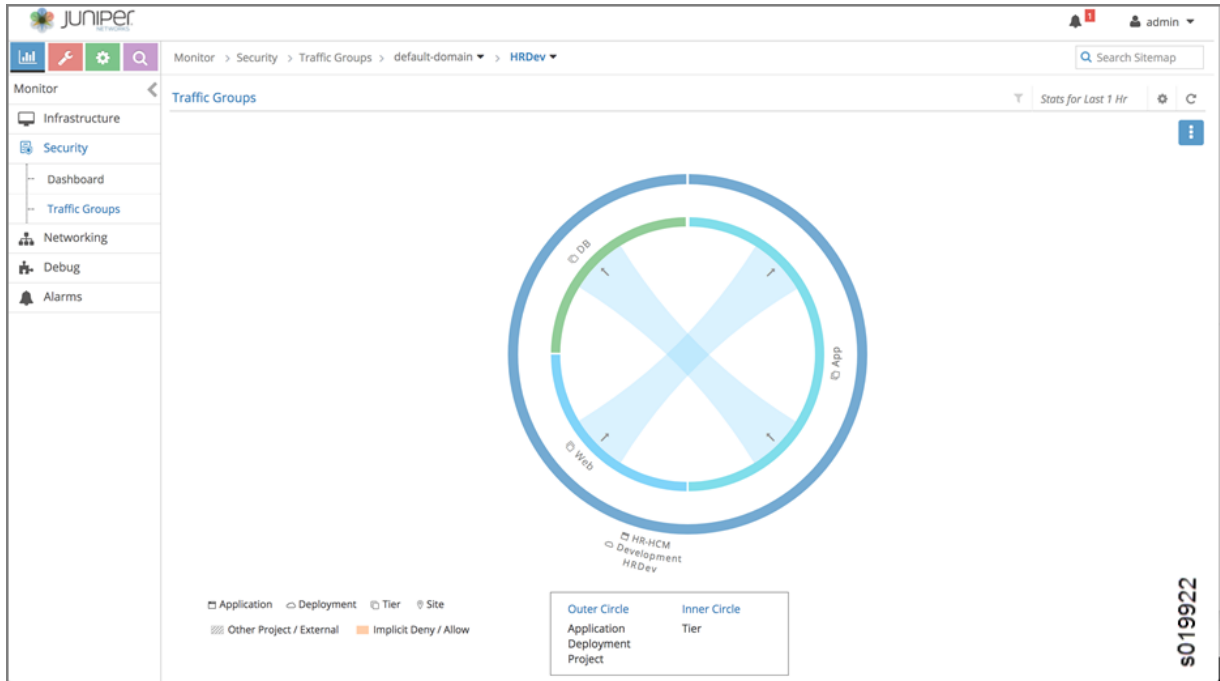
Figure 9: Service Groups



Visualizing Traffic Groups

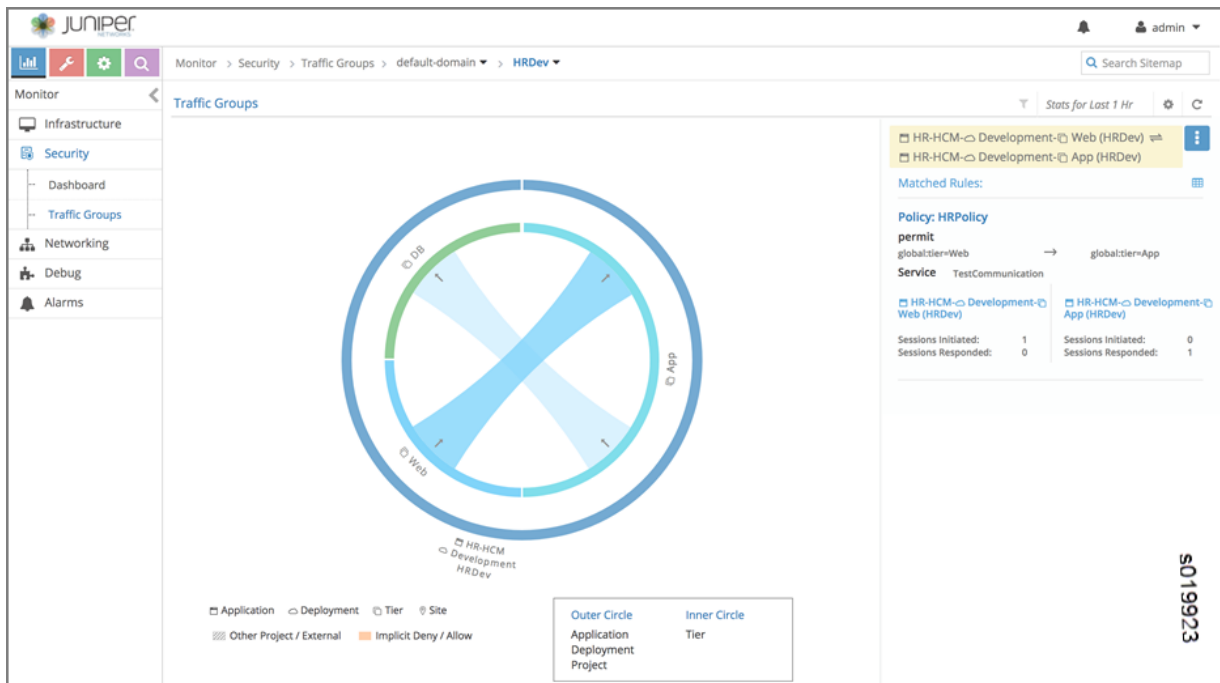
Use **Monitor > Security > Traffic Groups** to explore visual representations of how policies are applied to traffic groups. See [Figure 10 on page 29](#), which is a visual representation of the source and destination traffic for the past one hour of a traffic group named Traffic Groups. The outer circle represents traffic tagged with application, deployment, or project. The inner circle represents traffic tagged with tier. The center of the circle shows the traffic origination and destination.

Figure 10: Traffic Groups



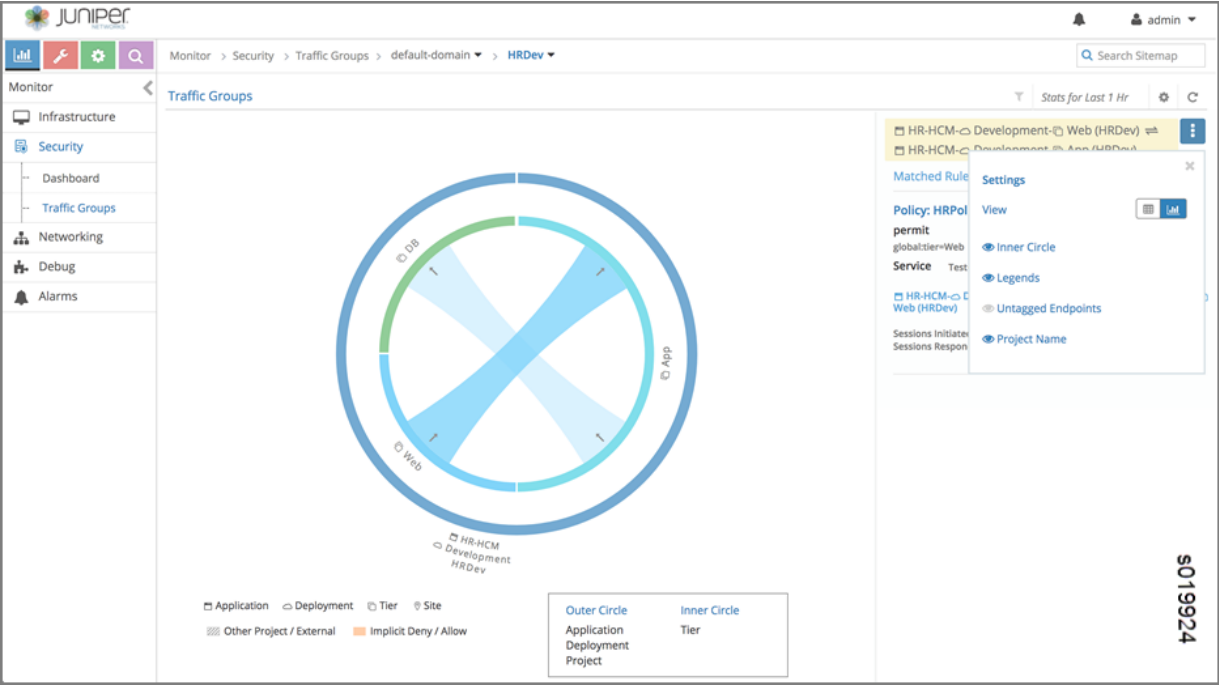
You can click in the right side of the screen to get details of the policy rules that have been matched by the selected traffic. See [Figure 11 on page 29](#).

Figure 11: Traffic Groups, Policy Details



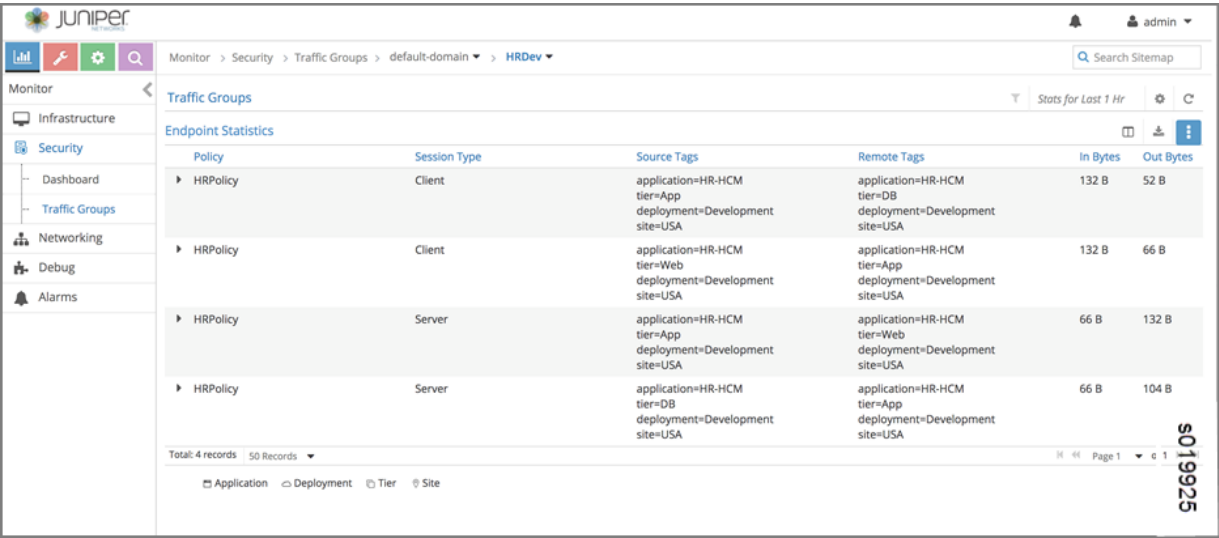
You can click in the right side of the screen to get to the **Settings** window, where you can change the type of view and change which items appear in the visual representation. See [Figure 12 on page 30](#).

Figure 12: Traffic Groups, Settings



You can click on the name of a policy that has been matched to view the endpoint statistics, including source tags and remote tags, of the traffic currently represented in the visual. See [Figure 13 on page 30](#).

Figure 13: Traffic Groups, Endpoint Statistics



You can click deeper through any linked statistic to view more details about that statistic, see [Figure 15 on page 31](#) and [Figure 15 on page 31](#).

Figure 14: Traffic Groups, Details

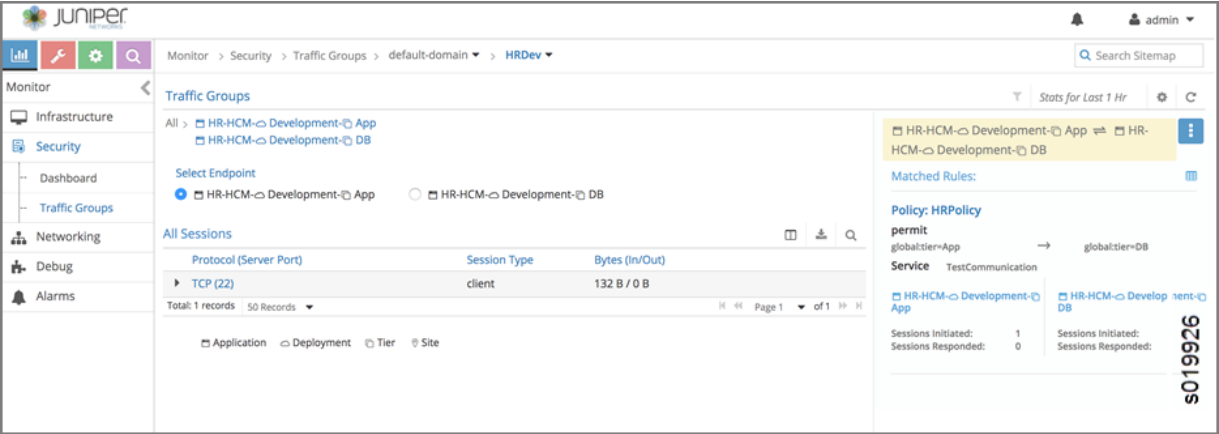
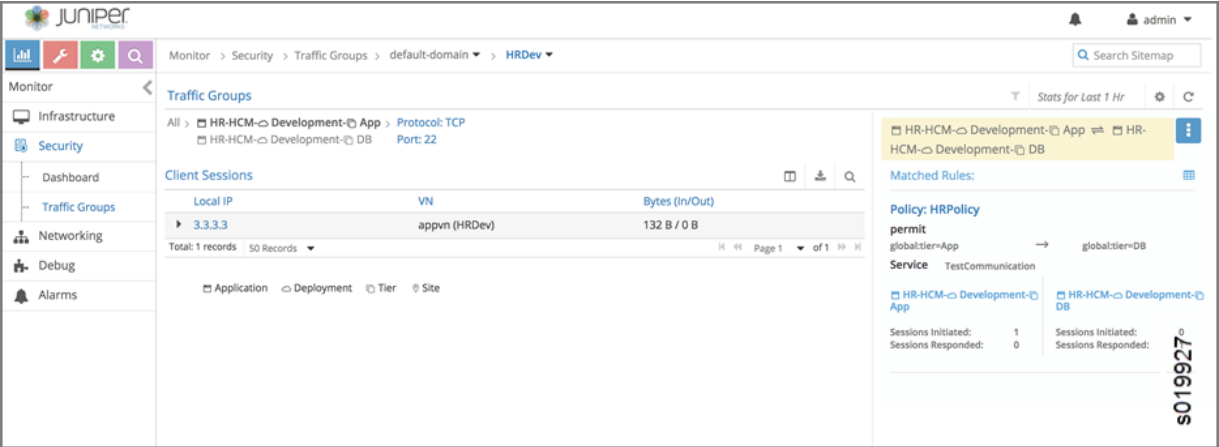
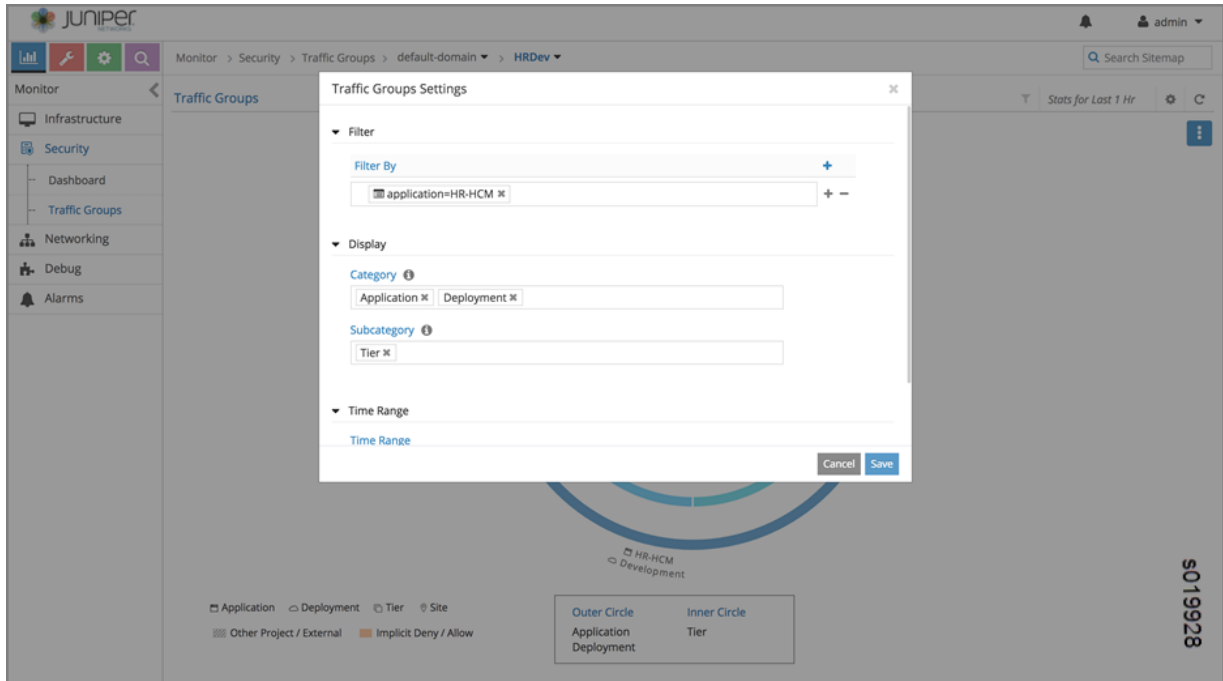


Figure 15: Traffic Groups, Details



You can change the settings of what statistics are displayed in each traffic group at the **Traffic Groups Settings** screen see [Figure 16 on page 32](#).

Figure 16: Traffic Groups Settings



RELATED DOCUMENTATION

[Security Policy Features in OpenStack | 32](#)

Security Policy Features in OpenStack

IN THIS SECTION

- [Overview of Existing Network Policy and Security Groups in OpenStack | 33](#)
- [Security Policy Enhancements | 34](#)

Overview of Existing Network Policy and Security Groups in OpenStack

Contrail virtual networks are isolated by default. Workloads in a virtual network cannot communicate with workloads in other virtual networks, by default. A neutron router or a Contrail network policy may be used to connect two virtual networks. In addition, Contrail network policy also provides security between two virtual networks by allowing or denying specified traffic.

OpenStack security groups allow access between workloads and instances for specified traffic types and any other types are denied.

A security policy model for any given customer first needs to map to the OpenStack network policy framework and security group constructs.

In modern cloud environments, workloads are moving from one server to another, one rack to another and so on. Therefore, users must rely less on using IP addresses or other network coordinates to identify the endpoints to be protected. Instead users must leverage application attributes to author policies, so that the policies don't need to be updated on account of workload mobility.

You might want to segregate traffic based on the different categories of data origination, such as:

- Protecting the application itself
- Segregating traffic for specific component tiers within the application
- Segregating traffic based on the deployment environment for the application instance
- Segregating traffic based on the specific geographic location where the application is deployed

There are many other possible scenarios where traffic needs to be segregated.

Additionally, you might need to group workloads based on combinations of tags. These intents are hard to express with existing network policy constructs or Security Group constructs. Besides, existing policy constructs leveraging the network coordinates, must continually be rewritten or updated each time workloads move and their associated network coordinates change.

Security Policy Enhancements

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

Starting with Release 5.1, Contrail Networking supports the OpenStack Neutron Firewall version 2 API extension known as Neutron FWaaS (Firewall as a Service). The Neutron API enhancements make the existing FWaaS more granular by giving you the ability to apply the firewall rules at the Layer 2 port level rather than at the Layer 3 router level, and to have different firewall policies with different rules applied to inbound versus outbound connections. Support is extended to various types of Neutron ports, including VM ports and SFC ports. This feature introduces the Firewall Group resource which enables firewall policies to be bound to Neutron ports instead of being bound tenant wide or on a set of routers in a tenant. The firewall groups consists of references to two firewall policies, one for ingress and one for egress traffic. The Firewall Group resource also references a collection of Neutron ports where the policies are applied. Firewall policies can be shared among multiple groups. The firewall policies are an ordered list of firewall rules. The firewall rules contain the filter specifications for the traffic including information such as source, destination, port, protocol, action, and so on. Firewall rules can be shared among multiple policies.

Related enhancements to the OpenStack Neutron and Contrail security groups API include:

- Firewall rules support deny, description, and admin status attributes
- A share attribute for firewall rules allow them to be shared between different projects
- Filtering based on the source and destination address prefix and port rather than just the remote destination
- Firewall groups reference firewall rules through a firewall policy, allowing reuse of shareable firewall policies that are referenced by multiple firewall groups

Limitations of the Neutron FWaaS feature include:

- Contrail Security does not distinguish between ingress and egress traffic, so ingress and egress firewall policy collections are always identical. Applying different policies is unsupported and error is displayed if you try.

Firewall rules **reject** action is not supported by Contrail Security.

Flexibility in Contrail Security flexibility enables you to apply multiple firewall groups to the same port.

As Contrail Networking does not really create an L3 port, we cannot apply firewall group to an L3 port

Firewall groups can be set as source or destination for a firewall rule

Configuration Objects

The following are the configuration objects for the new security features.

- firewall-policy
- firewall-rule
- policy-management
- application-policy
- service-group
- address-group
- tag
- global-application-policy

For more information on security policies in Contrail, see [“Security Policy Features” on page 15](#).

Configure Neutron as FWaaS

Consider an all-in-one setup, with OpenStack and Contrail installed.

1.

RELATED DOCUMENTATION

| [Security Policy Features](#) | 15

Policy Generation

The policy generation feature in Contrail Release 5.1 automates the task of policy creation based on observed traffic flows. Contrail creates and enforces intent-based policies. In many cases, Contrail Security is deployed in brownfield environments, in which inter-applications and intra-application traffic policies are pre-existing. However, in greenfield deployments and in complex environments, where many applications are communicating internally and externally, creating policies one-by-one is a tedious and time consuming task. Sometimes, manually created policies do not perform as per expectations in real-time traffic or sometimes you might create extra policies which are never used by the applications. The policy generation feature simplifies this process of creating policies by automating the generation of policies based on application communication.

The policy generation feature aids in the creation of policies based on observed traffic, without enforcing any new policies. In order to generate policies, workloads VMs or Containers need to be grouped within Contrail objects like virtual networks and Projects. Subsequently, tags must be created and associated with Projects, virtual networks or ports. In the policy generation mode, traffic from selected applications is allowed to pass for a selected period of time. The vRouter observes and forwards all traffic between the selected applications because the implicit rule is to allow all traffic to pass. On the basis of this observation, the vRouter generates a draft policy which is saved in the policy draft mode. You can review the draft policy and edit it as required before enforcing the policy. The policy generation feature significantly reduces the burden of policy creation from scratch.

To use the policy generation feature, the sequence of high level steps are as listed here:

- Create tags.
- Associate tags with projects, virtual networks (VNs) or ports.
- Run traffic.
- Edit the generated policies available in draft mode.
- Commit the (optionally) edited policies to enforce them

RELATED DOCUMENTATION

| [Configuring Policy Generation](#) | 37

Configuring Policy Generation

To configure policy generation:

1. Create tags.

In the Command UI, navigate to **Security > Tags**. Tags are key and value pairs. Create tags as appropriate for your environment. Tags can be created with a Project scope or a Global scope.

2. Associate tags.

Tags created can be associated with either individual virtual machines (VMs) or container ports or with groups of virtual machines or containers at either the VN level or the project level. Tags associated with a VN are inherited by all VMs or containers in that VN. Similarly, tags associated with a project get inherited by all VNs in that project and in turn by all VMs or containers in each VN in that project. To associate tags at any of these levels, perform the following steps:

- Project

Navigate to the **Projects** tab, double click on the project and associate the tags.

- Virtual Network

Navigate to **Overlay > Virtual Networks** and click **Edit** to add appropriate tags.

- Ports

Navigate to **Overlay > ports** and click **Edit** to add appropriate tags.

3. When running policy generation for the first time, you must provision the policy generator module using the following commands:

- a. Download the Contrail Security Apps tarball from the [Support - Software Downloads](#) site. Untar the .tgz file

```
untar contrail-security-apps-*.tgz
```

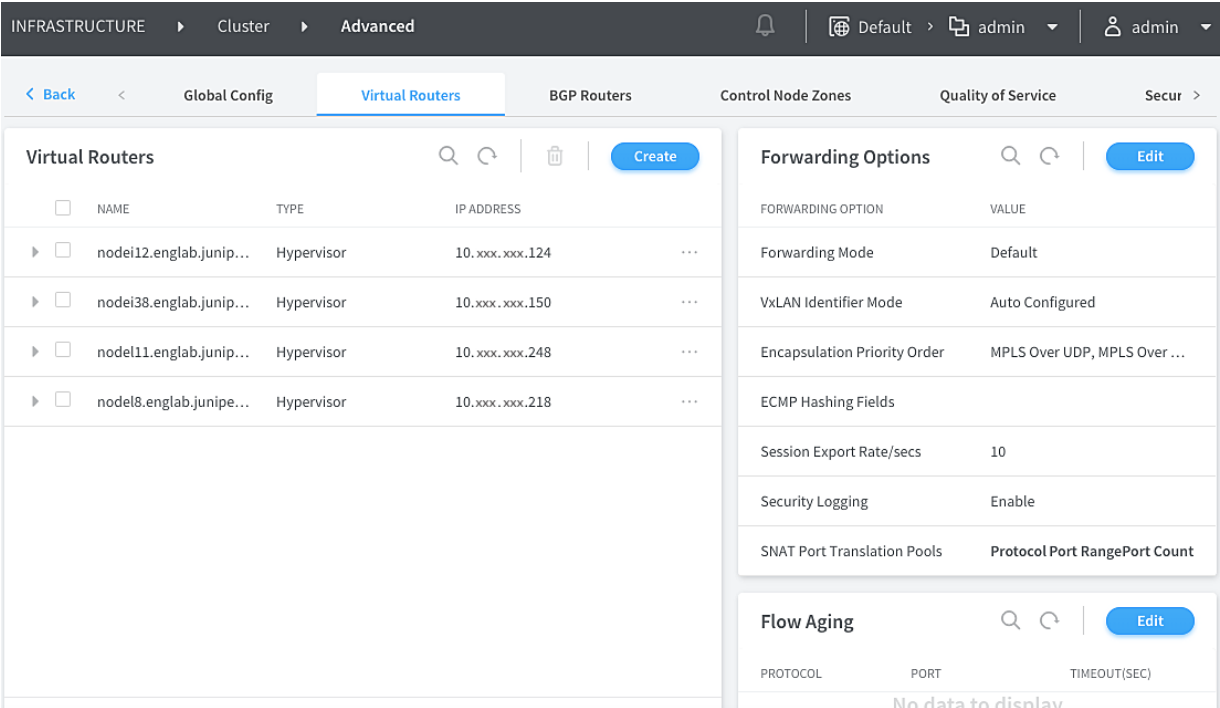
- b. `cd contrail-security-apps`

- c. `edit ansible/inventory/inventory.yml` file and specify the required values. For a sample of the `inventory.yml` file, see [Sample inventory.yml file on page 43](#).

- d. `ansible-playbook -i ansible/inventory/inventory.yml ansible/playbooks/deploy_and_run_all.yml`

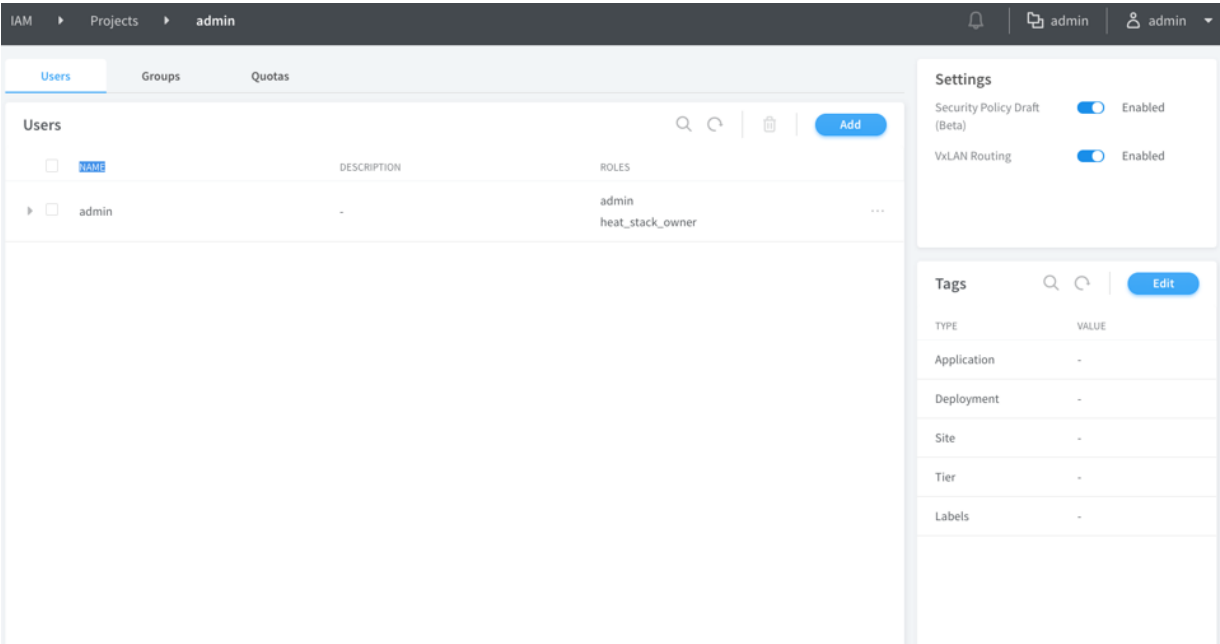
4. Specify the session export rate. Navigate to **INFRASTRUCTURE > Cluster > Advanced Options**.

Click the **Virtual Routers** tab, click **Edit** under **Forwarding Options** and enter the **Session Export Rate/secs** value.



5. Enable the security policy draft mode, either for the Project scope or for Global scope depending on your requirement.

For Project scope, navigate to **IAM > Projects**, select and click the project and enable the **Security Policy Draft** mode under **Settings**.



For Global scope, navigate to **INFRASTRUCTURE > Cluster > Advanced Options**, and click the **Global Config** tab. The **Edit System Configuration** page appears. Click **Edit** and enable the **Security Policy Draft** mode. Click **Save**.

6. Enable policy generation endpoint.

Navigate to the **INFRASTRUCTURE > Cluster > Advanced Options**, click the **Endpoints** tab, and click **Create**. The **Create Endpoint** page appears.

7. Enter **generation** under **Prefix** and enter the required URLs. Click **Create** to save the endpoint.

The screenshot shows the 'Create Endpoint' page within the 'INFRASTRUCTURE > Cluster > Advanced' navigation path. The page contains a form with the following fields and controls:

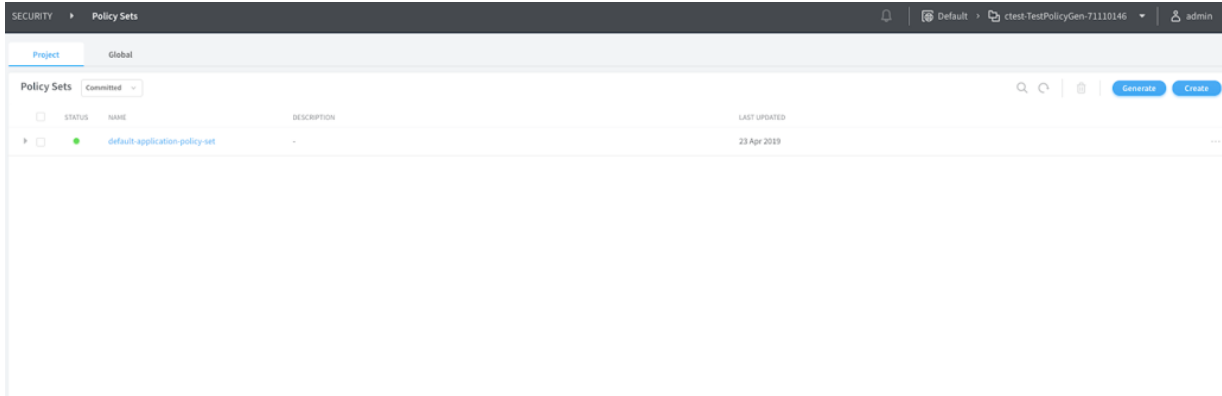
- Prefix:** A text input field containing the value 'generation'.
- Private url:** A text input field containing the value 'http://10.xxx.xxx.108:9093/'.
- Public url:** A text input field containing the value 'http://10.xxx.xxx.108:9093/'.
- User Name:** An empty text input field.
- Password:** An empty text input field.
- Enable proxy:** A checkbox that is currently unchecked.
- + Add:** A blue link to add a new endpoint.
- Create:** A blue button to save the endpoint.
- Cancel:** A light blue button to cancel the operation.

8. Generate traffic between the applications.

9. Generate policies.

Navigate to **Security > Policy Sets** and click **Generate Policy**. The **Generate Project Policy** page appears.

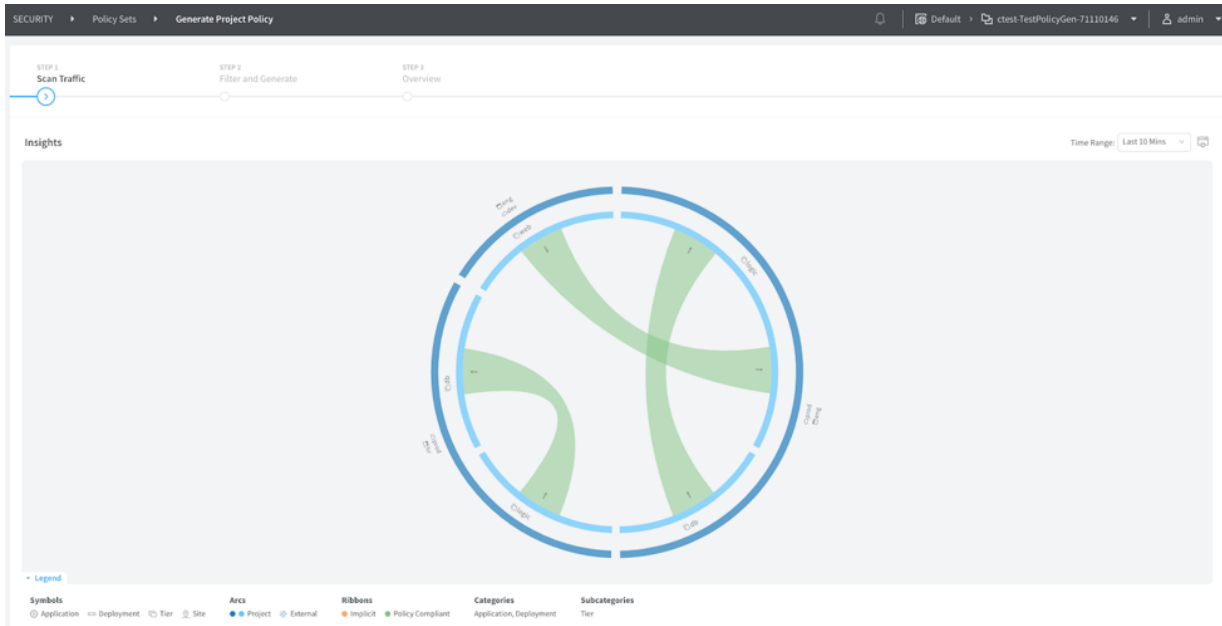
The **Generate Project Policy** page has three steps.



a. Step 1 Scan Traffic

Each vRouter scans the traffic it sees between and within applications. The controller analyzes the observed traffic patterns and displays the observed traffic in a graphical visualization. Arcs inside the circular graph represent the different observed flows. Mouse over the arcs to view additional details about the applications involved in that flow, other tags associated with the endpoints, and other flow characteristics.

The period of traffic considered as an input for policy generation can be customized by editing the **Time Range**. The default time range is 10 minutes. Click **Next** to proceed with policy generation.



b. Step 2 Filter and Generate

You can filter traffic and generate a draft policy based on the selected filters. By default, the current project is selected and the predefined tags, application, deployment, tier, and site are selected. Selected tags must be associated with the workloads, because not having these tags in flow records creates unknown flows.

If you haven't attached any of the predefined tags to workloads and you don't need the tags to be part of the policies, deselect them. However, it is mandatory to select at least the application tag, else flows designated with unknown are displayed.

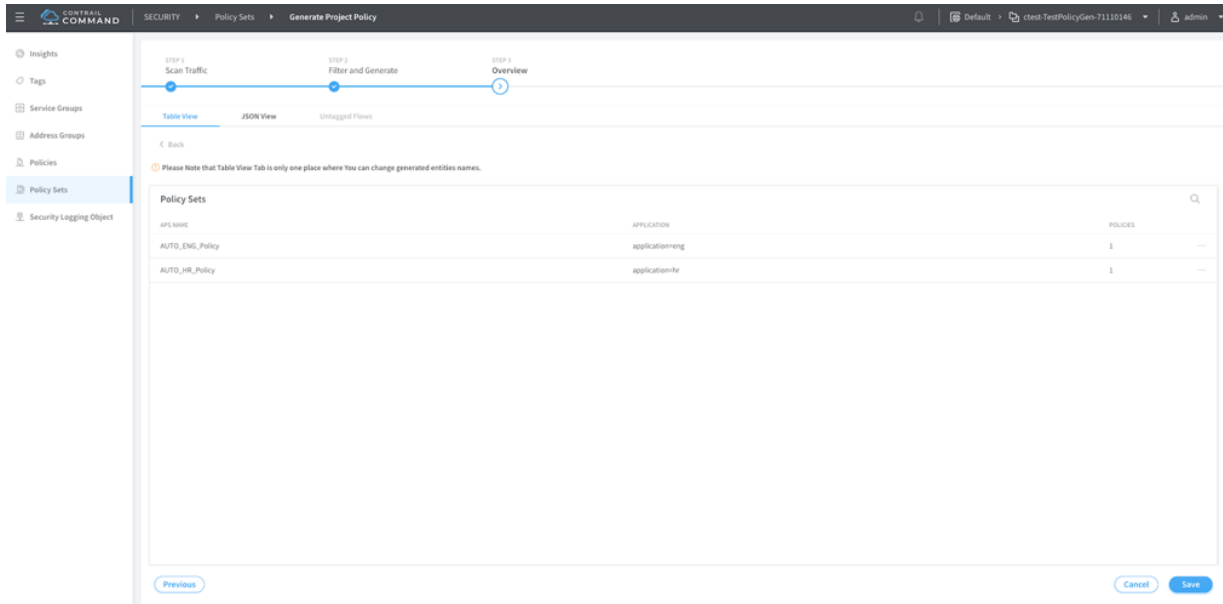
Click **Next**. The default firewall rules allow all traffic.

The screenshot shows the 'Generate Project Policy' wizard in the 'Filter and Generate' step. The breadcrumb trail at the top is 'SECURITY > Policy Sets > Generate Project Policy'. The top bar shows 'Default', 'ctest-TestPolicyGen-71110146', and 'admin'. The progress bar indicates three steps: 'STEP 1 Scan Traffic', 'STEP 2 Filter and Generate' (current), and 'STEP 3 Overview'. Under 'Filter By', there are three dropdown menus: 'Project' (selected: 'ctest-TestPolicyGen-7...'), 'Application', and 'Firewall Rules'. Under 'Select', there are checkboxes for 'Application' (checked), 'Site' (checked), 'Tier' (checked), 'Deployment' (checked), 'Local IP' (unchecked), and 'Remote IP' (unchecked). At the bottom, there are 'Previous', 'Cancel', and 'Next' buttons.

c. Step 3 Overview

You can view the application policy sets based on the selected input parameters. You can view the application policy sets in a tabular format as well as JSON format. You can also view traffic flows for untagged applications.

Click **Save** to save the application policy set and generate the draft policy.



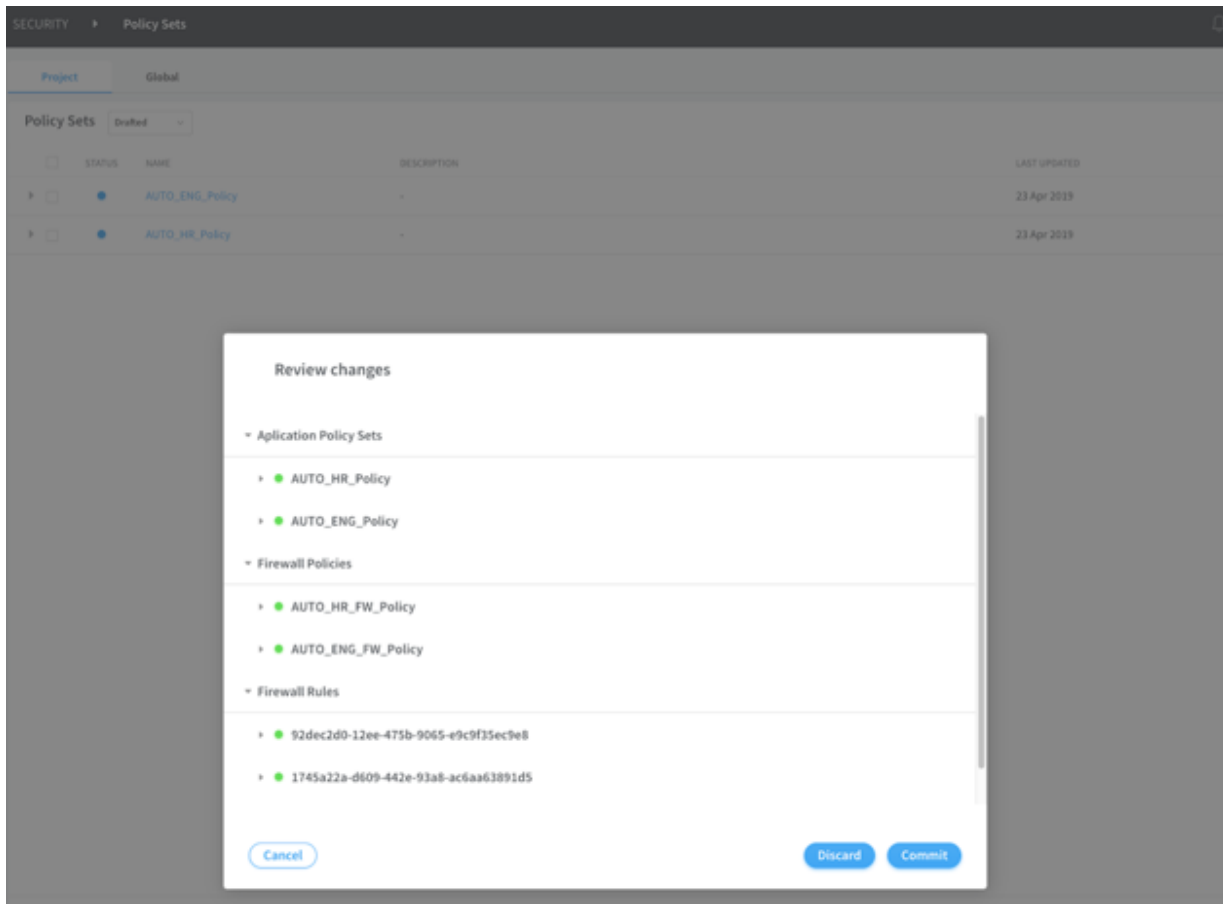
10. Review the draft policy.

The **Security > Policy Sets** page is displayed with the draft application policy set. Click the draft application policy set to view details about the policies.

Select the draft application policy set and click **Review**.

11. The **Review Changes** page appears listing the policies in the draft mode. You can review the draft policy, edit it as required and click **Commit** to enforce the application policy set.

Alternatively, click **Discard** to discard the generated application policy set.



Sample inventory.yml file

```
all:
  hosts:
    localhost:
      ansible_connection: local

  vars:
    PATH_POLICY_GENERATION_CONFIG: '/etc/contrail/securityapps/'
    # Policy generator API port
    rest_api_port: {{ pg_port }}
    # Policy generator API server ip address
    rest_api_ip: {{ pg_ip }}
    # Policy generator logs and log level
    log_file: '/var/log/contrail/policy-generation-server.log'
    log_console_level: logging.DEBUG

    # Registry inventory
    CONTAINER_REGISTRY: {{ pg_registry }}
```

```

CONTRAIL_VERSION: {{ pg_version }}
ansible_user: 'root'

# Config API server
api_server_listen_port: {{ config_port }}
api_server_ip: {{ config_ip }}

# Analytics IP & Port 8081 for remote host, 8181 for localhost
analytics_ip: {{ analytics_ip }}
admin_port: {{ analytics_port }}

# Keystone credentials with user
admin_password: {{ admin_password }}
admin_tenant_name: {{ admin_tenant }}
admin_user: {{ admin_user }}
auth_host: {{ auth_host }}
auth_port: {{ auth_port }}
auth_protocol: {{ auth_protocol }}
auth_version: {{ auth_version }}
auth_type: password
user_domain_name: Default
project_domain_name: Default
region_name: RegionOne

```

RELATED DOCUMENTATION

[Policy Generation](#) | 36

2

CHAPTER

Configuring Virtual Networks

Creating a Virtual Network with Juniper Networks Contrail | **46**

Creating a Floating IP Address Pool | **49**

Support for IPv6 Networks in Contrail | **51**

Configuring EVPN and VXLAN | **55**

Support for EVPN Route Type 5 | **66**

Support for EVPN Type 6 Selective Multicast Ethernet Tag Route | **67**

Support for L3VPN Inter AS Option C | **68**

Creating a Virtual Network with 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. You need to create an IP address management (IPAM) for your project for to create a virtual network. Select **Configure > Networking > IP Address Management**, then click the **Create** button.

The **Add IP Address Management** window appears, see [Figure 17 on page 46](#).

Figure 17: Add IP Address Management

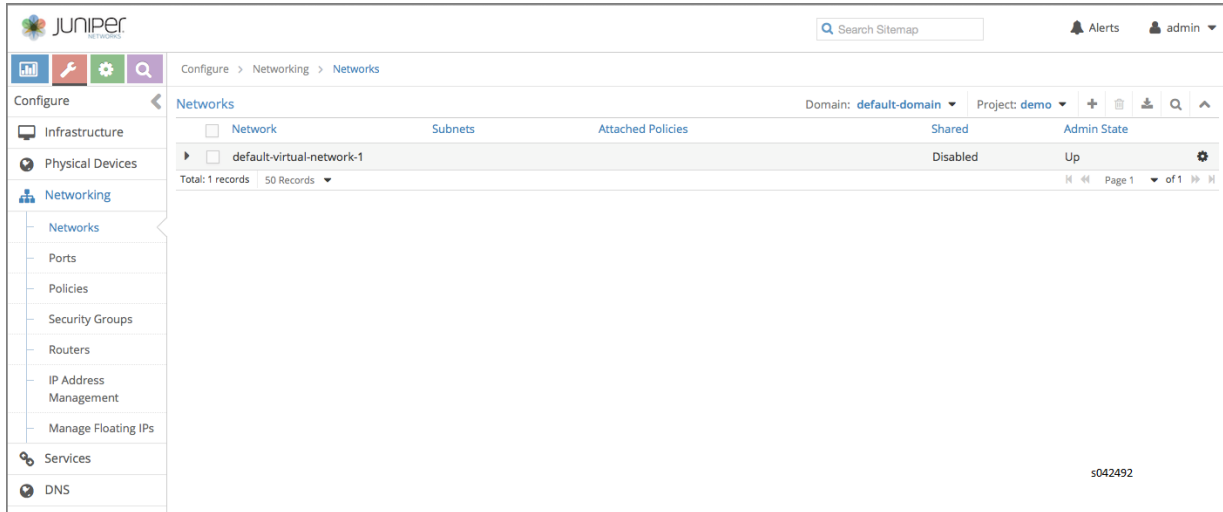
2. Complete the fields in **Add IP Address Management**: The fields are described in [Table 4 on page 46](#).

Table 4: Add IP Address Management Fields

Field	Description
Name	Enter a name for the IPAM you are creating.
DNS Method	Select from a list the domain name server method for this IPAM: Default, Virtual DNS, Tenant, or None.
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** page; see [Figure 18 on page 47](#).

Figure 18: 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 19 on page 47](#). Use the scroll bar to access all sections of this window.

Figure 19: 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 5 on page 48](#).

Table 5: Create Network Fields

Field	Description
Name	Enter a name for the virtual network you are creating.
Network Policy	Select the policy to be applied to this network from the list of available policies. You can select more than one policy by clicking each one needed.
Subnets	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. Alternatively, click the - icon to remove the selections.
Host Routes	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.
Advanced Options	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.
Floating IP Pools	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.
Route Target	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.

6. To save your network, click the **Save** button, or click **Cancel** to discard your work and start over.

RELATED DOCUMENTATION

Creating an Image for a Project in OpenStack Contrail

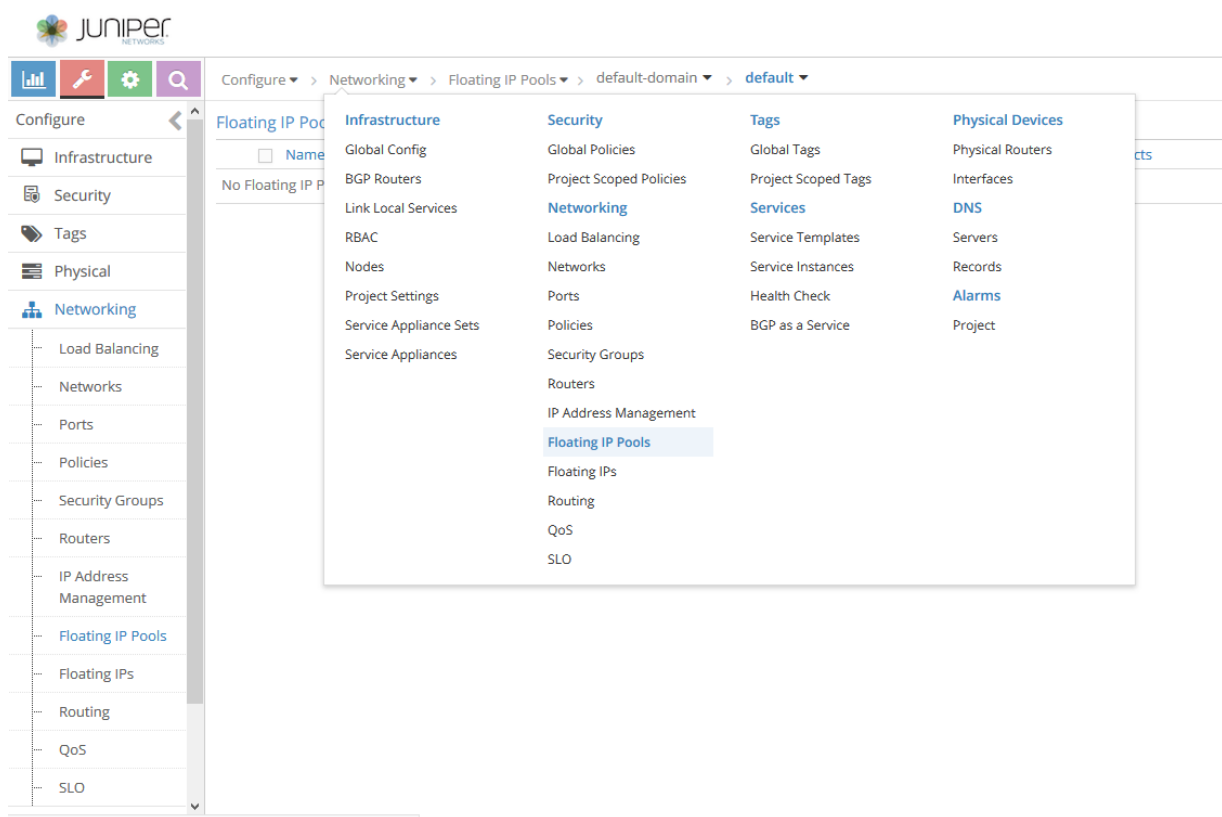
Creating a Floating IP Address Pool

A floating IP address is an IP address (typically public) that can be dynamically assigned to a running virtual instance.

To 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:

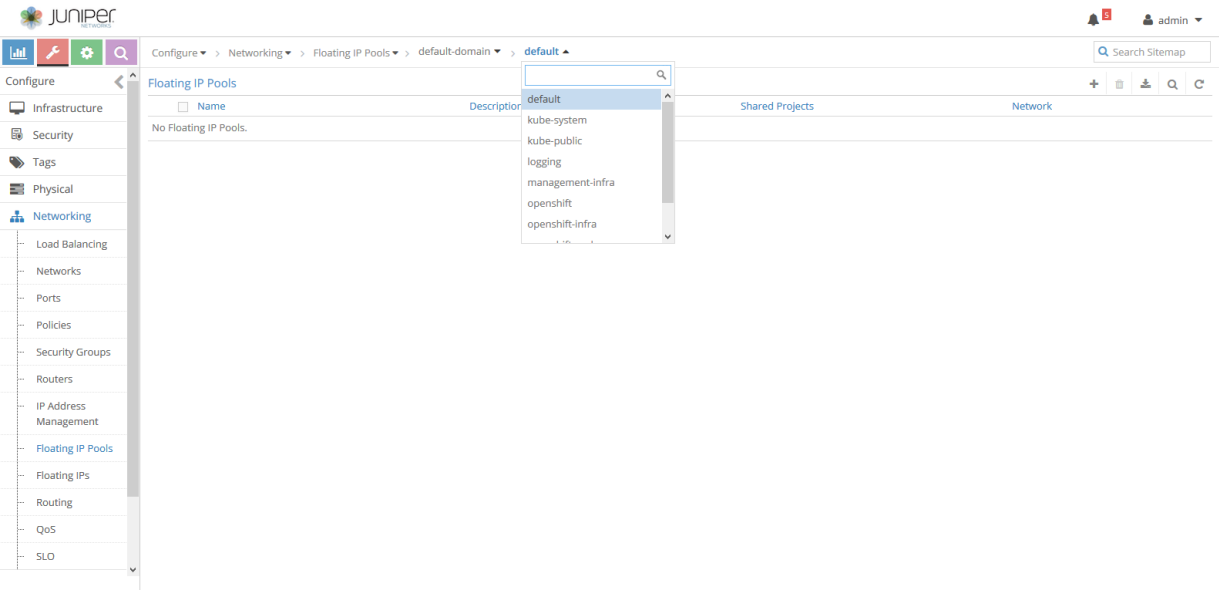
1. Select **Configure > Networking > Floating IP Pools**.

Figure 20: Floating IP Pools Selection



2. Select the network you want to associate with a floating IP pool.

Figure 21: Network Selection

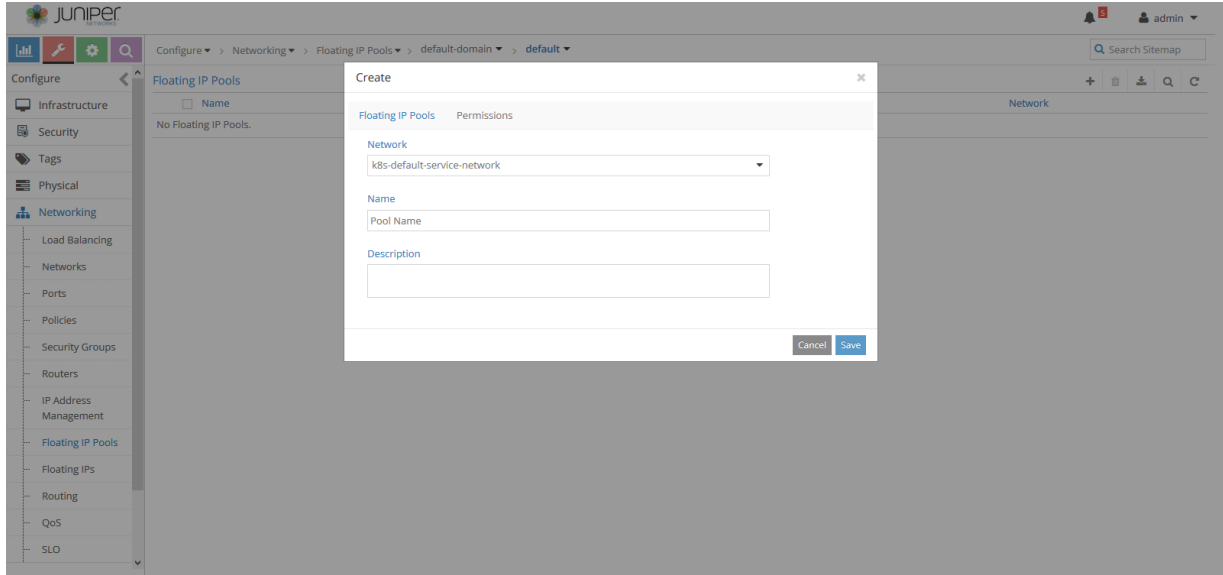


3. Click the add icon (+) to create a floating IP pool.

4. Add a **Name** and **Description** in the **Floating IP Pools** tab.

Click the **Permissions** tab to set **Owner Permissions** and **Global Share Permissions** for the floating IP pool. To associate the floating IP pool with multiple projects, click the add icon (+) in the **Share List**.

Figure 22: Create the Floating IP Pool



5. Click **Save** to create the floating IP address pool, or click **Cancel** to discard your changes and start over.

Support for IPv6 Networks in Contrail

IN THIS SECTION

- [Overview: IPv6 Networks in Contrail | 52](#)
- [Creating IPv6 Virtual Networks in Contrail | 53](#)
- [Adding IPv6 Peers | 54](#)

Starting with Contrail Release 2.0, support for IPv6 overlay networks is provided.

Overview: IPv6 Networks in Contrail

The following features are supported for IPv6 networks and overlay. The underlay network must be IPv4.

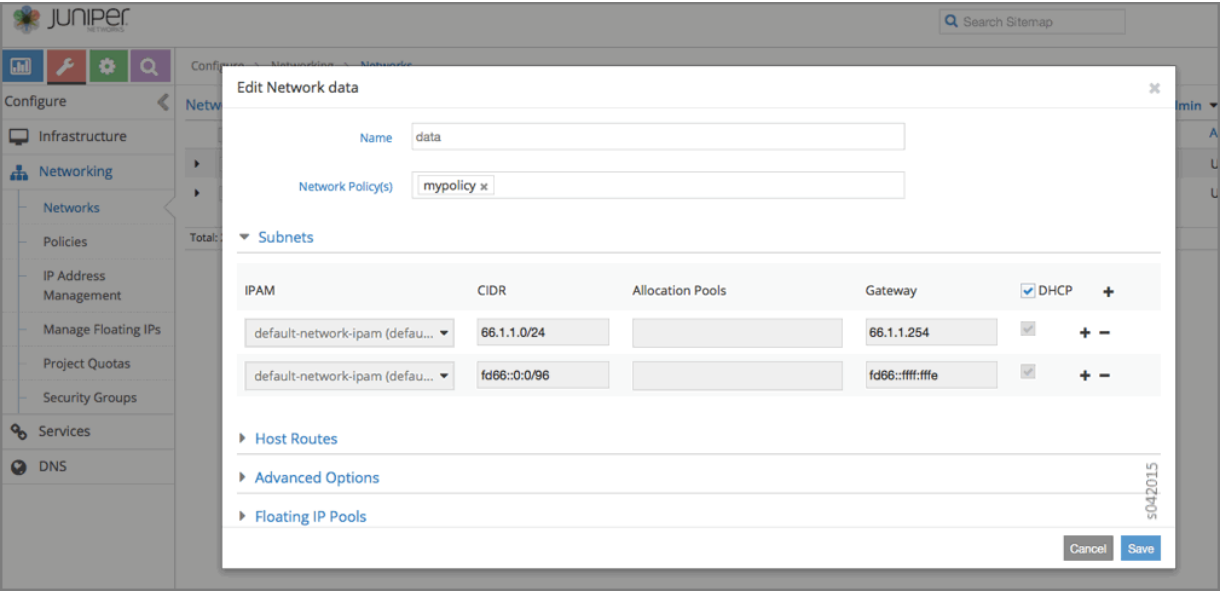
- Virtual machines with IPv6 and IPv4 interfaces
- Virtual machines with IPv6-only interfaces
- DHCPv6 and neighbor discovery
- Policy and Security groups
- IPv6 flow set up, tear down, and aging
- Flow set up and tear down based on TCP state machine
- Protocol-based flow aging
- Fat flow
- Allowed address pair configuration with IPv6 addresses
- IPv6 service chaining
- Equal Cost Multi-Path (ECMP)
- Connectivity with gateway (MX Series device)
- Virtual Domain Name Services (vDNS), name-to-IPv6 address resolution
- User-Visible Entities (UVEs)

NOT present is support for the following:

- Source Network Address Translation (SNAT)
- Load Balancing as a Service (LBaaS)
- IPv6 fragmentation
- Floating IP
- Link-local and metadata services
- Diagnostics for IPv6
- Contrail Device Manager
- Virtual customer premises equipment (vCPE)

Creating IPv6 Virtual Networks in Contrail

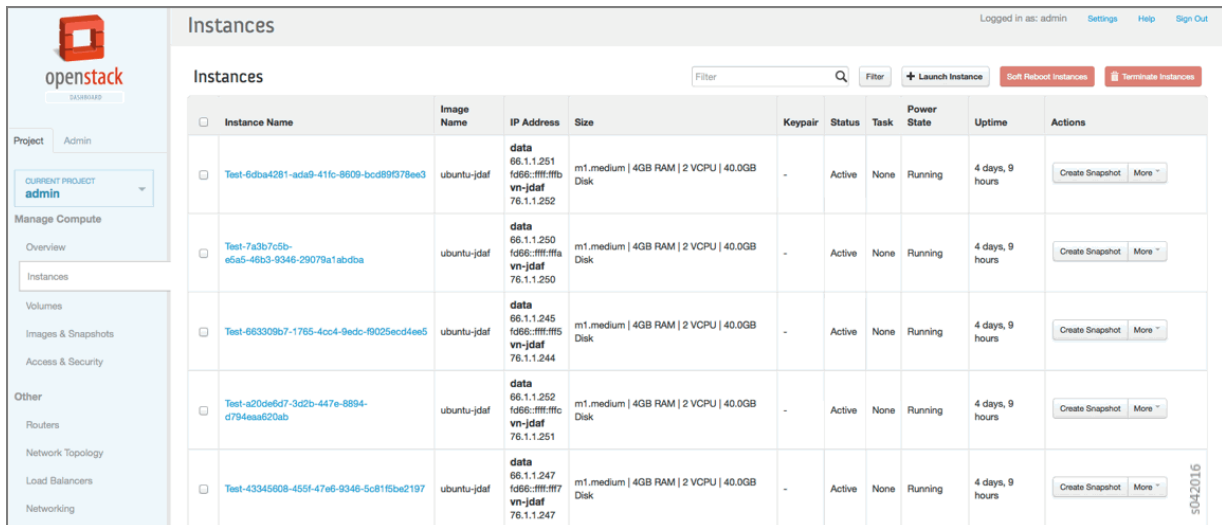
You can create an IPv6 virtual network from the Contrail user interface in the same way you create an IPv4 virtual network. When you create a new virtual network by selecting **Configure > Networking > Networks**, the Edit fields accept IPv6 addresses, as shown in the following image.



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 interface. On the left is a sidebar with the OpenStack logo and navigation links. The main area is titled 'Instances' and contains a table of five instances. Each instance is in a 'Running' state. The table columns are: Instance Name, Image Name, IP Address, Size, Keypair, Status, Task, Power State, Uptime, and Actions. The IP Address column shows IPv6 addresses for each instance.

Instance Name	Image Name	IP Address	Size	Keypair	Status	Task	Power State	Uptime	Actions
Test-6dbae4281-acda9-41fc-8609-bcd89f378ee3	ubuntu-jdof	66.1.1.251 fd96:ffff:fff:vn-jdaf 76.1.1.252	m1.medium 4GB RAM 2 VCPU 40.0GB Disk	-	Active	None	Running	4 days, 9 hours	Create Snapshot More ~
Test-7a3b7c5b-e6a5-48b3-9346-29079a1abda	ubuntu-jdof	66.1.1.250 fd96:ffff:fff:vn-jdaf 76.1.1.250	m1.medium 4GB RAM 2 VCPU 40.0GB Disk	-	Active	None	Running	4 days, 9 hours	Create Snapshot More ~
Test-663309b7-1765-4cc4-9edc-f9025ecd4ee5	ubuntu-jdof	66.1.1.245 fd96:ffff:fff:vn-jdaf 76.1.1.244	m1.medium 4GB RAM 2 VCPU 40.0GB Disk	-	Active	None	Running	4 days, 9 hours	Create Snapshot More ~
Test-a20de6d7-3d2b-447e-8894-d794eaa620ab	ubuntu-jdof	66.1.1.252 fd96:ffff:fff:vn-jdaf 76.1.1.251	m1.medium 4GB RAM 2 VCPU 40.0GB Disk	-	Active	None	Running	4 days, 9 hours	Create Snapshot More ~
Test-43345608-455f-47e6-9346-5c81f5be2197	ubuntu-jdof	66.1.1.247 fd96:ffff:fff:vn-jdaf 76.1.1.247	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:

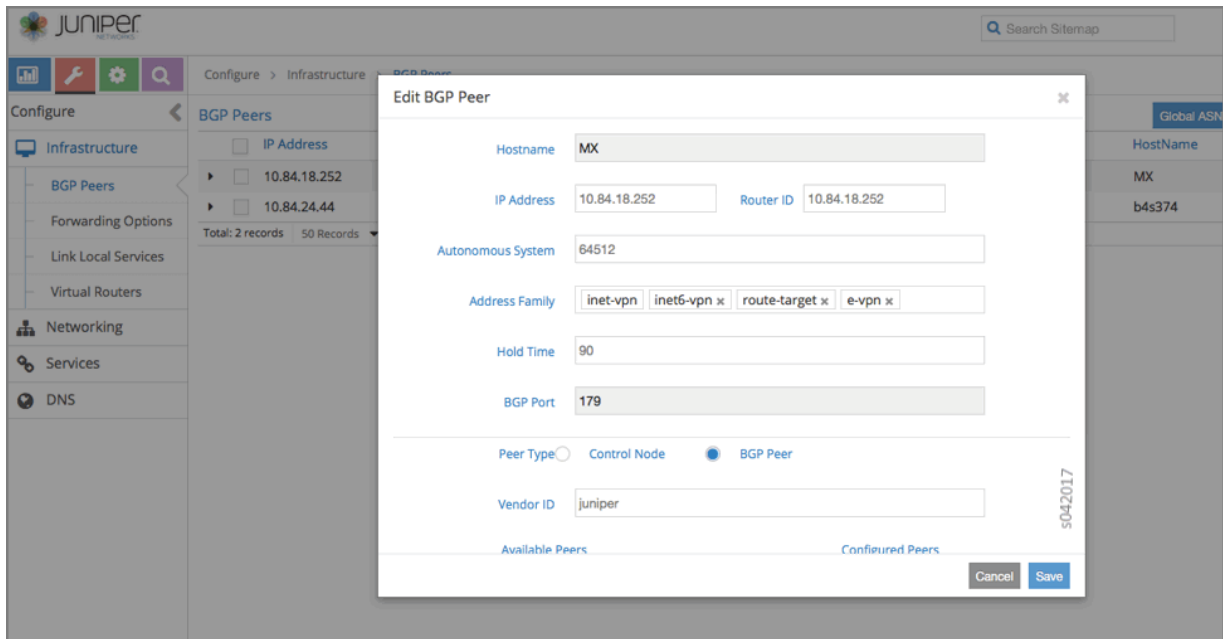
```
iface eth0 inet6 dhcp
```

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

IN THIS SECTION

- [Configuring the VXLAN Identifier Mode | 57](#)
- [Configuring Forwarding | 60](#)
- [Configuring the VXLAN Identifier | 62](#)
- [Configuring Encapsulation Methods | 63](#)

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

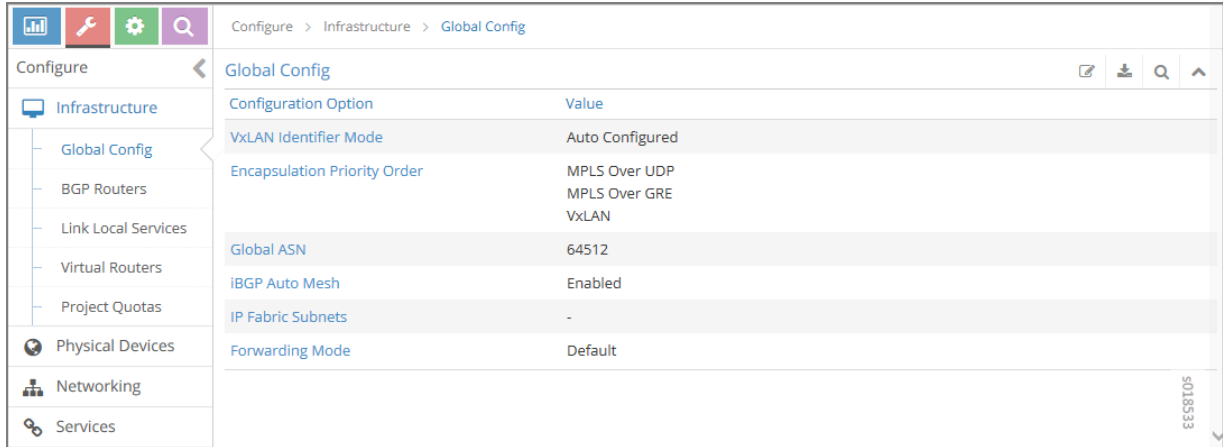
Configuring the VXLAN Identifier Mode

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 23: Global Config Window for VXLAN ID

Configuration Option	Value
VxLAN Identifier Mode	Auto Configured
Encapsulation Priority Order	MPLS Over UDP MPLS Over GRE VxLAN
Global ASN	64512
IBGP Auto Mesh	Enabled
IP Fabric Subnets	-
Forwarding Mode	Default

2. Click the edit icon .

The Edit Global Config window is displayed as shown in [Figure 24 on page 59](#).

Figure 24: 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/utls/encap.py` file as follows:

```
python encap.py <add | update | delete> <username> <password> <tenant_name> <config_node_ip>
```

Configuring Forwarding

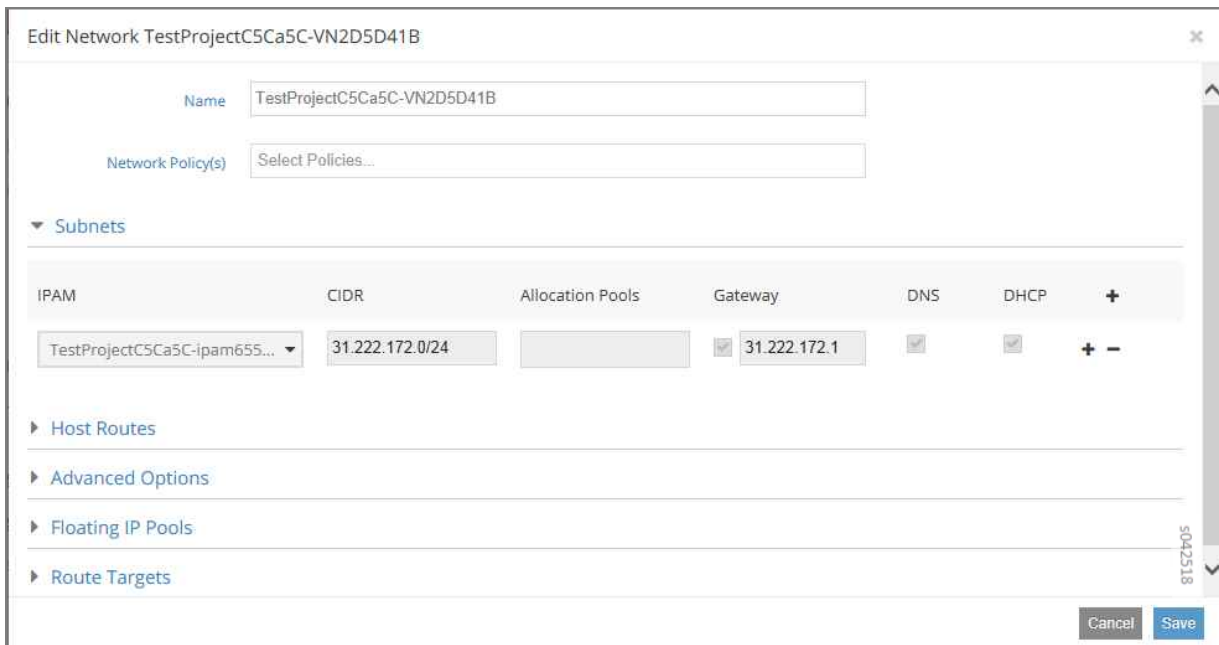
In Contrail, the default forwarding mode is enabled for fallback bridging (IP FIB and MAC FIB). The mode can be changed, either through the Contrail Web UI or by using python provisioning commands.

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 25 on page 60](#).

Figure 25: Edit Network Window



Edit Network TestProjectC5Ca5C-VN2D5D41B

Name: TestProjectC5Ca5C-VN2D5D41B

Network Policy(s): Select Policies...

▼ Subnets

IPAM	CIDR	Allocation Pools	Gateway	DNS	DHCP	+
TestProjectC5Ca5C-ipam655...	31.222.172.0/24		<input checked="" type="checkbox"/> 31.222.172.1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	+ -

▶ Host Routes

▶ Advanced Options

▶ Floating IP Pools

▶ Route Targets

Cancel Save

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 < I2_I3| I2 >
```

Options:


I2_I3 = Enable IP FIB and MAC FIB (fallback bridging)

I2 = Enable MAC FIB only (Layer 2 only)

Configuring the VXLAN Identifier

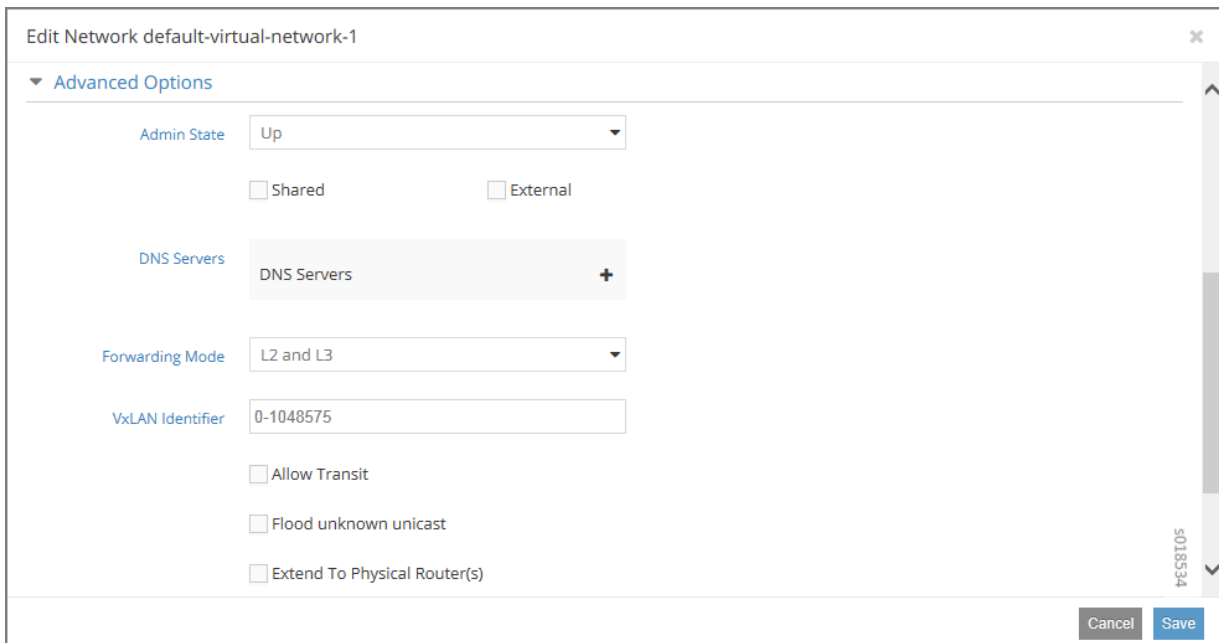
The VXLAN identifier can be set only if the VXLAN network identifier mode has been set to User Configured. You can then set the VXLAN ID by either using the Contrail Web UI or by using Python commands.

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 26 on page 62](#).

Figure 26: 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 27 on page 64](#).

Figure 27: Edit Global Config Window for Encapsulation Priority Order

The screenshot shows the 'Edit Global Config' window. It has a title bar with a close button. The main content is divided into two sections: 'Forwarding Options' and 'BGP Options'. Under 'Forwarding Options', there is a 'Forwarding Mode' dropdown set to 'Default'. Below that is the 'VxLAN Identifier Mode' with two radio buttons: 'Auto Configured' (unselected) and 'User Configured' (selected). Under 'User Configured', there is an 'Encapsulation Priority Order' section with a '+' button to its right. This section contains three dropdown menus: 'MPLS Over UDP', 'MPLS Over GRE', and 'VxLAN'. Each dropdown has a '+' and '-' button to its right. Under 'BGP Options', there is a 'Global ASN' field with the value '64512'. At the bottom right, there are 'Cancel' and 'Save' buttons. A vertical scrollbar is on the right side of the window.

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','M PLSoGRE'])
```

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.

Support for EVPN Route Type 5

Contrail Release 5.0.1 and later supports EVPN Route Type 5 messages as defined in the IETF specification *IP Prefix Advertisement in EVPN*. EVPN Route Type 5 is an extension of EVPN Route Type 2, which carries MAC addresses along with their associated IP addresses. EVPN Route Type 5 facilitates in inter-subnet routing.

Type 5 network layer reachability information (NLRI) contains information in the following format:

```
+-----+
|      RD      (8 octets)      |
+-----+
| Ethernet Segment Identifier (10 octets) |
+-----+
| Ethernet Tag ID (4 octets)      |
+-----+
| IP Prefix Length (1 octet)      |
+-----+
| IP Prefix (4 or 16 octets)      |
+-----+
| GW IP Address (4 or 16 octets)  |
+-----+
| MPLS Label (3 octets)          |
+-----+
```

When Type-5 EVPN prefix is received from a BGP peer, it is first installed into `bgp.evpn.0` like all other routes. From here, based on matching route targets, the route gets replicated into all `*.evpn.0` tables as applicable. From there, the routes are advertised over Extensible messaging and presence protocol (XMPP) to all interested agents.

NOTE: In Release 5.0.1, policy based route-leaking among different L3VRFs is not supported. Hence, service chaining for Type 5 L3VRFs is also not supported.

Support for EVPN Type 6 Selective Multicast Ethernet Tag Route


Contrail Release 5.1 supports EVPN Type 6 selective multicast Ethernet tag (SMET) route to selectively send or receive traffic based on the presence or absence of active receivers on a compute node. The EVPN Type-6 SMET route helps build and use multicast trees selectively on a per <*, G> basis.

Currently, all broadcast, unknown unicast, multicast (BUM) traffic is carried over the inclusive multicast ethernet tag (IMET) routes. This results in flooding of all compute nodes irrespective of whether an active receiver is present or not on each of those compute-nodes.

Configuring EVPN Type-6 SMET Routes

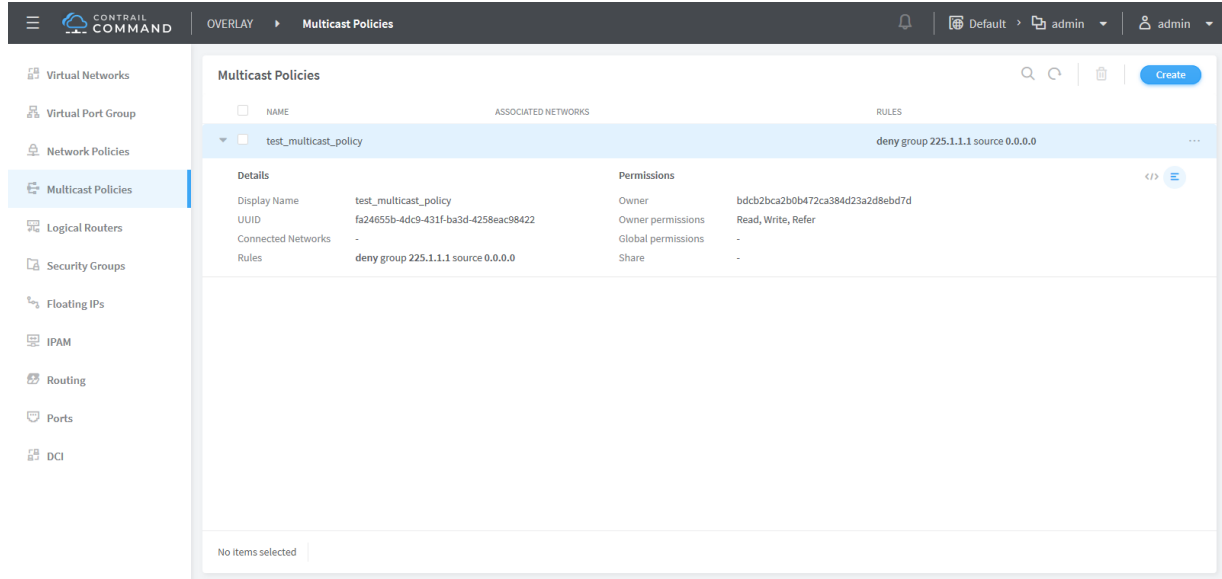
EVPN Type-6 SMET routes capability attaches a specific BGP community attribute **Ethernet Multicast flags Community (MF)** to the IMET routes. This community is advertised by default in Contrail release 5.1 and later. You must enable IGMP on the network as shown in [Figure 28 on page 67](#) as well as on the QFX device to which the multicast source is connected. You can configure IGMP at the global system configuration-level, at virtual network-level, or at VML-level. Configuring **ERB-UCAST-Gateway** role enables IGMP snooping on the QFX device.

Figure 28: Configure IGMP



You can allow or deny multicast traffic by attaching a policy at the virtual network-level as shown in [Figure 29 on page 68](#).

Figure 29: Define Multicast Policy




In Contrail Release 5.1, the receivers are always inside the contrail cluster and sender is always outside the cluster. This feature is supported only with `<*, G> /igmpv2`. The SMET feature is supported only on QFX10000 and QFX5110 devices running Junos OS Release 18.4R1 and later.

Support for L3VPN Inter AS Option C

Contrail Release 5.1 and later supports L3VPN inter AS Option C, which is used to interconnect multi-AS backbones as described in RFC 4364. Inter-AS option C uses BGP as the label distribution protocol. Labeled IPv4 routes are redistributed by eBGP between the neighboring autonomous systems. A new address family inet-labeled is added to maintain labeled unicast routes. This table is used for resolving L3VPN routes which are encapsulated using MPLS. [Figure 30 on page 69](#) shows the connectivity and the roles of different components in this architecture.

Figure 30: L3VPN Inter AS Option C Architecture



The controller maintains an eBGP session with the SDN-GW router and an iBGP session with the ASBR router. The controller exchanges labeled routes with the vRouters over XMPP. The vRouter uses MPLSoUDP or MPLSoGRE to reach the ASBR and encapsulates two labels within it - the inner VPN label and outer BGP-LU label. For the opposite direction, the vRouter advertises a labeled unicast route for its vhost address with a label 3 (implicit null), so the traffic from the ASBR delivered through the fabric to the vRouter over a UDP/GRE tunnel contains only the VPN label.

Configuring Inter AS Option C

This section describes how to configure L3VPN Inter AS Option C from Contrail Command UI.

1. Navigate to Infrastructure > Cluster > Advanced.
2. Click **BGP Routers** tab and click the **Edit** icon against the node that you want to configure.

The Edit BGP Router page is displayed.


3. Configure eBGP session for the SDN gateway. Make sure you include **inet-labeled** to the **Address Family** field.

4. Configure the iBGP session with the ASBR router. Address family must be set to **inet-labeled**.

5. Navigate to **Infrastructure > Cluster > Advanced**.
6. Click the **Virtual Routers** tab and on the Forwarding Options pane, click Edit to modify the encapsulation settings.

Contrail supports three encapsulation types. They are VXLAN, MPLS Over UDP, and MPLS Over GRE.

7. Set the Encapsulation priority order.



8. Click **Save**.

Changes made to Forwarding Options settings are saved.

3

CHAPTER

Deploying a Multi-Tier Web Application Using Contrail

Example: Deploying a Multi-Tier Web Application | **73**

Sample Network Configuration for Devices for Simple Tiered Web Application | **80**

Example: Deploying a Multi-Tier Web Application

IN THIS SECTION


- [Multi-Tier Web Application Overview | 73](#)
- [Example: Setting Up Virtual Networks for a Simple Tiered Web Application | 74](#)
- [Verifying the Multi-Tier Web Application | 76](#)
- [Sample Addressing Scheme for Simple Tiered Web Application | 77](#)
- [Sample Physical Topology for Simple Tiered Web Application | 78](#)
- [Sample Physical Topology Addressing | 79](#)

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 31 on page 74](#).

Figure 31: 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*.

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 a Virtual Network with OpenStack Contrail* or [“Creating a Virtual Network with Juniper Networks Contrail” on page 46.](#)

- 3. Create a floating IP pool named **public_pool** for the **public** network within the **demo** project; see [Figure 32 on page 75.](#)

Figure 32: Create Floating IP Pool

The screenshot shows a web-based configuration interface titled "Edit Network public". It contains several sections: "Network Name" with a text field containing "public"; "Network Policy(s)" with a button "Select Policies..."; "Address Management" with a dropdown menu showing "default-network..." and a text field "xxx.xxx.xxx.xxx/xx" with expand/collapse buttons; a table with headers "IPAM" and "IP Block" containing "default-network-ipam" and "10.84.41.0/24"; and "Floating IP Pools" with a text field "public_pool" and a dropdown menu showing "demo" with a close button. A dropdown menu is open below "demo", showing "admin" as a selection. At the bottom right are "Cancel" and "Save" buttons. A vertical label "s041841" is on the right edge.

IPAM	IP Block
default-network-ipam	10.84.41.0/24

Pool Name
admin

- 4. Allocate the floating IP pool **public_pool** to the **demo** project; see [Figure 33 on page 76.](#)

Figure 33: Allocate Floating IP

5. Verify that the floating IP pool has been allocated; see **Configure > Networking > Allocate Floating IPs**.

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.

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.

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 6 on page 77](#) as a guide for addressing devices in the simple tiered web example.

Table 6: 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


Table 6: Sample Addressing Scheme for Example (continued)

System Name	Address Allocation
public network (floating address)	10.84.41.0/24

Sample Physical Topology for Simple Tiered Web Application

Figure 34 on page 78 provides a guideline diagram for the physical topology for the simple tiered web application example.

Figure 34: Sample Physical Topology for Simple Tiered Web Application



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 73](#). 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 "xxxxxxxx"; ## 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 "xxxxxxxxxxxxx"; ## SECRET-DATA
}
login {
    class read {
        permissions [ clear interface view view-configuration ];
    }
    user admin {
        uid 2000;
        class super-user;
        authentication {
            encrypted-password "xxxxxxxxxxxxx"; ## SECRET-DATA
        }
    }
    user user1 {
        uid 2002;
        class read;
        authentication {
            encrypted-password "xxxxxxxxxxxxx"; ## 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;  
    }  
  }  
}
```

4

CHAPTER

Configuring Services

Configuring DNS Servers | **88**

Distributed Service Resource Allocation with Containerized Contrail | **98**

Support for Broadcast and Multicast | **107**

Configuring DNS Servers

IN THIS SECTION

- [DNS Overview | 88](#)
- [Defining Multiple Virtual Domain Name Servers in Contrail | 89](#)
- [IPAM and Virtual DNS | 89](#)
- [DNS Record Types | 90](#)
- [Configuring DNS on the User Interface | 91](#)
- [Configuring DNS Using Scripts | 97](#)

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 in Contrail

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 36 on page 89](#) shows examples of virtual DNS servers defined in **default-domain**, providing the name service for the DNS domains indicated.


Figure 36: 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 37 on page 89](#)

Figure 37: IPAM and Virtual DNS



DNS Record Types

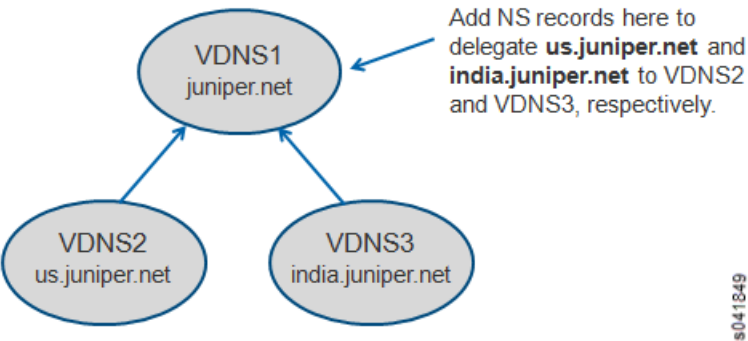
DNS records can be added statically. DNS record types **A**, **CNAME**, **PTR**, and **NS** are currently supported in the system. Each record includes the type, class (IN), name, data, and TTL values. See [Table 7 on page 90](#) for descriptions of the record types.

Table 7: DNS Record Types Supported

DNS Record Type	Description
A	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.
CNAME	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.
PTR	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.
NS	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 38 on page 90](#) shows an example usage for the DNS record type of **NS**.

Figure 38: Example Usage for NS Record Type



Configuring DNS on the User 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 39 on page 91](#).

Figure 39: Configure DNS Records

Configure > DNS > Servers

Search

Configure DNS Records

default-domain

admin

Configure Virtual DNS

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 40 on page 92](#)

Figure 40: Add DNS

Complete the fields for the new server; see [Table 8 on page 92](#).

Table 8: 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 None .
Load Balancing Order	Select the load-balancing order from a list— Random , Fixed , Round Robin . When a name has multiple records matching, the configured record order determines the order in which the records are sent in the response. Select Random to have the records sent in random order. Select Fixed to have records sent in the order of creation. Select Round Robin to have the record order cycled for each request to the record.
OK	Click OK to create the record.
Cancel	Click Cancel to clear the fields and start over.

- 3. 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 41 on page 93](#).

Figure 41: Add DNS Record

Add DNS Record

Type

A (IP Address Record)

Host Name

Host Name to be resolved

IP Address

Enter an IP Address

Class

IN (Internet)

Time To Live

TTL(86400 secs)

Cancel

Save

- 4. Complete the fields for the new record; see [Table 9 on page 93](#).

Table 9: Add DNS Record Fields

Field	Description
Record Name	Enter a name for this record.
Type	Select the record type from a list— A , CNAME , PTR , NS .
IP Address	Enter the IP address for the location for this record.
Class	Select the record class from a list— IN is the default.
Time To Live	Enter the TTL in seconds.
OK	Click OK to create the record.

Table 9: Add DNS Record Fields (continued)

Field	Description
Cancel	Click Cancel 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 42 on page 94](#).

Figure 42: Associate IPAMs to DNS

Edit DNS Server

Server Name

vdns1

Domain Name

juniper.net

DNS Forwarder

Enter Forwarder IP or Select a DNS Server

Record Resolution Order

Random

Time To Live

86400

Associate IPAMs

admin:ipam1

default-project:default-network-ipam

Cancel

Save

Complete the IPAM associations, using the field descriptions in [Table 10 on page 94](#).

Table 10: Associate IPAMs to DNS Fields

Field	Description
Associate to All IPAMs	Select this box to associate the selected DNS server to all available IPAMs.
Available IPAMs	This column displays the currently available IPAMs.
Associated IPAMs	This column displays the IPAMs currently associated with the selected DNS server.

Table 10: Associate IPAMs to DNS Fields (*continued*)

Field	Description
>>	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.
OK	Click OK to commit the changes indicated in the window.
Cancel	Click Cancel to clear all entries and start over.

6. Use the **IP Address Management** page (**Configure > Networking > IP Address Management**); see [Figure 43 on page 95](#)) 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 43: 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 44 on page 96](#).

Figure 44: DNS Server

Add IP Address Management

Name

DNS Method **Virtual DNS** ▼

Virtual DNS **Virtual DNS**

NTP Server IP **None**

Associate IP Blocks to Networks ☐

fip_vn IP Block Gateway + -

Network	IP Block	Gateway

Cancel Save

s041857

8. In the first field, select the **DNS Method** from a list (**None**, **Default DNS**, **Tenant DNS**, **Virtual DNS**; see [Table 11 on page 96](#)).

Table 11: DNS Modes

DNS Mode	Description
None	Select None when no DNS support is required for the VMs.
Default	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.
Tenant	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.

Table 11: DNS Modes (*continued*)

DNS Mode	Description
Virtual DNS	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

You can configure DNS by using scripts that are available in the `contrail-utils` RPM/DEB package in the `/opt/contrail/utils` directory. The scripts are copied to the `config_api_container` or `config node` when you install the `contrail-utils` RPM/DEB package. You can execute the scripts from either the `config_api` container or the `config node`. The scripts are described in [Table 12 on page 97](#).



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 12: DNS Scripts

Action	Script
Add a virtual DNS server	<p>Script: <code>add_virtual_dns.py</code></p> <p>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:vdns2</code></p>
Delete a virtual DNS server	<p>Script: <code>del_virtual_dns_record.py</code></p> <p>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></p>

Table 12: DNS Scripts (*continued*)

Action	Script
Add a DNS record	<p>Script: <code>add_virtual_dns_record.py</code></p> <p>Sample usage: <code>python add_virtual_dns_record.py --api_server_ip 10.204.216.21 --api_server_port 8082 --name rec1 --v dns_fqname default-domain:vdns1 --rec_name one --rec_type A --rec_class IN --rec_data 1.2.3.4 --rec_ttl 2400</code></p>
Delete a DNS record	<p>Script: <code>del_virtual_dns_record.py</code></p> <p>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></p>
Associate a virtual DNS server with an IPAM	<p>Script: <code>associate_virtual_dns.py</code></p> <p>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 --v dns_fqname default-domain:vdns1</code></p>
Disassociate a virtual DNS server with an IPAM	<p>Script: <code>disassociate_virtual_dns.py</code></p> <p>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 --v dns_fqname default-domain:vdns1</code></p>

Distributed Service Resource Allocation with Containerized Contrail

IN THIS SECTION

- [Replacement of Centralized Discovery Service | 99](#)
- [New Distributed Resource Allocation Manager | 99](#)
- [Changes in Configuration Files | 100](#)

Starting with Contrail Release 4.0, the existing centralized Contrail discovery service is replaced with a distributed method of allocating service resources.

Replacement of Centralized Discovery Service

In Contrail releases prior to Release 4.0, the Contrail discovery service is a centralized service resource allocation module with high availability, used primarily to automatically load-balance service resources in the system.

In the previous centralized discovery method, new service resources are registered (published) directly to the Contrail discovery module and allocated to the requester (subscriber) of the service resource, without disrupting the running state of the subscribers.

The centralized discovery method requires using a database to:

- synchronize across Contrail discovery nodes.
- maintain the list of publishers, subscribers, and the health of published services across reloads.
- provide a centralized view of the service allocation and health of the services.

This centralized discovery method resulted in unnecessary system churn when services were falsely marked as down, due to periodic health updates of services made to the database nodes, resulting in reallocation of healthy services.

Starting with Contrail 4.0, the Contrail discovery services centralized resource allocation manager has been removed. Its replacement is a distributed resource allocation list of service nodes, maintained in each module of the system.

New Distributed Resource Allocation Manager

Starting with Contrail Release 4.0, service resources are managed with a distributed allocation manager, with the following features:

- Each system module is provisioned with a list of service nodes (publishers).
- Each system module randomizes the list of service nodes and uses the resources. The randomized list is expected to be fairly load-balanced.
- When currently-used services are down, the system module detects the down immediately and reacts with no downtime by selecting another service from the list. This is distinctly different from the previous model, in which the module would need to contact the discovery service to check for available services, resulting in a finite time loss for allocation, distribution, and application of a new set of services.
- When service nodes are added or deleted, the system administrator updates the configuration file of all daemons using the service type of the service node added or deleted, sending a SIGHUP to the respective daemons.
- Each daemon randomizes the service list independently and reallocates the resources.

Deprecation of IF-MAP

In Contrail 4.0, the Interface for Metadata Access Points (IF-MAP) methodology has been deprecated. Contrail 4.0 uses CONFIGDB sections in configuration files instead of IF-MAP sections.

Changes in Configuration Files

Table 13 on page 100 lists configuration files in the Contrail system that have changes to enable the distributed service resource allocation system, starting with Contrail 4.0. In general, the changes include removing (deprecating) discovery server sections and subsections, and adding parameters needed to identify service resources in all modules.

Each daemon randomizes the published service list and uses the resources. Additionally, each daemon provides a SIGHUP handler to manage the addition or deletion of publishers.

Table 13: Contrail 4.0 Changes in Configuration Files

Configuration File	Configuration Parameter	Changes
<code>contrail-vrouter-agent.conf</code>	<code>[DISCOVERY]</code>	Section deprecated
	<code>[CONTROL-NODE].servers</code>	Provisioned list of control-node [role=control] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:5269 10.1.1.12:5269
	<code>[DNS].servers</code>	Provisioned list of DNS [role=control] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:53 10.1.1.2:5
	<code>[DEFAULT].collectors</code>	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
<code>contrail-control.conf</code>	<code>[DISCOVERY]</code>	Section deprecated

Table 13: Contrail 4.0 Changes in Configuration Files *(continued)*

Configuration File	Configuration Parameter	Changes
	[DEFAULT].collectors	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
	[CONFIGDB].rabbitmq_server_list	Provisioned list of config-node [role=cfgm] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:5672 10.1.1.2:5672
	[CONFIGDB].rabbitmq_user	guest (default string)
	[CONFIGDB].rabbitmq_password	guest (default string)
	[CONFIGDB].config_db_server_list	Provisioned list of Config DB [role=database] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:9042 10.1.1.2:9042 NOTE: Docker uses 9041 as port
	[CONFIGDB].certs_store	Deprecated
	[CONFIGDB].password	Deprecated
	[CONFIGDB].server_url	Deprecated
	[CONFIGDB].user	Deprecated
	[CONFIGDB].stale_entries_cleanup_timeout	Deprecated
	[CONFIGDB].end_of_rib_timeout	Deprecated
contrail-dns.conf		
	[DISCOVERY]	Deprecated

Table 13: Contrail 4.0 Changes in Configuration Files *(continued)*

Configuration File	Configuration Parameter	Changes
	[DEFAULT].collectors	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
	[CONFIGDB].rabbitmq_server_list	Provisioned list of config-node [role=cfgm] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:5672 10.1.1.2:5672
	[CONFIGDB].rabbitmq_user	guest (default string)
	[CONFIGDB].rabbitmq_password	guest (default string)
	[CONFIGDB].config_db_server_list	Provisioned list of Config DB [role=database] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:9042 10.1.1.2:9042 NOTE: Dockers use 9041 as port
	[CONFIGDB].certs_store	Deprecated
	[CONFIGDB].password	Deprecated
	[CONFIGDB].server_url	Deprecated
	[CONFIGDB].user	Deprecated
	[CONFIGDB].stale_entries_cleanup_timeout	Deprecated
	[CONFIGDB].end_of_rib_timeout	Deprecated
contrail-collector.conf	[DISCOVERY]	Deprecated
	[API_SERVER].api_server_list	Provisioned list of api-servers [role=config] in the format: ip-address:port Example: 10.1.1.1:8082 10.1.1.2:8082

Table 13: Contrail 4.0 Changes in Configuration Files *(continued)*

Configuration File	Configuration Parameter	Changes
contrail-alarm-gen.conf	[DISCOVERY]	Deprecated
	[DEFAULTS].collectors	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
	[API_SERVER].api_server_list	Provisioned list of api-servers [role=config] in the format: ip-address:port Example: 10.1.1.1:8082 10.1.1.2:8082
	[REDIS].redis_uve_list	Provisioned list of redis instances [role=collector] Example: 192.168.0.29:6379 192.168.0.30:6379
contrail-analytics-api.conf	[DISCOVERY]	Section deprecated
	[DEFAULTS].collectors	Provisioned list of collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
	[REDIS].redis_uve_list	Provisioned list of redis instances [role=collector] Example: 192.168.0.29:6379 192.168.0.30:6379
contrail-api.conf	[DISCOVERY]	Section deprecated
	[DEFAULTS].collectors	Provisioned list of collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086

Table 13: Contrail 4.0 Changes in Configuration Files *(continued)*

Configuration File	Configuration Parameter	Changes
contrail-schema.conf	[DISCOVERY]	Section deprecated
	[DEFAULTS].collectors	Provisioned list of Collector [role=collector] service providers in ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
contrail-svc-monitor.conf	[DISCOVERY]	Section deprecated
	[DEFAULTS].collectors	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
contrail-device-manager.conf	[DISCOVERY]	Section deprecated
	[DEFAULTS].collectors	Provisioned list of Collector [role=collector] service providers in ip-address:port ip-address2:port format Example: 10.1.1.1:8086 10.1.1.2:8086
contrail-analytics-nodemgr.conf	[DISCOVERY]	Section deprecated
	[COLLECTOR].server_list	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
contrail-config-nodemgr.conf	[DISCOVERY]	Section deprecated
	[COLLECTOR].server_list	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
contrail-control-nodemgr.conf	[DISCOVERY]	Section deprecated

Table 13: Contrail 4.0 Changes in Configuration Files *(continued)*

Configuration File	Configuration Parameter	Changes
	[COLLECTOR].server_list	Provisioned list of Collector [role=collector] service providers in ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
contrail-database-nodemgr.conf	[DISCOVERY]	Section deprecated
	[COLLECTOR].server_list	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
contrail-vrouter-nodemgr.conf	[DISCOVERY]	Section deprecated
	[COLLECTOR].server_list	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
contrail-query-engine.conf	[DISCOVERY]	Section deprecated
	[COLLECTOR].server_list	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
contrail-snmp-collector.conf	[DISCOVERY]	Section deprecated
	[DEFAULTS].collectors	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086

Table 13: Contrail 4.0 Changes in Configuration Files *(continued)*

Configuration File	Configuration Parameter	Changes
	[API_SERVER].api_server_list	Provisioned list of api-servers [role=config] in the format: ip-address:port Example: 10.1.1.1:8082 10.1.1.2:8082
contrail-topology.conf	[DISCOVERY]	Section deprecated
	[DEFAULTS].collectors	Provisioned list of Collector [role=collector] service providers in the format: ip-address:port ip-address2:port Example: 10.1.1.1:8086 10.1.1.2:8086
	[API_SERVER].api_server_list	Provisioned list of api-servers [role=config] in ip-address:port Example: 10.1.1.1:8082 10.1.1.2:8082
Contrail Web UI		
config.global.js	config.discovery.server	Discovery subsection deprecated
	config.discovery.port	Discovery subsection deprecated
	config.cnfg.server_ip	Provisioned list of Config [role=cfgm] service providers as list of ip-address Example: ['10.1.1.1 10.1.1.2']
	config.cnfg.server_port	Server port as a string Example: '8082'
	config.analytics.server_ip	Provisioned list of Collector [role=collector] service providers as a list of ip-address Example: ['10.1.1.1 10.1.1.2']
	config.analytics.server_port	Server port as a string Example: '8081'

Table 13: Contrail 4.0 Changes in Configuration Files *(continued)*

Configuration File	Configuration Parameter	Changes
	<code>config.dns.server_ip</code>	Provisioned list of Controller [role=control] service providers as a list of ip-address Example: ['10.1.1.1 10.1.1.2']
	<code>config.dns.server_port</code>	Server port as a string Example: '8092'

Support for Broadcast and Multicast

IN THIS SECTION

- Subnet Broadcast | 108
- All-Broadcast/Limited-Broadcast and Link-Local Multicast | 108
- Host Broadcast | 109

This section describes how the Contrail Controller supports broadcast and multicast.

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.

5

CHAPTER

Configuring Service Chaining

Service Chaining | **111**

Service Chaining MX Series Configuration | **115**

ECMP Load Balancing in the Service Chain | **116**

Service Chain Version 2 with Port Tuple | **118**

Service Chain Route Reorigination | **122**

Example: Creating an In-Network Service Chain by Using Contrail Command | **144**

Example: Creating an In-Network-NAT Service Chain | **153**

Example: Creating a Transparent Service Chain by Using Contrail Command | **161**

Using Static Routes with Services | **171**

Configuring Metadata Service | **176**

Service Chaining

IN THIS SECTION

- [Service Chaining Basics | 111](#)
- [Service Chaining Configuration Elements | 112](#)


Contrail Controller supports chaining of various Layer 2 through Layer 7 services such as firewall, NAT, IDP, and so on.

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 45 on page 111](#) 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 45: 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 46 on page 112](#) 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 46: 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 Service Chain by Using Contrail Command | 144](#)

[Example: Creating an In-Network-NAT Service Chain | 153](#)

[Example: Creating a Transparent Service Chain by Using Contrail Command | 161](#)

[ECMP Load Balancing in the Service Chain | 116](#)

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 80.

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 47 on page 117](#), 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 47: Load Balancing a Service Chain



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 | 111](#)

Customized Hash Field Selection for ECMP Load Balancing

Service Chain Version 2 with Port Tuple

IN THIS SECTION

- [Overview of Port Tuple | 118](#)
- [Service Chain Version 2 Sample Workflow | 119](#)
- [Service Chain with Health Check | 121](#)

Starting with Contrail 3.0, the user can create a port-tuple object for binding service instances to ports.

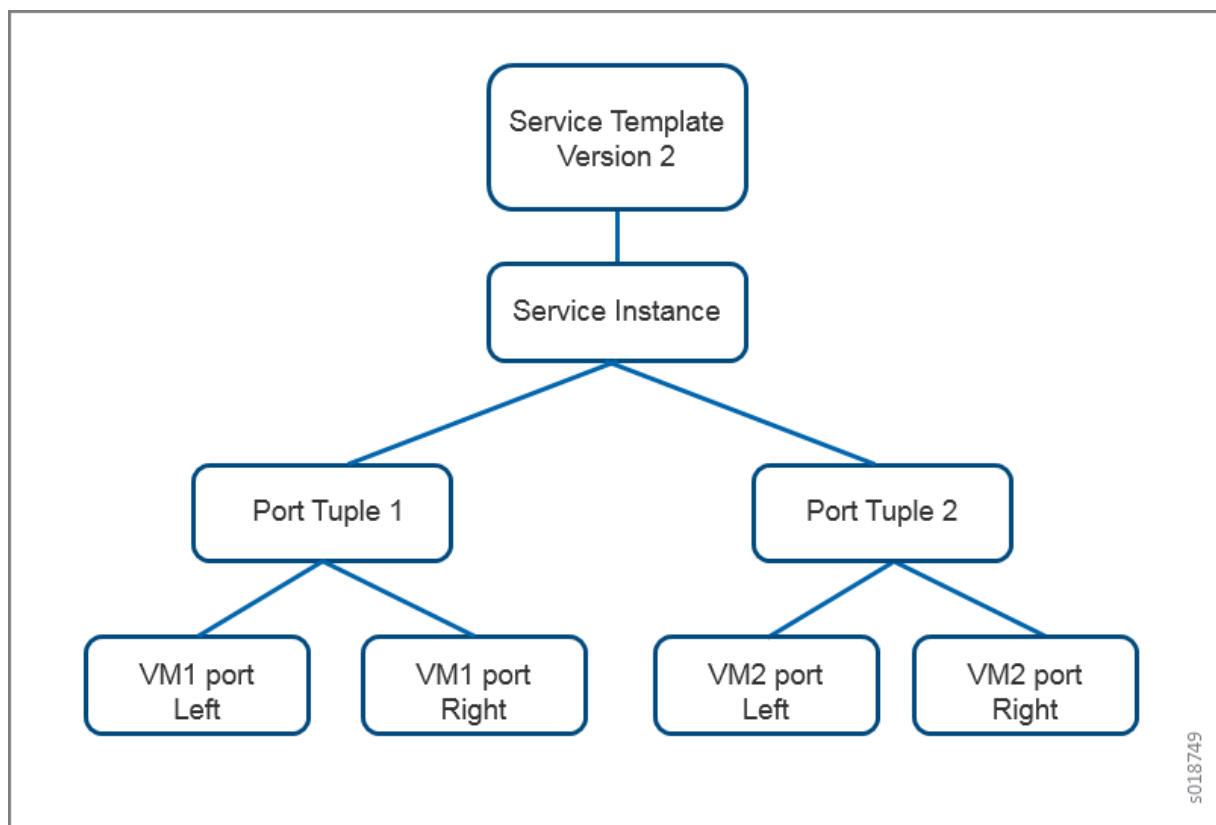
Overview of Port Tuple

In previous versions of Contrail, when a service instance is created for a virtual machine (VM)-based service, the service monitor creates one or more VM objects and creates a port for each VM object. Each VM object is a placeholder for binding a service instance to a port. The VM object also acts as a placeholder for the instance ID when using equal-cost multipath (ECMP).

Using the VM object as a placeholder doesn't add value beyond binding information between the service instance object and the port objects. By using a port-tuple object, the service instance can be linked directly to the port objects, eliminating the need to create a VM object. This simplifies the implementation of service instance objects, and also allows integration with Heat templates.

With a port-tuple object, the user can create ports and pass the port information when creating a service instance. The ports can be created as part of a VM launch from Nova or without using a VM launch. The ports are linked to a port-tuple object that is a child of a service instance. This functionality can also be leveraged in Heat stacks. See [Figure 48 on page 119](#).

Figure 48: Port Tuple Overview



Service Chain Version 2 Sample Workflow

With Contrail service templates Version 2, the user can create ports and bind them to a VM-based or container-based service instance, by means of a port-tuple object. All objects created with the Version 2 service template are visible to the Contrail Heat engine, and are 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 **left** and **right** are used for forwarding.
6. Link the port tuple to a service instance.
7. Launch the service instance. This creates the necessary objects in the Contrail database.

NOTE: Port-tuple is *not* supported on transparent service instance, whether active/active, active/standby, or scale-out.

Service Chain with Equal-Cost Multipath in Active-Active Mode

Equal-cost multipath (ECMP) can be used to distribute traffic across VMs. To support ECMP in the service chain, create multiple port tuples within the same service instance. The labels should be the same for the VM ports in each port tuple. For example, if port tuple 1 uses the labels **left** and **right**, then port tuple 2 in the same service instance should also use the labels **left** and **right** for its ports.

When there are multiple port tuples, the default mode of operation is **active-active**.

Service Chain Active-Standby Mode with Allowed Address Pair

To support **active-standby** mode, you must configure an allowed address pair on the interfaces. The **active-standby** is used as the high availability mode in the allowed address pair. The allowed address pair is configured as part of the service instance for a particular VM port label. For example, if the allowed address pair is configured in a service instance for the port with the label **left**, then all of the port-tuple VM ports with the label **left** will use the allowed address pair high availability mode.

Allowed Address Pair

An allowed address pair extension is an OpenStack feature supported by Contrail.

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

In Contrail, you can configure allowed address pairs in the service instance configuration, using **Configure > Services > Service Instances > Allowed Address Pair**, see [Figure 49 on page 121](#).

Figure 49: Edit Service Instance, Allowed Address Pair

Edit

Interface Type **Virtual Machine Interface**

left (1.1.1.4) - f88f8960-7897-4df5-b... ▼

right (2.2.2.4) - 1fab8367-a4ed-43cb-... ▼

▶ Service Health Check

▶ Routing Policy

▶ Route Aggregate

▼ Allowed Address Pair

Interface Type	IP	MAC	+
left ▼	1.1.1.254	00:00:5e:00:01:03	+ -
right ▼	2.2.2.254	00:00:5e:00:01:04	+ -

Cancel Save

s019908

For more information about OpenStack allowed address pairs, see

https://specs.openstack.org/openstack/neutron-specs/specs/api/allowed_address_pairs.html.

Service Chain with Static Route Table

The service chain Version 2 also supports static route tables. A static route table is configured similar to how the allowed address pair is configured, except with using the label **right**. The route table will be attached to the correct VM ports of the port tuples, based on the configuration of the port with the label **right**.

Service Chain with Health Check

Service chain Version 2 also allows service instance health check configuration on a per interface label. This is used to monitor the health of the service.

For more information about the service instance health check, see *Health Check Object*.

RELATED DOCUMENTATION

| *Health Check Object*

Service Chain Route Reorigination

IN THIS SECTION

- [Overview: Service Chaining in Contrail | 122](#)
- [Route Aggregation | 124](#)
- [Routing Policy | 132](#)
- [Control for Route Reorigination | 142](#)

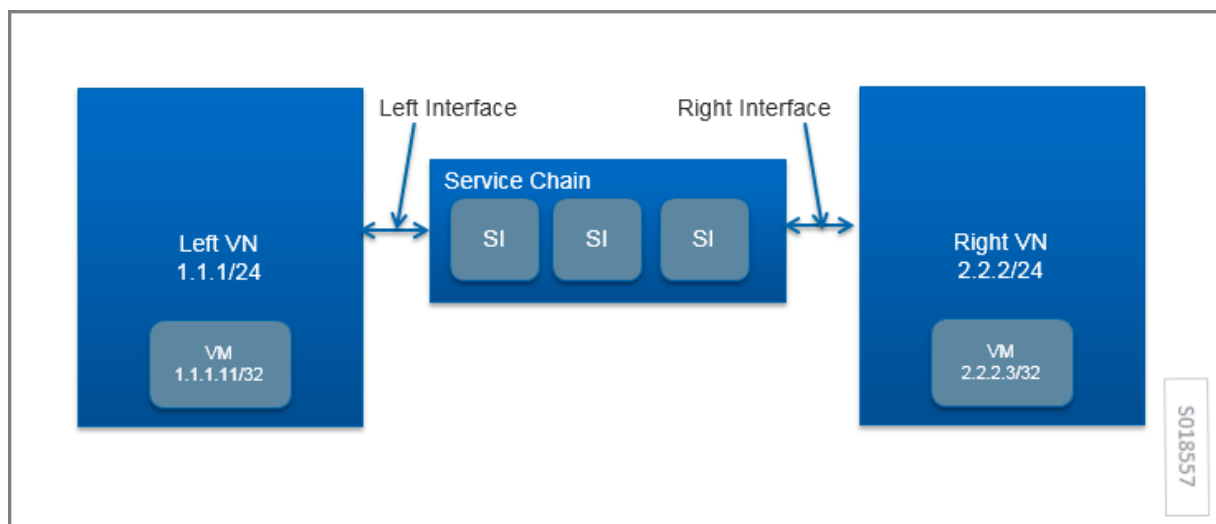
Overview: Service Chaining in Contrail

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 50 on page 123](#), 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 50: 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.

Starting in Contrail Release 5.0, when one or more than one service instance in a service chain fails, reorigination of routes on both sides of the service chain is stopped and routes automatically converge to

a backup service chain that is part of another Contrail cluster. For more information, see *Service Instance Health Checks*.

Route Aggregation

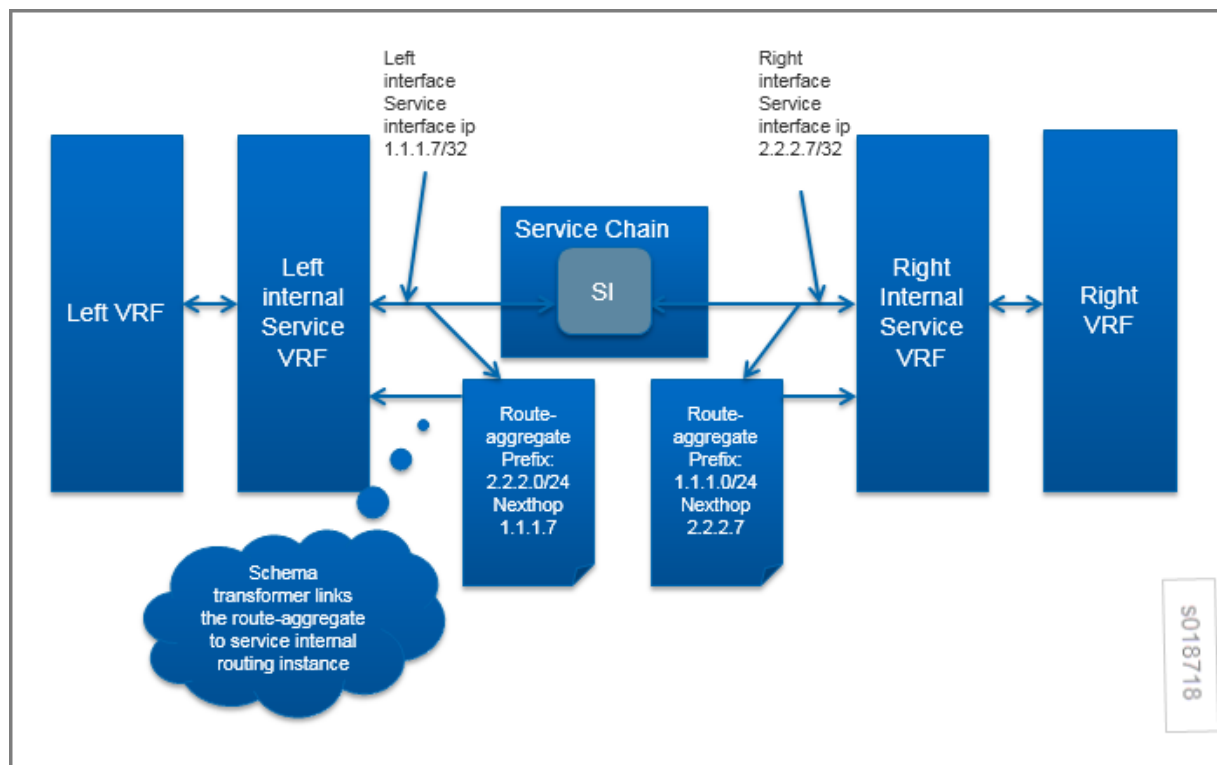
The route aggregation configuration object contains a list of prefixes to aggregate. The next-hop field in the route aggregate object contains the address of the route whose next hop is stitched as a next hop of the aggregate route.

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 51 on page 124](#).

Figure 51: 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

IN THIS SECTION

- [Route Aggregate Object | 125](#)
- [Service Instance Link to Route Aggregate Object | 126](#)
- [Routing Instance Link to Route Aggregate Object | 126](#)

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

IN THIS SECTION

- [Configure Route Aggregate Object | 127](#)
- [Configuring a Service Instance | 128](#)
- [Create a Virtual Network and Network Policy | 129](#)
- [Validate the Route Aggregate Object in the API Server | 129](#)
- [Validate the Route Aggregate Object in the Control Node | 130](#)

Configure Route Aggregate Object

You can use the Contrail UI, **Configure > Networking > Routing > Create > Route Aggregate** screen to name the route aggregate object and identify the routes to aggregate. See [Figure 52 on page 127](#).

Figure 52: Create Route Aggregate

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 53 on page 128](#).

Figure 53: Create Service Instance

Create Service Instance [X]

si-aggregate st-with-aggregate - [transparent (left, right)...

▼ **Interface Details**

Interface Type	Virtual Network
left	Auto Configured
right	Auto Configured

▼ **Advanced Options**

► Routing Policy

▼ **Route Aggregate**

Interface Type	Route Aggregate
right	left-to-right x

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 checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Service Instance
si-aggregate

+ Add Rule

Cancel Save

Attach the network policy to create the service chain between the left and right VNs. See the following example.

Edit Network

Name
left

Network Policy(s)
default-domain:admin:service-chain-policy

Subnets

IPAM	CIDR	Allocation Pools	Gateway	DNS	DHCP
default-network-ip...	1.1.1.0/24	start-end	1.1.1.1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Host Route(s)

Route Prefix	Next Hop
--------------	----------

Advanced Options

Cancel Save

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_next_hop: "1.1.1.3",
    - routing_instance_refs: {
      - {
        - to: {
          "default-domain",
          "admin",
          "right",
          "service-ace7ae00-56e3-42d1-96ec-7fe77088d97f-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"
  }
}
```

5018723

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_ShowBgplInstanceConfigReq?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	
service_chain_infos					static_routes		aggregate_routes	
family	routing_instance	chain_address	prefixes	service_instance			prefix	nexthop
inet	default-domain:admin:left:left	1.1.1.3	prefixes 1.1.1.0/24	default-domain:admin:si-aggregate			1.1.1.0/24	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

See the following example.

service_chain_infos					static_routes		aggregate_routes	
service_chain_infos					static_routes		aggregate_routes	
family	routing_instance	chain_address	prefixes	service_instance			prefix	nexthop
inet	default-domain:admin:left:left	1.1.1.3	prefixes 1.1.1.0/24	default-domain:admin:si-aggregate			1.1.1.0/24	1.1.1.3

You can also check the route table for the aggregate route in the right VN BGP table. For example:

http://<control-node>:8083/Snh_ShowRouteReq?x=default-domain:admin:right:right.inet.0

See the following example.

routes									
routes									
prefix	last_modified		paths						
1.1.1.0/24	2016-Feb-18 05:00:29.211876		paths						
protocol	last_modified		local_preference	local_as	peer_as	peer_router_id	source_as	path	next_hop
Aggregate	2016-Feb-18 05:00:29.211876		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.

ERROR: Unresolved graphic fileref="s018727.png" not found in
 "-//cmsxml/default/main/supplemental/STAGING/images/".

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"> • exact • orlonger • longer <p>Example: 1.1.0.0/16 orlonger</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 no export or no reoriginate, 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.</p> <p>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	<p>As part of the policy update, the following actions can be taken for community:</p> <ul style="list-style-type: none"> • Add a list of community to the existing community. • Set a list of community. • Remove a list of community (if present) from the existing community.
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

IN THIS SECTION

- [Create Routing Policy | 137](#)
- [Configure Service Instance | 138](#)
- [Configure the Network Policy for the Service Chain | 138](#)

This section shows how to create a routing policy for service chains and how to validate the policy.

Create Routing Policy

First, create the routing policy, **Configure > Networking > Routing > Create > Routing Policy**. See the following example.

The screenshot shows a 'Create Routing Policy' dialog box. The 'Name' field is set to 'failover'. Under 'Term(s)', there is one term: 'from: { prefix 2.2.2.0/24 orlonger } then: { local-preference 200 }'. The 'From' section shows 'prefix' selected, with '2.2.2.0/24' in the adjacent field, and 'orlonger' selected in a dropdown. The 'Then' section shows 'local-preference' selected, with '200' in the adjacent field. At the bottom right are 'Cancel' and 'Save' buttons. A vertical label 's018729' is visible on the right side of the dialog.

NOTE: The Contrail UI and REST APIs enable you to configure a BGP routing policy and then assign it to a virtual network, but the routing policy will not be applied if the virtual network is attached to an L3VPN.

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.

Action	Protocol	Source	Ports	Direction	Destination	Ports	Log	Services	Mirror
PASS	ANY	left	ANY	<=>	right	ANY	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

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-7fe7708d97f-default-domain_admin_ha-chain"
      ],
      href: "http://nodes27.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-7fe7708d97f-default-domain_admin_ha-chain"
      ],
      href: "http://nodes27.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://nodes27.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

See the following example.

5018745

routing_policies													
name	generation	ref_count	terms	deleted									
default-domain:admin:failover	0	2	<table><tr><td>terminal</td><td>matches</td><td>actions</td></tr><tr><td>true</td><td><table><tr><td>prefix [2.2.2.0/24 orlonger]</td></tr></table></td><td><table><tr><td>accept</td></tr><tr><td>local-pref 200</td></tr></table></td></tr></table>	terminal	matches	actions	true	<table><tr><td>prefix [2.2.2.0/24 orlonger]</td></tr></table>	prefix [2.2.2.0/24 orlonger]	<table><tr><td>accept</td></tr><tr><td>local-pref 200</td></tr></table>	accept	local-pref 200	false
terminal	matches	actions											
true	<table><tr><td>prefix [2.2.2.0/24 orlonger]</td></tr></table>	prefix [2.2.2.0/24 orlonger]	<table><tr><td>accept</td></tr><tr><td>local-pref 200</td></tr></table>	accept	local-pref 200								
prefix [2.2.2.0/24 orlonger]													
accept													
local-pref 200													
default-domain:default-project:default-routing-policy	0	0	<table><tr><td>terms</td></tr></table>	terms	false								
terms													

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

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>

See the following example.

service_chain_info					static_routes aggregate_routes routing_policies		
2 service_chain_info					static_routes	aggregate_routes	routing_policies
family	routing_instance	chain_address	prefixes	service_instance	static_routes	aggregate_routes	routing_policies
inet	default-domain:admin:right:right	1.1.1.6	prefixes	default-domain:admin:ha-chain			policy_name sequence
			2.2.2.0/24				default-domain:admin:failover 1

You can also verify the routing policy on the routing instance operational object.

Point your browser to:

http://<control-node>:8083/Snh_ShowRoutingInstanceReq?x=<name-of-internal-vrf>

See the following example.

routing_policies	
routing_policies	
policy_name	generation
default-domain:admin:failover	0

Control for Route Reorigination

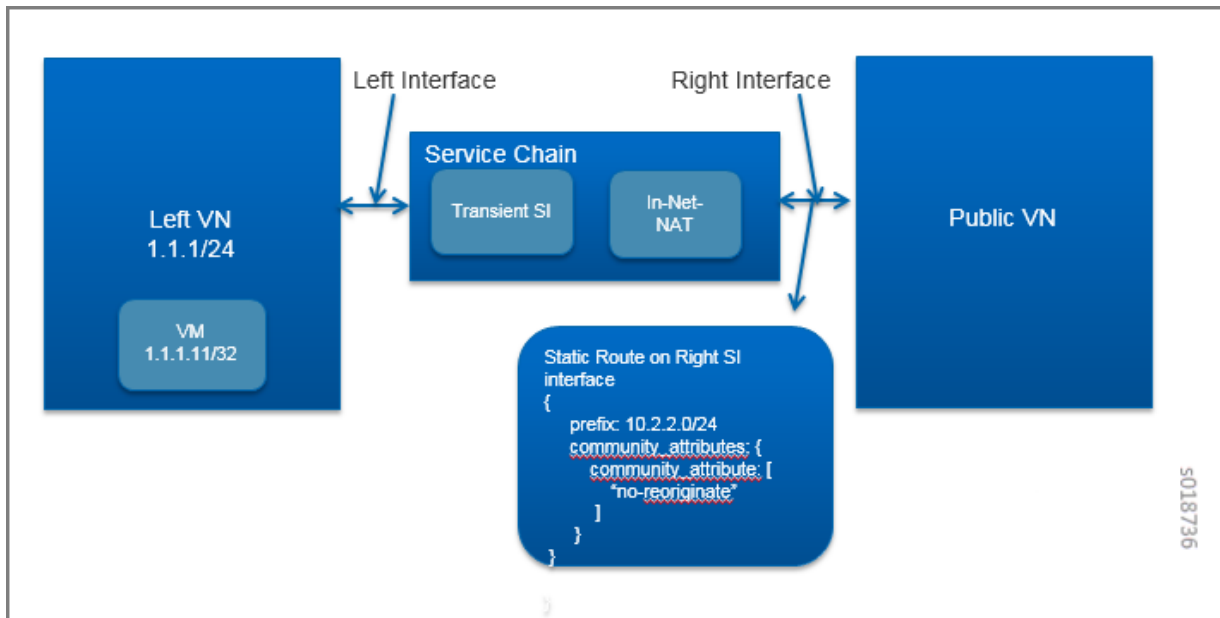
The ability to prevent reorigination of interface static routes is typically required when routes are configured on an interface that belongs to a service VM.

As an example, the following image shows a service chain that has multiple service instances, with an **in-net-nat** service instance as the last service VM, also with the right VN as the public VN.

The last service instance performs NAT by using a NAT pool. The right interface of the service VM must be configured with an interface static route for the NAT pool so that the destination in the right VN knows how to reach addresses in the NAT pool. However, the NAT pool prefix should not be reoriginated into the left VN.

To prevent route reorigination, the interface static route is tagged with a well-known BGP community called **no-reoriginate**.

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.

Create Service Instance

Name:

Service Template:

Interface Type:

Virtual Network:

Interface Type:

Virtual Network:

+ Add Static Routes

Prefix	Next Hop	Community
<input type="text" value="10.2.2.0/24"/>	<input type="text" value="Interface 2"/>	<input type="text" value="no-reoriginate"/>

Routing Policy

Interface Type:

Cancel Save

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_href: "http://nodeg2:8082/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

Example: Creating an In-Network Service Chain by Using Contrail Command

IN THIS SECTION

- [Hardware and Software Requirements | 145](#)
- [Overview | 145](#)
- [Configuration | 145](#)

This example provides instructions to create an in-network service chain by using the Contrail Command user interface (UI).

Hardware and Software Requirements

The following are the minimum requirements needed:

Hardware

- Processor: 4 core x86
- Memory: 32GB RAM
- Storage: at least 128GB hard disk

Software

- Contrail Release 3.2 or later

NOTE: For Contrail Networking Release 3.2 through Release 4.1, you use the Contrail Web UI. For more information, see [Example: Creating an In-Network Service Chain by Using Contrail Web UI](#).

Overview

A service chain is a set of services that are connected across networks. A service chain consists of service instances, left and right virtual networks, and a service policy attached to the networks. A service chain can have in-network services, in-network-nat services, and transparent services.

In an in-network service chain, packets are routed between service instance interfaces. When a packet is routed through the service chain, the source address of the packet entering the left interface of the service chain and source address of the packet exiting the right interface is the same. For more information, see [“Service Chaining” on page 111](#).

Configuration

IN THIS SECTION

- [Create Virtual Network | 146](#)
- [Create Virtual Machine | 147](#)

- [Configure Service Template | 148](#)
- [Add Service Instance | 149](#)
- [Create Service Policy | 150](#)
- [Attach Service Policy | 151](#)
- [Launch Virtual Machine | 152](#)

These topics provide instructions to create an in-network service chain.

Create Virtual Network

Step-by-Step Procedure

Use the Contrail Command UI to create a left virtual network, right virtual network, and management virtual network.

To create a left virtual network:

1. Click **Overlay>Virtual Networks**.

The All Networks page is displayed.

2. Click **Create** to create a network.

The Create Virtual Network page is displayed.

3. In the **Name** field enter **test-left-VN** for the left virtual network.

4. Select **(Default) User defined subnet only** from the **Allocation Mode** list.

5. Click **+Add** in the Subnets section to add subnets.

In the row that is displayed,

- a. Click the arrow in the Network IPAM field and select **left-ipam** for the left virtual network.

For the right virtual network, select **right-ipam** and for the management network, select **mgmt-ipam**.

NOTE: Management network is not used to route packets. This network is used to help debug issues with the virtual machine.

6. Enter **192.0.2.0/24** in the **CIDR** field.

7. Click **Create**.

The All Networks page is displayed. All virtual networks that you created are displayed in this page.

Repeat steps 2 through 7 to create the right virtual network (**test-right-VN**) and management virtual network (**test-mgmt-VN**).

Create Virtual Machine

Step-by-Step Procedure

Follow these steps to create a left virtual machine by using the Contrail Command UI.

1. Click **Workloads > Instances**.

The Instances page is displayed.

2. Click **Create**.

The Create Instance page is displayed.

3. Select **Virtual Machine** option button as the serve type.

4. Enter **test-left-VM** for the left virtual machine in the **Instance Name** field.

5. Select **Image** as the boot source from the **Select Boot Source** list.

NOTE: vSRX image with M1.large flavor is recommended for in-network virtual machine.

6. Select **vSRX image** file from the **Select Image** list.
7. Select **M1.large** flavor from the **Select Flavor** list.
8. Select the network you want to associate with the left virtual machine by clicking > next to the name of the virtual machine listed in the Available Networks table.

For the left virtual machine, select **test-left-VN**. For the right virtual machine, select **test-right-VN**. For the management virtual machine, select **test-mgmt-VN**.

The network is added to the Allocated Networks table.

9. Select **nova** from the **Availability Zone** list.

NOTE: You can choose any other availability zone.

10. Select **5** from the **Count (1-10)** list.

NOTE: You can choose any value from 1 through 10.

11. Click **Create** to launch the left virtual machine instance.

The Instances page is displayed. The virtual machine instances that you created are listed on the Instances page.

Repeat steps 2 through 11 to create right virtual machine instance (**test-right-VM**) and management virtual machine instance (**test-mgmt-VM**).

Configure Service Template

Step-by-Step Procedure

Follow these steps to create a service template by using the Contrail Command UI:

1. Click **Services>Catalog**.

The VNF Service Templates page is displayed.

2. Click **Create**.

The Create VNF Service Template page is displayed.

3. Enter **test-service-template** in the **Name** field.
4. Select **v2** as the version type.

NOTE: Starting with Release 3.2, Contrail supports only *Service Chain Version 2 (v2)*.

5. Select **Virtual Machine** as the virtualization type.
6. Select **In-Network** as the service mode.
7. Select **Firewall** as the service type.
8. From the Interface section,
 - Select **left** as the interface type from the **Interface Type** list.
 - Click + **Add**.

The Interface Type list is added to the table.

Select **right** as the interface type.

- Click + **Add** again.
- Another Interface Type list is added to the table.
- Select **management** as the interface type.

NOTE: The interfaces created on the virtual machine must follow the same sequence as that of the interfaces in the service template.

9. Click **Create** to create the service template.

The VNF Service Templates page is displayed. The service template that you created is displayed in the VNF Service Templates page.

Add Service Instance

Step-by-Step Procedure

Follow these steps to add a service instance by using the Contrail Command UI:

1. Click **Services>Deployments**.

The VNF Service Instances page is displayed.

2. Click **Create**.

The Create VNF Service Instance page is displayed.

3. Enter **test-service-instance** in the **Name** field.

4. Select **test-service-template - [in-network, (left, right, management)] - v2** from the **Service Template** list.

The **Interface Type** and **Virtual Network** fields are displayed.

5. Select the virtual network for each interface type as given below.

- **left**—Select the left virtual network (**test-left-VN**) that you created.
- **right**—Select the right virtual network (**test-right-VN**) that you created.
- **management**—Select the management virtual network (**test-management-VN**) that you created.

6. Click the **Port Tuples** section and click **+Add**.

Select the virtual machine instance for each interface type as given below.

- **left**—Select the left virtual machine instance that you created.
- **right**—Select the right virtual machine instance that you created.
- **management**—Select the management virtual machine instance that you created.

7. Click **Create** to create the service instance.

The VNF Service Instances page is displayed. The service instance that you created is displayed in the VNF Service Instances page.

Create Service Policy

Step-by-Step Procedure

Follow these steps to create a service policy by using the Contrail Command UI.

1. Click **Overlay > Network Policies**.

The Network Policies page is displayed.

2. Click **Create**.

The Network Policy tab of the Create Network Policy page is displayed.

3. Enter **test-network-policy** in the **Policy Name** field.
4. In the **Policy Rule(s)** section,
 - Select **pass** from the **Action** list.
 - Select **ANY** from the **Protocol** list.
 - Select **Network** from the **Source Type** list.
 - Select the **test-left-VN** from the **Source** list.
 - In the **Source Port** field, leave the default option, **Any**, as is.
 - Select **< >** from the **Direction** list.
 - Select **Network** from the **Destination Type** list.
 - Select the **test-right-VN** from the **Destination** list.
 - In the **Destination Ports** field, leave the default option, **Any**, as is.
5. Click **Create** to create the service policy.

The Network Policies page is displayed. All policies that you created are displayed in the Network Policies page.

Attach Service Policy

Step-by-Step Procedure

Follow these steps to attach a service policy:

1. Click **Overlay>Virtual Networks**.

The All networks page is displayed.

2. Attach service policy to the left virtual network (**test-left-VN**) and right virtual network (**test-right-VN**) that you created.

To attach service policy,

- a. Select the check box next to the name of the virtual network.
- b. Hover over to the end of the selected row and click the **Edit** icon.

The Edit Virtual Network page is displayed.

- c. Select the network policy from the Network Policies list.

3. Click **Save** to save the changes.

The Virtual Networks page is displayed.

Launch Virtual Machine

Step-by-Step Procedure

You can launch virtual machines from Contrail Command and test the traffic through the service chain by doing the following:

1. Launch the left virtual machine in left virtual network. See [“Create Virtual Machine” on page 147](#).
2. Launch the right virtual machine in right virtual network. See [“Create Virtual Machine” on page 147](#).
3. Ping the left virtual machine IP address from the right virtual machine.

Follow these steps to ping a virtual machine:

- a. Click **Workloads>Instances**.

The Instances page is displayed.

- b. Click the open console icon next to **test-right-VM**.

The Console page is displayed.

- c. Log in using root user credentials.

- d. Ping the left virtual machine IP address (**190.0.2.3**) from the Console.

See [Figure 54 on page 152](#) for a sample output.

Figure 54: Ping test-left-VM

```
root@test-right-vm:~# ping -c 5 192.0.2.3
PING 192.0.2.3 (192.0.2.3) 56(84) bytes of data:
64 bytes from 192.0.2.3: icmp_seq=1 ttl=63 time=0.238 ms
64 bytes from 192.0.2.3: icmp_seq=2 ttl=63 time=0.208 ms
64 bytes from 192.0.2.3: icmp_seq=3 ttl=63 time=0.231 ms
64 bytes from 192.0.2.3: icmp_seq=4 ttl=63 time=0.210 ms
64 bytes from 192.0.2.3: icmp_seq=5 ttl=63 time=0.210 ms

--- 192.0.2.3 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4006ms
rtt min/avg/max/mdev = 0.208/0.219/0.238/0.018 ms
root@test-right-vm:~#
```


RELATED DOCUMENTATION

[Service Chaining | 111](#)

[Example: Creating an In-Network-NAT Service Chain | 153](#)

[Example: Creating a Transparent Service Chain by Using Contrail Command | 161](#)

Example: Creating an In-Network-NAT Service Chain

IN THIS SECTION

- [Prerequisites | 153](#)
- [Overview | 154](#)
- [Configuration | 155](#)

This example provides instructions to create an in-network-nat service chain by using the Contrail Command user interface (UI).

Prerequisites

• Hardware and Software Requirements

Hardware

- Processor: 4 core x86
- Memory: 32GB RAM
- Storage: at least 128GB hard disk

Software

- Contrail Release 3.2 or later

NOTE: For Contrail Networking Release 3.2 through Release 4.1, you use the Contrail Web UI. For more information, see [Example: Creating a In-Network-NAT Service Chain by Using Contrail Web UI](#).

- **Create Network IPAM (IP Address Management)**

1. Click **Overlay>IPAM**.

The IP Address Management page is displayed.

2. Click **Create** to create a new network IPAM.

3. Enter a name for the IPAM in the name field.

4. Select **Default** from the DNS list.

5. Enter valid IP address in the NTP Server IP field.

6. Enter domain name in the Domain Name field.

7. Click **Create**.

The IP Address Management page is displayed.

Overview

A service chain is a set of services that are connected across networks. A service chain consists of service instances, left and right virtual networks, and a service policy attached to the networks. A service chain can have in-network services, in-network-nat services, and transparent services.

In an in-network-nat service chain, packets are routed between service instance interfaces. In-network-nat service chain does not require return traffic to be routed to the source network. When a packet is routed through the service chain, the source address of the packet entering the left interface of the service chain is updated and is not the same as the source address of the packet exiting the right interface. For more information, see [“Service Chaining” on page 111](#).

Configuration

IN THIS SECTION

- [Create Virtual Network | 155](#)
- [Create Virtual Machine | 156](#)
- [Configure Service Template | 157](#)
- [Add Service Instance | 158](#)
- [Create Service Policy | 159](#)
- [Attach Service Policy | 160](#)
- [Launch Virtual Machine | 160](#)

These topics provide instructions to create an in-network-nat service chain.

Create Virtual Network

Step-by-Step Procedure

Use the Contrail Command UI to create a left virtual network, right virtual network, and management virtual network.

To create a left virtual network:

1. Click **Overlay>Virtual Networks**.

The All Networks page is displayed.

2. Click **Create** to create a network.

The Create Virtual Network page is displayed.

3. In the **Name** field enter **test-left-VN** for the left virtual network.

4. Select **(Default) User defined subnet only** from the **Allocation Mode** list.

5. Click **+Add** in the Subnets section to add subnets.

In the row that is displayed,

- a. Select an IPAM for the virtual network from the Network IPAM list.

- b. Enter **192.0.2.0/24** in the **CIDR** field.

6. Click **Create**.

The All Networks page is displayed. All virtual networks that you created are displayed in this page.

NOTE: Management network is not used to route packets. This network is used to help debug issues with the virtual machine.

Repeat steps 2 through 7 to create the right virtual network (**test-right-VN**) and management virtual network (**test-mgmt-VN**).

Create Virtual Machine

Step-by-Step Procedure

Follow these steps to create a left virtual machine by using the Contrail Command UI.

1. Click **Workloads > Instances**.

The Instances page is displayed.

2. Click **Create**.

The Create Instance page is displayed.

3. Select **Virtual Machine** option button as the serve type.

4. Enter **test-left-VM** for the left virtual machine in the **Instance Name** field.

5. Select **Image** as the boot source from the **Select Boot Source** list.

NOTE: vSRX image with M1.large flavor is recommended for in-network-nat virtual machine.

6. Select **vSRX image** file from the **Select Image** list.

7. Select **M1.large** flavor from the **Select Flavor** list.

8. Select the network you want to associate with the left virtual machine by clicking > next to the name of the virtual machine listed in the Available Networks table.

For the left virtual machine, select **test-left-VN**. For the right virtual machine, select **test-right-VN**. For the management virtual machine, select **test-mgmt-VN**.

The network is added to the Allocated Networks table.

9. Select **nova** from the **Availability Zone** list.

NOTE: You can choose any other availability zone.

10. Select **5** from the **Count (1-10)** list.

NOTE: You can choose any value from 1 through 10.

11. Click **Create** to launch the left virtual machine instance.

The Instances page is displayed. The virtual machine instances that you created are listed on the Instances page.

Repeat steps 2 through 11 to create right virtual machine instance (**test-right-VM**) and management virtual machine instance (**test-mgmt-VM**).

Configure Service Template

Step-by-Step Procedure

Follow these steps to create a service template by using the Contrail Command UI:

1. Click **Services>Catalog**.

The VNF Service Templates page is displayed.

2. Click **Create**.

The Create VNF Service Template page is displayed.

3. Enter **test-service-template** in the **Name** field.

4. Select **v2** as the version type.

NOTE: Starting with Release 3.2, Contrail supports only *Service Chain Version 2 (v2)*.

5. Select **Virtual Machine** as the virtualization type.
6. Select **In-Network Nat** as the service mode.
7. Select **Firewall** as the service type.
8. From the Interface section,
 - Select **left** as the interface type from the **Interface Type** list.
 - Click + **Add**.

The Interface Type list is added to the table.

Select **right** as the interface type.

- Click + **Add** again.

Another Interface Type list is added to the table.

Select **management** as the interface type.

NOTE: The interfaces created on the virtual machine must follow the same sequence as that of the interfaces in the service template.

9. Click **Create** to create the service template.

The VNF Service Templates page is displayed. The service template that you created is displayed in the VNF Service Templates page.

Add Service Instance

Step-by-Step Procedure

Follow these steps to add a service instance by using the Contrail Command UI:

1. Click **Services>Deployments**.

The VNF Service Instances page is displayed.

2. Click **Create**.

The Create VNF Service Instance page is displayed.

3. Enter **test-service-instance** in the **Name** field.

4. Select **test-service-template - [in-network-nat, (left, right, management)] - v2** from the **Service Template** list.

The **Interface Type** and **Virtual Network** fields are displayed.

5. Select the virtual network for each interface type as given below.
 - **left**—Select the left virtual network (**test-left-VN**) that you created.
 - **right**—Select the right virtual network (**test-right-VN**) that you created.
 - **management**—Select the management virtual network (**test-management-VN**) that you created.

6. Click the **Port Tuples** section and click **+Add**.

Select the virtual machine instance for each interface type as given below.

- **left**—Select the left virtual machine instance that you created.
- **right**—Select the right virtual machine instance that you created.
- **management**—Select the management virtual machine instance that you created.

7. Click **Create** to create the service instance.

The VNF Service Instances page is displayed. The service instance that you created is displayed in the VNF Service Instances page.

Create Service Policy

Step-by-Step Procedure

Follow these steps to create a service policy by using the Contrail Command UI.

1. Click **Overlay > Network Policies**.

The Network Policies page is displayed.

2. Click **Create**.

The Network Policy tab of the Create Network Policy page is displayed.

3. Enter **test-network-policy** in the **Policy Name** field.

4. In the **Policy Rule(s)** section,
 - Select **pass** from the **Action** list.
 - Select **ANY** from the **Protocol** list.
 - Select **Network** from the **Source Type** list.
 - Select the **test-left-VN** from the **Source** list.

- In the **Source Port** field, leave the default option, **Any**, as is.
- Select < > from the **Direction** list.
- Select **Network** from the **Destination Type** list.
- Select the **test-right-VN** from the **Destination** list.
- In the **Destination Ports** field, leave the default option, **Any**, as is.

5. Click **Create** to create the service policy.

The Network Policies page is displayed. All policies that you created are displayed in the Network Policies page.

Attach Service Policy

Step-by-Step Procedure

Follow these steps to attach a service policy:

1. Click **Overlay>Virtual Networks**.

The All networks page is displayed.

2. Attach service policy to the left virtual network (**test-left-VN**) and right virtual network (**test-right-VN**) that you created.

To attach service policy,

- a. Select the check box next to the name of the virtual network.
- b. Hover over to the end of the selected row and click the **Edit** icon.

The Edit Virtual Network page is displayed.

- c. Select the network policy from the Network Policies list.

3. Click **Save** to save the changes.

The Virtual Networks page is displayed.

Launch Virtual Machine

Step-by-Step Procedure

You can launch virtual machines from Contrail Command and test the traffic through the service chain by doing the following:

1. Launch the left virtual machine in left virtual network. See [“Create Virtual Machine” on page 156](#).

2. Launch the right virtual machine in right virtual network. See [“Create Virtual Machine” on page 156](#).

3. Ping the left virtual machine IP address from the right virtual machine.

Follow these steps to ping a virtual machine:

a. Click **Workloads>Instances**.

The Instances page is displayed.

b. Click the open console icon next to **test-right-VM**.

The Console page is displayed.

c. Log in using root user credentials.

d. Ping the left virtual machine IP address (**190.0.2.3**) from the Console.

RELATED DOCUMENTATION

[Service Chaining | 111](#)

[Example: Creating an In-Network Service Chain by Using Contrail Command | 144](#)

[Example: Creating a Transparent Service Chain by Using Contrail Command | 161](#)

Example: Creating a Transparent Service Chain by Using Contrail Command

IN THIS SECTION

- [Prerequisites | 162](#)
- [Overview | 163](#)
- [Configuration | 163](#)

This example provides step-by-step instructions to create a transparent service chain by using the Contrail Command user interface (UI).

Prerequisites

- **Hardware and Software Requirements**

Hardware

- Processor: 4 core x86
- Memory: 32GB RAM
- Storage: at least 128GB hard disk

Software

- Contrail Release 3.2 or later

NOTE: For Contrail Networking Release 3.2 through Release 4.1, you use the Contrail Web UI. For more information, see [Example: Creating a Transparent Service Chain by Using Contrail Web UI](#).

- **Create Network IPAM (IP Address Management)**

1. Click **Overlay>IPAM**.

The IP Address Management page is displayed.

2. Click **Create** to create a new network IPAM.

3. Enter a name for the IPAM in the name field.

4. Select **Default** from the DNS list.

5. Enter valid IP address in the NTP Server IP field.

6. Enter domain name in the Domain Name field.

7. Click **Create**.

The IP Address Management page is displayed.

Overview

A service chain is a set of services that are connected across networks. A service chain consists of service instances, left and right virtual networks, and a service policy attached to the networks. A service chain can have in-network services, in-network-nat services, and transparent services. A transparent service chain is used for services that do not modify packets that are bridged between service instance interfaces. For more information, see [“Service Chaining” on page 111](#).

Configuration

IN THIS SECTION

- [Create Primary Virtual Networks | 163](#)
- [Create Secondary Virtual Network | 164](#)
- [Create Service Virtual Machine | 165](#)
- [Create Virtual Machine | 166](#)
- [Configure Service Template | 167](#)
- [Add Service Instance | 168](#)
- [Create Service Policy | 169](#)
- [Attach Service Policy | 170](#)
- [Launch Virtual Machine | 171](#)

These topics provide instructions to create a transparent service chain.

Create Primary Virtual Networks

Step-by-Step Procedure

Use the Contrail Command UI to create three primary virtual networks: left virtual network, right virtual network, and management virtual network. You attach service policies to the primary virtual networks that you create.

Follow these steps To create a left virtual network:

1. Click **Overlay>Virtual Networks**.

The All Networks page is displayed.

2. Click **Create** to create a network.

The Create Virtual Network page is displayed.

3. In the **Name** field enter **test-left-VN** for the left virtual network.

4. Select **(Default) User defined subnet only** from the **Allocation Mode** list.

5. Click **+Add** in the Subnets section to add subnets.

In the row that is displayed,

- a. Select an IPAM for the virtual network from the Network IPAM list.
- b. Enter **192.0.2.0/24** in the **CIDR** field.

6. Click **Create**.

The All Networks page is displayed. All virtual networks that you created are displayed in this page.

NOTE: Management network is not used to route packets. This network is used to help debug issues with the virtual machine.

Repeat steps 2 through 7 to create the right virtual network (**test-right-VN**) and management virtual network (**test-mgmt-VN**).

Create Secondary Virtual Network

Step-by-Step Procedure

Use the Contrail Command UI to create three secondary virtual networks: left virtual network (**trans-left-VN**), right virtual network (**trans-right-VN**), and management virtual network (**trans-mgmt-VN**). You associate the secondary virtual network to the transparent service instance that you create. For more information on creating virtual networks, see [“Create Primary Virtual Networks” on page 163](#).

Create Service Virtual Machine

Step-by-Step Procedure

Follow these steps to create a service virtual machine (SVM) by using the Contrail Command UI.

1. Click **Workloads > Instances**.

The Instances page is displayed.

2. Click **Create**.

The Create Instance page is displayed.

3. Select **Virtual Machine** option button as the serve type.

4. Enter **test-SVM** in the **Instance Name** field.

5. Select **Image** as the boot source from the **Select Boot Source** list.

NOTE: vSRX image with M1.large flavor is recommended for in-network virtual machine.

6. Select **vSRX image** file from the **Select Image** list.

7. Select **M1.large** flavor from the **Select Flavor** list.

8. From the Available Networks table, select **trans-left-VN**, **trans-right-VN**, and **trans-mgmt-VN** networks that you want to associate with the SVM by clicking > next to the name of the virtual machine.

The network is added to the Allocated Networks table.

9. Select **nova** from the **Availability Zone** list.

NOTE: You can choose any other availability zone.

10. Select **5** from the **Count (1-10)** list.

NOTE: You can choose any value from 1 through 10.

11. Click **Create** to launch the left virtual machine instance.

The Instances page is displayed. The virtual machine instances that you created are listed on the Instances page.

Create Virtual Machine

Use the Contrail Command UI to create two virtual machines: left virtual machine (**test-left-VM**) and right virtual machine (**test-right-VM**).

Step-by-Step Procedure

Follow these steps to create a left virtual machine by using the Contrail Command UI.

1. Click **Workloads > Instances**.

The Instances page is displayed.

2. Click **Create**.

The Create Instance page is displayed.

3. Select **Virtual Machine** option button as the serve type.

4. Enter **test-left-VM** for the left virtual machine in the **Instance Name** field.

5. Select **Image** as the boot source from the **Select Boot Source** list.

NOTE: vSRX image with M1.large flavor is recommended for in-network virtual machine.

6. Select **vSRX image** file from the **Select Image** list.

7. Select **M1.large** flavor from the **Select Flavor** list.

8. From the Available Networks table, select **test-left-VN** network that you want to associate with the left virtual machine by clicking > next to the name of the virtual machine.

For the right virtual machine, select **test-right-VN**.

The network is added to the Allocated Networks table.

9. Select **nova** from the **Availability Zone** list.

NOTE: You can choose any other availability zone.

10. Select **5** from the **Count (1-10)** list.

NOTE: You can choose any value from 1 through 10.

11. Click **Create** to launch the left virtual machine instance.

The Instances page is displayed. The virtual machine instances that you created are listed on the Instances page.

Repeat steps 2 through 11 to create right virtual machine instance (**test-right-VM**).

Configure Service Template

Step-by-Step Procedure

Follow these steps to create a service template by using the Contrail Command UI:

1. Click **Services>Catalog**.

The VNF Service Templates page is displayed.

2. Click **Create**.

The Create VNF Service Template page is displayed.

3. Enter **test-service-template** in the **Name** field.

4. Select **v2** as the version type.

NOTE: Starting with Release 3.2, Contrail supports only *Service Chain Version 2 (v2)*.

5. Select **Virtual Machine** as the virtualization type.

6. Select **Transparent** as the service mode.
7. Select **Firewall** as the service type.
8. From the Interface section,
 - Select **left** as the interface type from the **Interface Type** list.
 - Click + **Add**.

The Interface Type list is added to the table.

Select **right** as the interface type.

- Click + **Add** again.
- Another Interface Type list is added to the table.
- Select **management** as the interface type.

NOTE: The interfaces created on the virtual machine must follow the same sequence as that of the interfaces in the service template.

9. Click **Create** to create the service template.

The VNF Service Templates page is displayed. The service template that you created is displayed in the VNF Service Templates page.

Add Service Instance

Step-by-Step Procedure

Follow these steps to add a service instance by using the Contrail Command UI:

1. Click **Services>Deployments**.
The VNF Service Instances page is displayed.
2. Click **Create**.
The Create VNF Service Instance page is displayed.
3. Enter **test-service-instance** in the **Name** field.
4. Select **test-service-template - [transparent, (left, right, management)] - v2** from the **Service Template** list.

The **Interface Type** and **Virtual Network** fields are displayed.

5. Select the virtual network for each interface type as given below.

- **left**—Select **trans-left-VN** virtual network that you created.
- **right**—Select the **trans-right-VN** virtual network that you created.
- **management**—Select the **trans-mgmt-VN** virtual network that you created.

6. Click the **Port Tuples** section and click **+Add**.

Select the virtual machine instance for each interface type as given below. The port tuples should match the interfaces of the SVM. See [“Create Service Virtual Machine” on page 165](#).

- **left**—Select the left virtual machine instance that you created.
- **right**—Select the right virtual machine instance that you created.
- **management**—Select the management virtual machine instance that you created.

7. Click **Create** to create the service instance.

The VNF Service Instances page is displayed. The service instance that you created is displayed in the VNF Service Instances page.

Create Service Policy

Step-by-Step Procedure

Follow these steps to create a service policy by using the Contrail Command UI.

1. Click **Overlay > Network Policies**.

The Network Policies page is displayed.

2. Click **Create**.

The Network Policy tab of the Create Network Policy page is displayed.

3. Enter **test-network-policy** in the **Policy Name** field.

4. In the **Policy Rule(s)** section,

- Select **pass** from the **Action** list.
- Select **ANY** from the **Protocol** list.
- Select **Network** from the **Source Type** list.
- Select the **test-left-VN** from the **Source** list.
- In the **Source Port** field, leave the default option, **Any**, as is.

- Select < > from the **Direction** list.
 - Select **Network** from the **Destination Type** list.
 - Select the **test-right-VN** from the **Destination** list.
 - In the **Destination Ports** field, leave the default option, **Any**, as is.
5. Click **Create** to create the service policy.
- The Network Policies page is displayed. All policies that you created are displayed in the Network Policies page.

Attach Service Policy

Step-by-Step Procedure

Follow these steps to attach a service policy:

1. Click **Overlay>Virtual Networks**.
- The All networks page is displayed.
2. Select the **test-left-VN** network that you want to edit, and click the **Edit** icon.
- The Edit Virtual Network page is displayed.

NOTE: For the right virtual network, edit **test-right-VN**.

3. Select **test-network-policy** from the Network Policies list.
 4. Click **Save** to save the changes.
- The Virtual Networks page is displayed.
- Repeat steps 2 through 4 to attach the service policy to **test-right-VN**.

Launch Virtual Machine

You launch the SVM from the virtual networks that are not defined in the service policy.

- Launch the SVM in the **trans-left-VN** virtual network.
- Launch the SVM in the **trans-right-VN** virtual network.

RELATED DOCUMENTATION

[Service Chaining | 111](#)

[Example: Creating an In-Network-NAT Service Chain | 153](#)

[Example: Creating an In-Network Service Chain by Using Contrail Command | 144](#)

Using Static Routes with Services

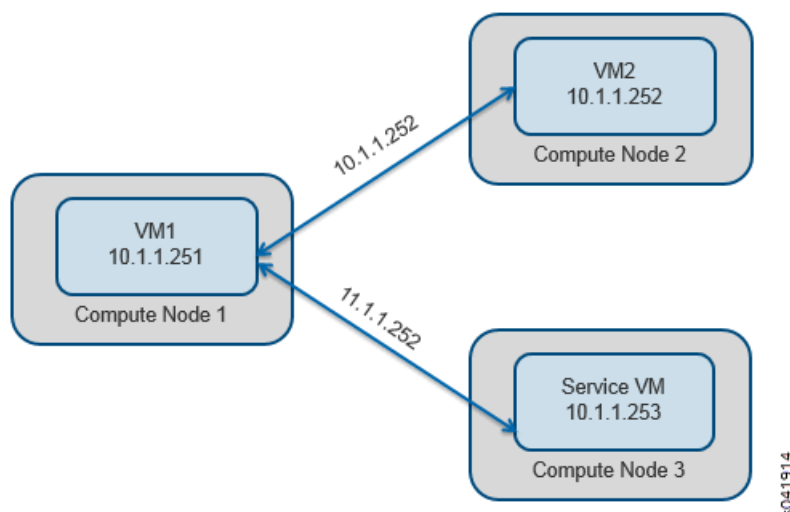
IN THIS SECTION

- [Static Routes for Service Instances | 171](#)
- [Configuring Static Routes on a Service Instance | 172](#)
- [Configuring Static Routes on Service Instance Interfaces | 173](#)
- [Configuring Static Routes as Host Routes | 175](#)

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.

Create Service Instances

Instance Name

nat

Services Template

nat - [in-network (management, left, right)]

Interface 1

Management

Auto Configured

Interface 2

Left

vn1

Static Routes

Prefix

Next hop

+

11.1.1.0/24

Interface 2

+ -

Interface 3

Right

vn2

Static Routes

Cancel

Save

5041916

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 with two interface configurations:

- Interface 2 (Left):**
 - Static Routes table:

Prefix	Next hop	
11.1.1.0/24	Interface 2	+ -
- Interface 3 (Right):**
 - Static Routes table:

Prefix	Next hop	
10.1.1.0/24	Interface 3	+ -

Buttons at the bottom: Cancel, Save. A vertical scrollbar on the right shows the value s041917.

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.

Create Network

Address Management

ipam1

IP Block

Gateway

+

-

IPAM	IP Block	Gateway
ipam1	1.2.3.0/24	1.2.3.254

Route Targets

Floating IP Pools

Host Routes

IPAM	Route Prefix	+	-
ipam1	1.1.1.0/24	+	-
ipam1	2.2.2.0/24	+	-

Cancel

Save

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
```

6

CHAPTER

Optimizing Contrail

Source Network Address Translation (SNAT) | 179

Source Network Address Translation (SNAT)

IN THIS SECTION

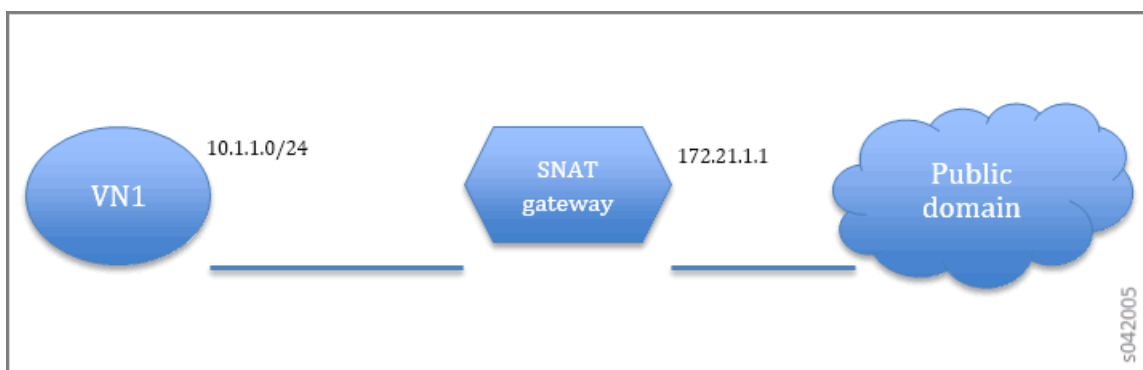
- [Overview | 179](#)
- [Neutron APIs for Routers | 180](#)
- [Network Namespace | 181](#)
- [Using the Web UI to Configure Routers with SNAT | 181](#)
- [Using the Web UI to Configure Distributed SNAT | 182](#)

Overview

Source Network Address Translation (source-nat or SNAT) allows traffic from a private network to go out to the internet. Virtual machines launched on a private network can get to the internet by going through a gateway capable of performing SNAT. The gateway has one arm on the public network and as part of SNAT, it replaces the source IP of the originating packet with its own public side IP. As part of SNAT, the source port is also updated so that multiple VMs can reach the public network through a single gateway public IP.

The following diagram shows a virtual network with the private subnet of 10.1.1.0/24. The default route for the virtual network points to the SNAT gateway. The gateway replaces the source-ip from 10.1.1.0/24 and uses its public address 172.21.1.1 for outgoing packets. To maintain unique NAT sessions the source port of the traffic also needs to be replaced.

Figure 55: Virtual Network With a Private Subnet



Neutron APIs for Routers

OpenStack supports SNAT gateway implementation through its Neutron APIs for routers. The SNAT flag can be enabled or disabled on the external gateway of the router. The default is True (enabled).

The OpenContrail plugin supports the Neutron APIs for routers and creates the relevant service-template and service-instance objects in the API server. The service scheduler in OpenContrail instantiates the gateway on a randomly-selected virtual router. OpenContrail uses network namespace to support this feature.

Example Configuration: SNAT for Contrail

The SNAT feature is enabled on OpenContrail through Neutron API calls.

The following configuration example shows how to create a test network and a public network, allowing the test network to reach the public domain through the SNAT gateway.

1. Create the public network and set the router external flag.

```
neutron net-create public
neutron subnet-create public 172.21.1.0/24
neutron net-update public -- --router:external=True
```

2. Create the test network.

```
neutron net-create test
neutron subnet-create --name test-subnet test 10.1.1.0/24
```

3. Create the router with one interface in test.

```
neutron router-create r1
neutron router-interface-add r1 test-subnet
```

4. Set the external gateway for the router.

```
neutron router-gateway-set r1 public
```

Network Namespace

Setting the external gateway is the trigger for OpenContrail to set up the Linux network namespace for SNAT.

The network namespace can be cleared by issuing the following Neutron command:

```
neutron router-gateway-clear r1
```

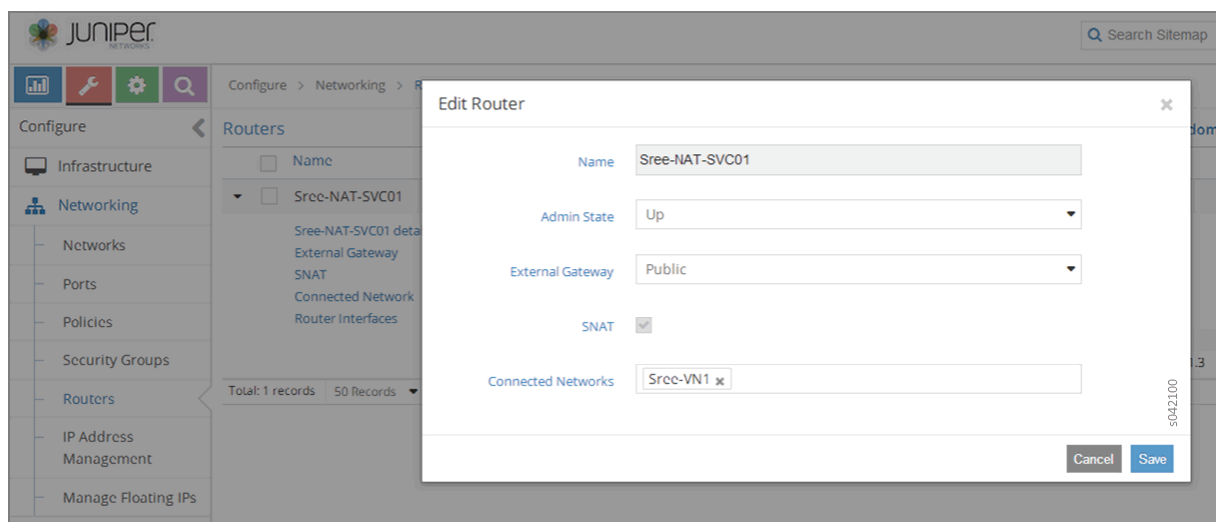
Using the Web UI to Configure Routers with SNAT

You can use the Contrail user interface to configure routers for SNAT and to check the SNAT status of routers.

To enable SNAT for a router, go to **Configure > Networking > Routers**. In the list of routers, select the router for which SNAT should be enabled. Click the Edit cog to reveal the **Edit Routers** window. Click the check box for SNAT to enable SNAT on the router.

The following shows a router for which SNAT has been **Enabled**.

Figure 56: Edit Router Window to Enable SNAT



When a router has been **Enabled** for SNAT, the configuration can be seen by selecting **Configure > Networking > Routers**. In the list of routers, click open the router of interest. In the list of features for that router, the status of SNAT is listed. The following shows a router that has been opened in the list. The status of the router shows that SNAT is **Enabled**.

Figure 57: Router Status for SNAT

Name	External Gateway	Connected Network	Admin State
Sree-NAT-SVC01	Public	Sree-VN1	Up

UUID	Network	IP
62a171f5-d323-4c70-942f-9cc02e973d53	Sree-VN1	100.1.1.3

You can view the real time status of a router with SNAT by viewing the instance console, as in the following.

Figure 58: Instance Details Window

```

Connected (unencrypted) to: QEMU (instance-00000057)
Inet addr:100.1.1.2 Bcast:100.1.1.255 Mask:255.255.255.0
Inet6 addr: fe80::9c:9aff:fe20:64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:5611 errors:0 dropped:0 overruns:0 frame:0
TX packets:5790 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:551725 (551.7 KB) TX bytes:564998 (564.9 KB)

lo
Link encap:Local Loopback
Inet addr:127.0.0.1 Mask:255.0.0.0
Inet6 addr: ::1/128 Scope:Host
UP LOOPBACK RUNNING MTU:16436 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)

ubuntu@sree-vm1:~$ ping 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_req=1 ttl=44 time=19.1 ms
^C
--- 8.8.8.8 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 19.107/19.107/19.107/0.000 ms
ubuntu@sree-vm1:~$

```

Using the Web UI to Configure Distributed SNAT

The distributed SNAT feature allows virtual machines to communicate with the IP fabric network using the existing forwarding infrastructure for compute node connectivity. This functionality is achieved through port address translation of virtual machine traffic using the IP address of the compute node as the public address.

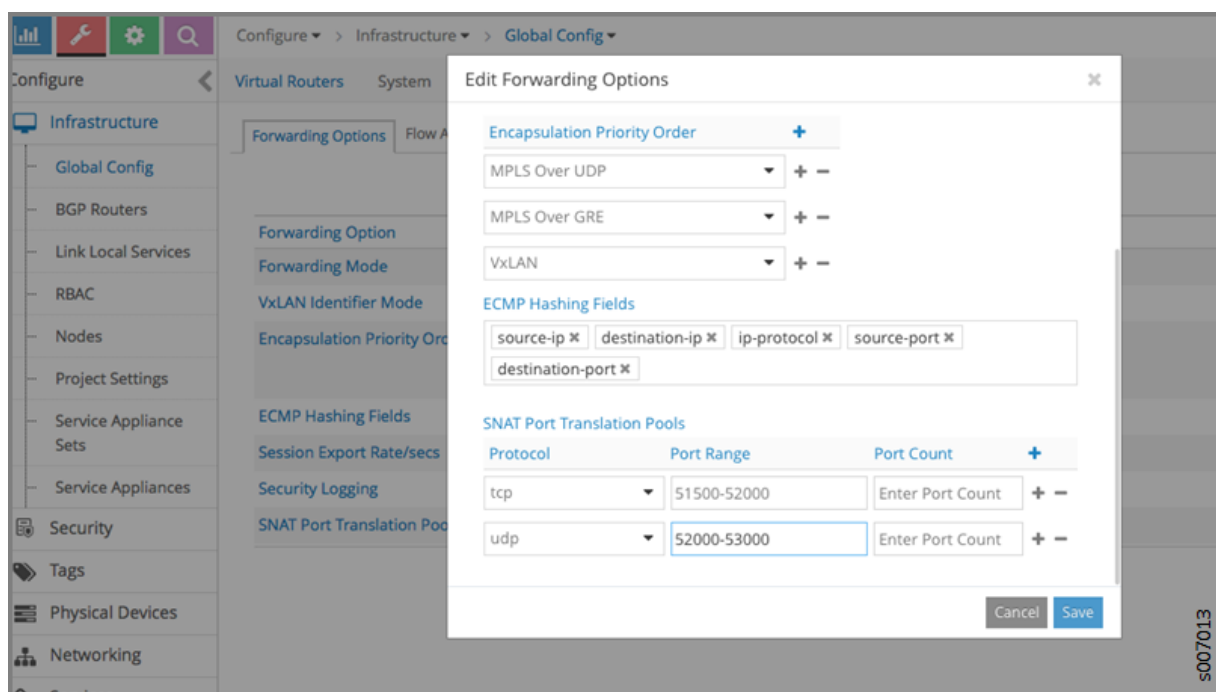
The following distributed SNAT use case is supported:

- Virtual networks with distributed SNAT enabled can communicate with the IP fabric network. The session must be initiated from a virtual machine. Sessions initiated from the external network are not supported.

Distributed SNAT is supported only for TCP and UDP, and you can configure discrete port ranges for both protocols.

A pool of ports is used for distributed SNAT. To create a pool of ports, go to **Configure > Infrastructure > Global Config**. The following shows an example of a port range used for port address translation.

Figure 59: Edit Forwarding Options Window



To use distributed SNAT, you must enable SNAT on the virtual network. To enable SNAT on the virtual network, go to **Configure > Networking > Networks**. The following shows a virtual network for which SNAT has been enabled under Advanced Options.

Figure 60: Create Window

