

Contrail® Networking

Fabric Lifecycle Management Guide

Published
2024-08-14

RELEASE
21.4

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

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

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

Contrail® Networking Fabric Lifecycle Management Guide

21.4

Copyright © 2024 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 This Guide | x

1

Overview

Understanding Underlay Management | 2

Fabric Lifecycle Management | 3

Fabric Overview | 4

2

Zero-Touch-Provisioning

Create a Fabric | 7

Provisioning Option - New Fabric | 8

Provisioning Option - Existing Fabric | 14

Discover a Device | 21

Assign a Role to a Device | 25

Assign Telemetry Profiles | 29

Delete a Fabric | 30

Provision Fabric Devices Using End-to-End ZTP | 32

3

Fabric Configuration

Image Management | 54

Upload a New Device Image | 54

Onboard Brownfield Devices | 57

Onboard Greenfield Devices | 67

Device Import | 78

Create Virtual Network | 82

Create Logical Routers | 90

Create Network Policy | 92

Create Network IPAM | 94

Reconfigure Roles | 96

Managing Custom Roles | 99

Adding Custom Roles | 99

Backup and Restore Custom Roles | 104

Backup Custom Roles | 105

Restore Custom Roles | 105

View Node Profile Information | 109

Monitoring Fabric Jobs | 110

Terminating Ongoing Fabric Jobs | 113

Adding a Leaf or Spine Device to an Existing Fabric Using ZTP | 115

Grouping Fabric Devices and Roles Using Device Functional Groups | 118

Creating Layer 3 PNF Service Chains for Inter-LR Traffic | 121

Onboard Fabric Devices | 123

Configure Virtual Networks | 123

Configure Virtual Port Groups | 124

Configure Logical Routers | 124

Configure PNF | 124

View Service Appliance Sets and Service Appliances | 127

Creating VNF Service Chains for Inter-LR Traffic | 128

Onboard Brownfield Devices | 132

Create Virtual Network | 142

Configuring Virtual Port Groups | 150

Create Logical Routers | 158

Configure the Internal Virtual Networks | 160

Create the Service Virtual Machine | 161

Create VNF Service Template | 161

Create VNF Service Instance | 162

Create the Network Policy | 163

Retaining the AS Path Attribute in a Service Chain | 164

Assisted Replication of Broadcast, Unknown Unicast, and Multicast Traffic | 165

Running Generic Device Operations Commands In Contrail Command | 168

Adding DHCP Server Information for Virtual Networks and Logical Routers | 173

Topology | 174

Steps to Add DHCP Server Information | 176

Adding DHCP Server Information to an Existing Logical Router | 176

Adding DHCP Server Information while Creating a Logical Router | 177

Steps to Remove CSN Information | 178

Return Material Authorization | 179

Move a Device to RMA State | 180

Replace a Device in RMA State with a New Device | 181

Getting Started with a New Device | 182

Approaches to Enable External Connectivity for Overlay Networks | 183

Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles | 185

Hardware Platforms and Associated Roles | 185

Hardware Platforms and Associated Node Profiles and Roles | 191

Managing Data Center Devices

Data Center Interconnect | 206

Understanding Data Center Interconnect | 206

Data Center Interconnect Deployment Topologies | 207

Creating Data Center Interconnect | 208

Onboard Brownfield Devices | 209

Create Virtual Network | 209

Create Logical Routers | 210

| Create DCI | 210

Logical Router Interconnect | 213

 Understanding Logical Router Interconnect | 213

 Creating Logical Router Interconnect | 215

 | Create a Fabric and Deploy Logical Routers on the Fabric Devices | 215

 | Create a Routing Policy for QFX Series Devices | 215

 | Creating Logical Router Interconnect | 217

Configuring Data Center Gateway | 220

 Configuring QFX Series Devices as Data Center Gateway | 220

 | Onboard Brownfield Devices | 221

 | Add Bare Metal Server | 221

 | Create Tenant Virtual Network | 223

 | Add CSN Nodes | 230

 | Create Logical Routers | 231

 | Verification | 233

 Configuring MX Series Routers as Data Center Gateway | 234

 | Onboard Brownfield Devices | 234

 | Create Virtual Network | 235

Virtual Port Groups | 236

Configuring Virtual Port Groups | 238

How to Detach and Attach a VMI to VPG | 246

Using Static, eBGP, PIM, and OSPF Protocols to Connect to Third-Party Network Devices | 258

 Overview | 259

 Steps to Connect to a Third-Party Device | 260

 | Topology | 260

 | Before You Begin | 261

 | Create Routed Virtual Networks | 262

 | Create Routed Virtual Port Groups | 264

 | Create Logical Routers | 267

Configuring Storm Control on Interfaces | 278

Creating Port Profiles, Storm Control Profiles, sFlow Profiles, or Telemetry Profiles by Cloning | 285

Configuring EVPN VXLAN Fabric with Multitenant Networking Services | 289

Edge-Routed Bridging for QFX Series Switches | 291

Activating Maintenance Mode on Data Center Devices | 293

Viewing the Network Topology | 295

Viewing Hardware Inventory of Data Center Devices | 303

Viewing Configuration of Devices Deployed in Contrail Fabric | 305

Detecting and Managing Manual CLI Configuration Changes | 308

Detecting a CLI Change | 308

Accept, Ignore, or Reject a CLI Change | 312

Certificate Lifecycle Management Using Red Hat Identity Management | 314

Fully Qualified Domain Names | 315

Performing Lifecycle Management of Certificates using Identity Management | 315

Collapsed Spine Architecture | 318

Support for Superspine Role | 320

5

High Availability in Contrail Networking

Using HA Cluster to Manage Fabric | 323

Hitless Software Upgrade of Data Center Devices Overview | 325

Performing Hitless Software Upgrade on Data Center Devices | 326

Fast Routing Convergence with Contrail Networking | 337

What is Convergence | 337

Fast Network Convergence in a Network Managed by Contrail Networking | 338

Configuring Fast Convergence from Contrail Command | 342

6

Integrating VMware with Contrail Networking Fabric

Understanding VMware-Contrail Networking Fabric Integration | 345

Deploying Contrail vCenter Fabric Manager Plug-in | 348

Prerequisites | 349

Deploying CVFM Plug-in while Provisioning Contrail Command | 349

Deploying CVFM Plug-in after Provisioning Contrail Command | 349

Troubleshooting Information | 350

Fabric Discovery and ESXi Discovery by Using Contrail Command | 351

Fabric Discovery | 352

ESXi Discovery | 357

Adding Distributed Port Groups | 358**Updating vCenter Credentials on Contrail Command | 360**

7

Integrating OpenStack with Contrail Networking Fabric**Understanding OpenStack-Contrail Networking Fabric Integration | 363****Deploying ML2 Plug-in with Red Hat OpenStack | 366**

Deploy Contrail Command and CFM without Orchestrator | 367

Configure Fabric by using Contrail Command | 368

Deploy RHOSP13 with ML2 Plug-in | 373

Configure Connectivity between RHOSP Internal API Network and Contrail Command Virtual Machines | 378

Add Red Hat OpenStack Orchestrator | 379

Create Swift Containers in OpenStack | 380

(Optional) Deploy AppFormix and sFlows | 381

Sample Network Files | 384

8

Extending Contrail Networking to Bare Metal Servers**Bare Metal Server Management | 392**

Understanding Bare Metal Server Management | 392

Features of the Bare Metal Server Management Framework | 394

How Bare Metal Server Management Works | 396

LAG and Multihoming Support | 398

Adding Bare Metal Server to Inventory | 400

Launching a Bare Metal Server | 402

Onboarding and Discovery of Bare Metal Servers | 404

Launching and Deleting a Greenfield Bare Metal Server | 405

Destination Network Address Translation for Bare Metal Servers | 406

Enabling DNAT in a Data Center Gateway | 407

Extending a Public Virtual Network to the Data Center Gateway | 408

Creating a Floating IP Address Pool | 408

Mapping Floating IP Address to the Fixed IP address of the BMS Private Network | 409

Troubleshooting Bare Metal Servers | 411

About This Guide

Use this guide to understand Contrail Networking underlay management and managing data center devices. This guide also provides information on integrating VMware with Contrail Networking fabric and extending Contrail Networking to bare metal servers.

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

Table 1: Contrail Networking Guides

Guide Name	Description
Contrail Networking Installation and Upgrade Guide	Provides step-by-step instructions to install and bring up Contrail and its various components.
Contrail Networking for Container Networking Environments User Guide	Provides information about installing and using Contrail Networking in containerized environments using Kubernetes orchestration.
Contrail Networking Fabric Lifecycle Management Guide	Provides information about Contrail underlay management and data center automation.
Contrail Networking and Security User Guide	Provides information about creating and orchestrating highly secure virtual networks.
Contrail Networking Service Provider Focused Features Guide	Provides information about the features that are used by service providers.
Contrail Networking Monitoring and Troubleshooting Guide	Provides information about Contrail Insights and Contrail analytics.

RELATED DOCUMENTATION

[README Access to Contrail Networking Registry 21xx](#)

[Contrail Networking Release Notes 21xx](#)

Tungsten Fabric Architecture Guide

Juniper Networks TechWiki: Contrail Networking

1

CHAPTER

Overview

[Understanding Underlay Management](#) | 2

[Fabric Lifecycle Management](#) | 3

[Fabric Overview](#) | 4

Understanding Underlay Management

IN THIS SECTION

- [Benefits of Underlay Management | 3](#)

A private cloud data center is a critical business infrastructure that enterprise customers and service providers need. These private cloud data centers help deliver automated application networking services to internal departments. Today, most enterprises and service providers are moving from a vendor proprietary fabric to a standard-based EVPN-VXLAN data center built on IP Clos technology. In an EVPN-VXLAN data center, the underlay network is the physical infrastructure (switches, routers, firewall) on which overlay network services are built.

An EVPN-VXLAN data center fabric relies on a standard model that consists of tenants. These tenants are a group of endpoints, where,

- groups are subnets that are routed to other groups.
- endpoints are bridged within a group.
- tenants are routed to other tenants depending on the overlay architecture.
- tenants, groups, and endpoints may have services such as security, transit, multihoming, and QoS associated with them.
- tenants and groups are implemented in the network as IP and Ethernet Virtual Private Networks (VPNs) and Virtual Tunnel End Points (VTEPs).

EVPN-VXLAN is used in a data center fabric to deliver multi-tenant networking services. The following network virtualization overlay architectures can be deployed in an EVPN-VXLAN IP fabric.

- Centrally-Routed Bridging overlay design—inter-VN routing occurs in either the spine switch or border leaf switch.
- Edge-Routed Bridging overlay design— inter-VN routing occurs natively in the leaf switch that workloads and servers are attached to.
- Ethernet overlays—Layer 2 reachability and workload mobility across endpoints are the main services that the data center fabric provides.
- IP overlay—traffic in a tenant is routed using IP routes.

Contrail Networking Release 5.0.1 supports the automation and management of EVPN-VXLAN data center IP fabric as well as the automation of layer 2 and layer 3 multi-tenant services on the IP fabric. The existing Contrail Networking configuration node can provide intent driven automation capabilities on physical network elements such as ToR and EoR switches, Spines, SDN gateway, and VPN gateways in the data center. In addition, you can perform basic device management functions such as image upgrade, device discovery, device underlay configuration, assigning roles to devices, and viewing node profile information from the node.

Benefits of Underlay Management

- Enables basic device management functions from the Contrail Networking configuration node.
- Enables underlay network automation.
- Supports zero-touch-provisioning (ZTP) of factory-default devices to form an IP Clos network.

NOTE: ZTP allows you to provision new devices in your network automatically, with minimal manual intervention.

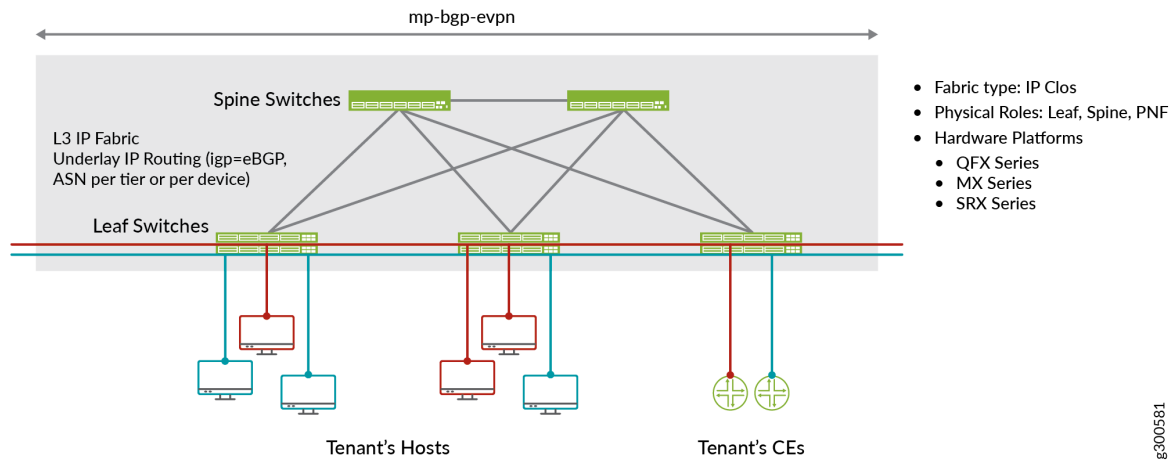
RELATED DOCUMENTATION

| [Fabric Overview](#) | 4

Fabric Lifecycle Management

You can onboard, configure, and manage a set of devices, and physical network functions (PNF) in Contrail Networking as an IP fabric. A fabric is a set of devices, and PNFs that fall under the same data center administrator responsibility area. The fabric is linked to different role-based access control (RBAC) profiles for ease of administration and management.

Figure 1: Sample Layer 3 IP Clos Fabric



Contrail Networking helps you provision both greenfield and brownfield devices to form IP CLOS networks. You can bring up all factory-default greenfield devices using zero-touch-provisioning to form an operational IP CLOS network with underlay connectivity. However, unlike greenfield devices, brownfield devices are manually provisioned before device onboarding.

RELATED DOCUMENTATION

[Understanding Underlay Management | 2](#)

[Understanding Bare Metal Server Management | 392](#)

[Configuring Data Center Gateway | 220](#)

Fabric Overview

You can manage a set of devices, and physical network functions (PNF) in Contrail Networking as a fabric. A fabric is a set of data center devices, and PNFs that fall under the same data center administrator responsibility area. The fabric is linked to different role-based access control (RBAC) profiles for ease of administration and management.

You can provision greenfield devices and brownfield devices by using the Contrail Command user interface (UI).

Greenfield devices

You can provision new devices to form an IP Clos network. These devices are connected to a management network that is provisioned before device onboarding. The greenfield fabric workflow then zero-touch-provisions all factory-default devices to form an operational IP Clos network with underlay connectivity.

This greenfield fabric workflow includes playbooks that automate the fabric data model creation in the database, DHCP server configuration, generating device bootstrap configuration, uploading device bootstrap configuration to TFTP server, device discovery, node profile auto-assignment, device role assignment, and role-based auto configuration.

Brownfield devices

You can provision legacy devices or existing devices to form an IP Clos network. Unlike greenfield devices, brownfield devices are manually provisioned before device onboarding. The brownfield fabric workflow includes playbooks that automate the fabric data model creation in the database. You can perform basic device management functions such as image upgrade, device discovery, device underlay configuration, assign roles to devices, and view node profile information.

You can use the Contrail Command UI to:

- ["Create a Fabric" on page 7](#)
- ["Discover a Device" on page 21](#)
- ["Assign a Role to a Device" on page 25](#)
- ["View Node Profile Information" on page 109](#)
- ["Delete a Fabric" on page 30](#)

2

CHAPTER

Zero-Touch-Provisioning

[Create a Fabric | 7](#)

[Discover a Device | 21](#)

[Assign a Role to a Device | 25](#)

[Assign Telemetry Profiles | 29](#)

[Delete a Fabric | 30](#)

[Provision Fabric Devices Using End-to-End ZTP | 32](#)

Create a Fabric

IN THIS SECTION

- [Provisioning Option - New Fabric | 8](#)
- [Provisioning Option - Existing Fabric | 14](#)

You can create a fabric by using the Contrail Command UI.

Follow these steps to create a fabric:

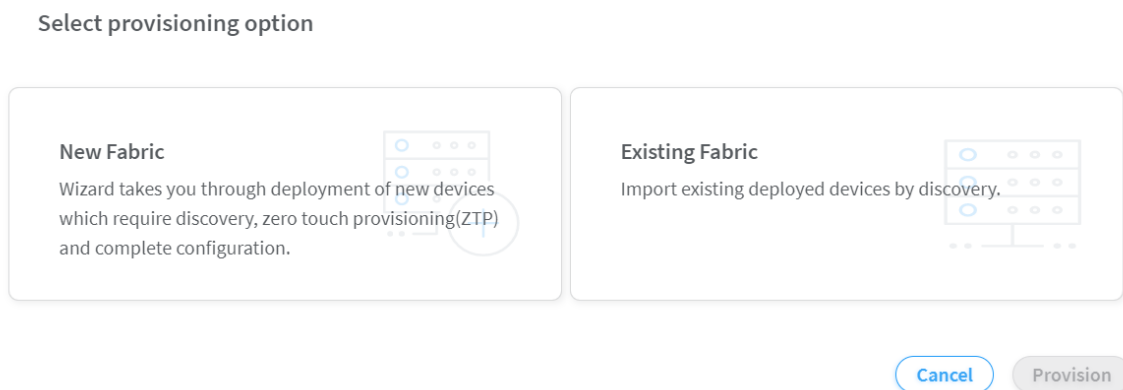
1. Click **Infrastructure**>**Fabrics**.

The Fabrics page is displayed.

2. Click **Create**.

You are prompted to select a provisioning option. See [Figure 2 on page 7](#).

Figure 2: Select Provisioning Option



- Click **New Fabric** to deploy new (greenfield) devices. See [Figure 3 on page 14](#).
- Click **Existing Fabric** to import existing (brownfield) devices by discovery. See [Figure 4 on page 20](#).

Click **Provision**.

The Create Fabric page is displayed.

If you select **New Fabric** as the provisioning option, see ["Provisioning Option - New Fabric" on page 8](#).

If you select **Existing Fabric** as the provisioning option, see ["Provisioning Option - Existing Fabric" on page 14](#).

Provisioning Option - New Fabric

You can use zero-touch-provisioning (ZTP) to deploy greenfield devices by using the Contrail Command UI.

Enter the information given in [Table 2 on page 8](#) if you have selected New Fabric as the provisioning option.

Table 2: Provisioning Option - New Fabric

Field	Action
Name	<p>Enter a name for the fabric.</p> <p>The name identifies the fabric on all fabric configuration and monitoring pages.</p>
Device credentials	<p>Enter root user password.</p> <p>The password entered in this field becomes the root password to access every device in the fabric.</p>
Overlay ASN (iBGP)	<p>Enter autonomous system (AS) number in the range of 1-65,535.</p> <p>If you enable 4 Byte ASN in Global Config, you can enter 4-byte AS number in the range of 1-4,294,967,295.</p>

Table 2: Provisioning Option - New Fabric (Continued)

Field	Action
Device Info	<p>Upload YAML file.</p> <p>This YAML file contains the serial numbers of each device in the fabric for device discovery. Click browse and navigate to the local directory and select the YAML file. Click Open to confirm.</p> <p>Alternatively, you can drag and drop the .yaml or .yml file in the Device Info box.</p> <p>To create this YAML file, click (* .yaml) in the Template field, download the file, modify the file to include the serial numbers and hostnames for your fabric devices, and save the file.</p> <p>For a sample YAML file, see "No Link Title" on page 13.</p>
Node profiles	<p>Add node profiles.</p> <p>You can add more than one node profile.</p> <p>All preloaded node profiles are added to the fabric by default. You can remove a node profile by clicking X on the node profile. For more information, see "View Node Profile Information" on page 109.</p> <p>For more information on supported hardware platforms, associated node profiles and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185.</p>
Upgrade devices during the process?	<p>Select the Upgrade devices during the process? check box as given in Figure 3 on page 14 to enable the OS Version list.</p> <p>Starting with Contrail Networking Release 1907, you can upgrade a device during the ZTP process.</p>

Table 2: Provisioning Option - New Fabric (Continued)

Field	Action
OS Version	<p>Select the OS version you want to upgrade the device to, from the OS Version list.</p> <p>The OS Version list is enabled when you select the Upgrade devices during the process? check box.</p> <p>NOTE: The options in the OS Version list are the OS versions of the images that you uploaded.</p>
Disable VLAN-VN Uniqueness Check	<p>Select this check box when you are using the enterprise style of configuration but want to disable the requirement that every VLAN ID must have a 1:1 mapping with a VNI. Enterprise style of configuration is enabled by selecting the VLAN-ID Fabric-Wide Significance check box.</p>
VLAN-ID Fabric Wide Significance	<p>Select the VLAN-ID Fabric Wide Significance check box to enable enterprise style of configuration for the CRB-Access role on QFX devices. Deselect the check box to enable service provider style of configuration for the CRB-Access role. The check box is selected by default since enterprise style is the default setting.</p> <p>Once configured you can modify the enterprise style setting to service provider style of configuration. However, you cannot modify the service provider style to enterprise style of configuration without having to recreate the fabric.</p> <p>NOTE: Contrail Networking Release 1909 supports QFX10002-60C devices running Junos OS Release 19.1R2 and later. QFX10002-60C device works only if enterprise style of configuration is enabled. To enable enterprise style of configuration, select the VLAN-ID Fabric Wide Significance check box when onboarding the QFX10002-60C device. For more information on enterprise style of configuration, see "Configuring EVPN VXLAN Fabric with Multitenant Networking Services" on page 289.</p> <p>For more information on supported hardware platforms and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185.</p>

Table 2: Provisioning Option - New Fabric (Continued)

Field	Action
Underlay ASNs (eBGP)	<p>Enter autonomous system (AS) number in the range of 1-65,535.</p> <p>If you enable 4 Byte ASN in Global Config, you can enter 4-byte AS number in the range of 1-4,294,967,295.</p> <ul style="list-style-type: none"> • Enter minimum value in ASN From field. • Enter maximum value in ASN To field.
Management subnets	<p>Enter the following information to auto-assign management IP addresses to devices:</p> <p>CIDR—Enter the block of IP addresses that will be assigned as management IP addresses. The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Gateway—Enter gateway address for the devices in the management subnet that connect to the fabric.</p>
Fabric subnets (CIDR)	<p>Enter fabric CIDR address. The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.</p>
Loopback subnets (CIDR)	<p>Enter loopback subnet (lo0) address.</p> <p>The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.</p>

Table 2: Provisioning Option - New Fabric *(Continued)*

Field	Action
LR Loopback subnets	<p>Enter an IP subnet to be assigned as loopback interface (lo0) addresses used in Logical Routers (LR). The LR loopback interface IP address is required for eBGP peering to external or unmanaged devices.</p> <p>The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p>
PNF Servicechain subnets	<p>Enter the IP subnet for allocating IP addresses in the PNF Servicechain subnets field to establish EBGP session between PNF device and SPINE switch. This is an optional field that should be left blank when you are not creating service chains.</p>
Advanced interface filters	<p>Create an interface filter to filter the interfaces to include in the fabric. By default, all interfaces identified as participating in Contrail are imported into the fabric during the fabric provisioning process. If an interface filter is set, the fabric provisioning process includes the interfaces that are participating in Contrail and that match the interface filter in the fabric.</p> <p>To create an interface filter, choose the operation as regex and enter the filter characters in the Expression field. The Expression field supports all characters - including metacharacters - allowed in Python regex filters. For example, you can enter ^xe in the Expression field to filter out all 10Gbps xe interfaces from the fabric.</p>

Table 2: Provisioning Option - New Fabric (*Continued*)

Field	Action
Import configured interfaces	<p>Choose this option if configured interfaces need to be imported into the fabric in addition to runtime interfaces. With some exceptions, a configured interface is generally an interface that has been configured in the Junos OS software.</p> <p>A runtime interface is generally an interface that has not been configured in Junos OS. You can confirm which interfaces are configured interfaces by entering the show interfaces command at the configuration mode prompt(#) in Junos. Only runtime interfaces are imported into the fabric by default.</p>

Sample YAML File Snippet

```

supplemental_day_0_cfg:
  - name: 'cfg1'
    cfg: |
      set system ntp server 167.XX.XX.XX
device_to_ztp:
  - serial_number: 'serial number'
    supplemental_day_0_cfg: 'cfg1'
    hostname: '<host name>'
    device_functional_group: 'dfg1'
  - serial_number: 'serial number'
    supplemental_day_0_cfg: 'cfg1'
  - serial_number: 'serial number'
  - serial_number: 'serial number'

```

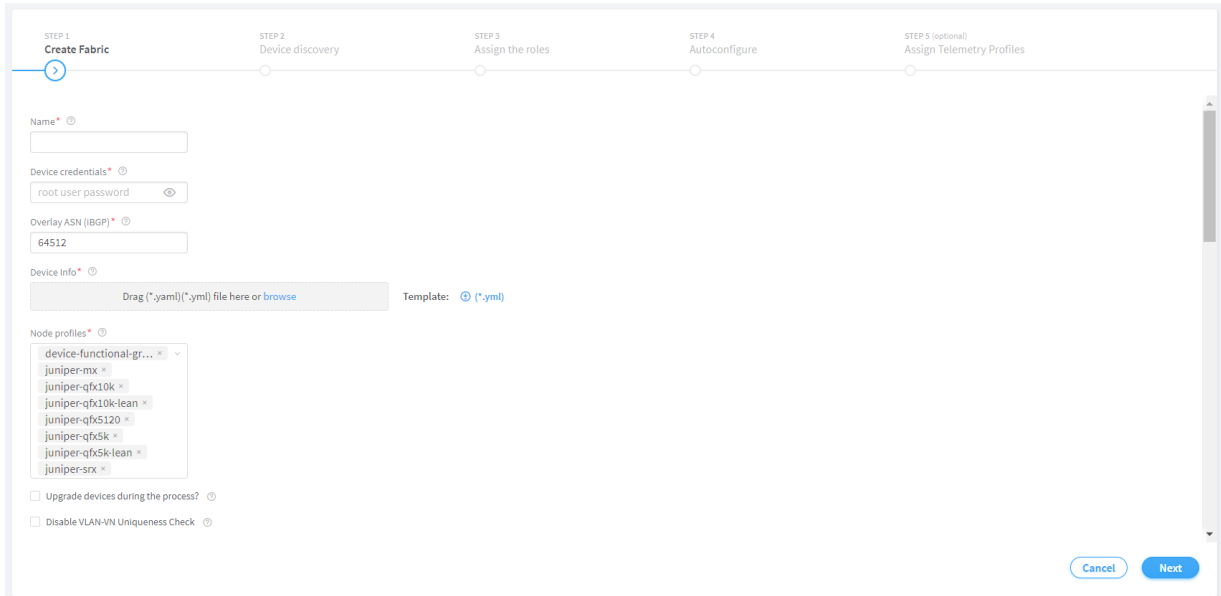
where,

supplemental_day_0_cfg is the additional configuration that is pushed on to the device during ZTP.

serial_number is the serial number of the device that is added to the fabric.

hostname is the device host name. If host name is not set, the serial number of the device is set as the device host name by default.

Figure 3: Deploy Greenfield Devices



Click **Next**.

The Discovered devices page is displayed.

Provisioning Option - Existing Fabric

Enter the information as given in [Table 3 on page 14](#) if you have selected Existing Fabric as the provisioning option.

Table 3: Provisioning Option - Existing Fabric

Field	Action
Name	Enter a name for the fabric.
Overlay ASN (iBGP)	Enter autonomous system (AS) number in the range of 1-65,535. If you enable 4 Byte ASN in Global Config , you can enter 4-byte AS number in the range of 1-4,294,967,295.

Table 3: Provisioning Option - Existing Fabric *(Continued)*

Field	Action
Node profiles	<p>Add node profiles.</p> <p>You can add more than one node profile.</p> <p>All preloaded node profiles are added to the fabric by default. You can remove a node profile by clicking X on the node profile. For more information, see "View Node Profile Information" on page 109.</p> <p>For more information on supported hardware platforms, associated node profiles and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185.</p>
Disable VLAN-VN Uniqueness Check	<p>Select this check box when you are using the enterprise style of configuration but want to disable the requirement that every VLAN ID must have a 1:1 mapping with a VNI. Enterprise style of configuration is enabled by selecting the VLAN-ID Fabric-Wide Significance check box.</p>

Table 3: Provisioning Option - Existing Fabric (Continued)

Field	Action
VLAN-ID Fabric Wide Significance	<p>Select the check box to enable enterprise style of configuration for the CRB-Access role on QFX devices. De-select the check box to enable service provider style of configuration for the CRB-Access role. The check box is selected by default since enterprise style is the default setting.</p> <p>Once configured you can modify the enterprise style setting to service provider style of configuration. However, you cannot modify the service provider style to enterprise style of configuration without having to recreate the fabric.</p> <p>The service provider style of configuration allows for customization of Ethernet-based services at the logical interface level. Each logical interface is bound to a unique VLAN ID. With the enterprise style of configuration, logical interfaces are placed into Layer 2 mode by specifying ethernet-switching as the interface family. The ethernet-switching family can be configured only on a single logical unit, unit 0. For more information on enterprise and service provider type of configurations, see Flexible Ethernet Services Encapsulation.</p> <p>NOTE: Contrail Networking Release 1909 supports QFX10002-60C device running Junos OS Release 19.1R2 and later. QFX10002-60C device works only if enterprise style of configuration is enabled. To enable enterprise style of configuration, select the VLAN-ID Fabric Wide Significance check box when onboarding the QFX10002-60C device. For more information on enterprise style of configuration, see "Configuring EVPN VXLAN Fabric with Multitenant Networking Services" on page 289. For more information on supported hardware platforms and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185.</p>

Table 3: Provisioning Option - Existing Fabric (Continued)

Field	Action
Device credentials	Enter the device credentials to access the fabric devices for discovery. If your fabric devices have different username and password combinations for device access, click the + Add option to add additional username and password credentials.
Management subnets	Enter the following information to auto-assign management IP addresses to devices: CIDR —Enter the block of IP addresses that will be assigned as management IP addresses. The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24. Gateway —Enter gateway address for the devices in the management subnet that connect to the fabric.
Loopback subnets	Enter loopback subnet (lo0) address. The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24. Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.
Underlay ASNs (eBGP)	Enter autonomous system (AS) number in the range of 1-65,535. If you enable 4 Byte ASN in Global Config , you can enter 4-byte AS number in the range of 1-4,294,967,295. <ul style="list-style-type: none"> • Enter minimum value in ASN From field. • Enter maximum value in ASN To field.

Table 3: Provisioning Option - Existing Fabric *(Continued)*

Field	Action
Fabric subnets	<p>Enter fabric CIDR address. The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.</p>
LR Loopback subnets	<p>Enter an IP subnet to be assigned as loopback interface (lo0) addresses used in Logical Routers (LR). The LR loopback interface IP address is required for eBGP peering to external or unmanaged devices.</p> <p>The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p>
Loopback subnets (CIDR)	<p>Enter loopback address.</p> <p>Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.</p> <p>If you assign the AR-Replicator and AR-Client roles to enable assisted replication on the QFX10000 devices in a datacenter, you must enter loopback address. For more information, see "Assign a Role to a Device" on page 25.</p>
PNF Servicechain subnets	<p>Enter the IP subnet for allocating IP addresses in the PNF Servicechain subnets field to establish EBGP session between PNF device and SPINE switch. This is an optional field that should be left blank when you are not creating service chains.</p>

Table 3: Provisioning Option - Existing Fabric *(Continued)*

Field	Action
Advanced interface filters	<p>Create an interface filter to filter the interfaces to include in the fabric. By default, all interfaces identified as participating in Contrail are imported into the fabric during the fabric provisioning process. If an interface filter is set, the fabric provisioning process includes the interfaces that are participating in Contrail and that match the interface filter in the fabric.</p> <p>To create an interface filter, choose the operation as regex and enter the filter characters in the Expression field. The Expression field supports all characters - including metacharacters - allowed in Python regex filters. For example, you can enter <code>^xe</code> in the Expression field to filter out all 10Gbps xe interfaces from the fabric.</p>
Import configured interfaces	<p>Choose this option if configured interfaces need to be imported into the fabric in addition to runtime interfaces. With some exceptions, a configured interface is generally an interface that has been configured in the Junos OS software.</p> <p>A runtime interface is generally an interface that has not been configured in Junos OS. You can confirm which interfaces are configured interfaces by entering the <code>show interfaces</code> command at the configuration mode prompt(<code>#</code>) in Junos. Only runtime interfaces are imported into the fabric by default.</p>

Figure 4: Import Brownfield Devices

Click **Next**.

The Device discovery page is displayed.

For more information on device discovery, see "[Discover a Device](#)" on page 21.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1909	Contrail Networking Release 1909 supports QFX10002-60C devices running Junos OS Release 19.1R2 and later.
1908	Select the VLAN-ID Fabric Wide Significance check box to enable enterprise style of configuration for the CRB-Access role on QFX devices. Deselect the check box to enable service provider style of configuration for the CRB-Access role.
1907	Starting with Contrail Networking Release 1907, you can upgrade a device during the ZTP process.

RELATED DOCUMENTATION

[Discover a Device](#) | 21

[Assign a Role to a Device](#) | 25

[View Node Profile Information | 109](#)

[Delete a Fabric | 30](#)

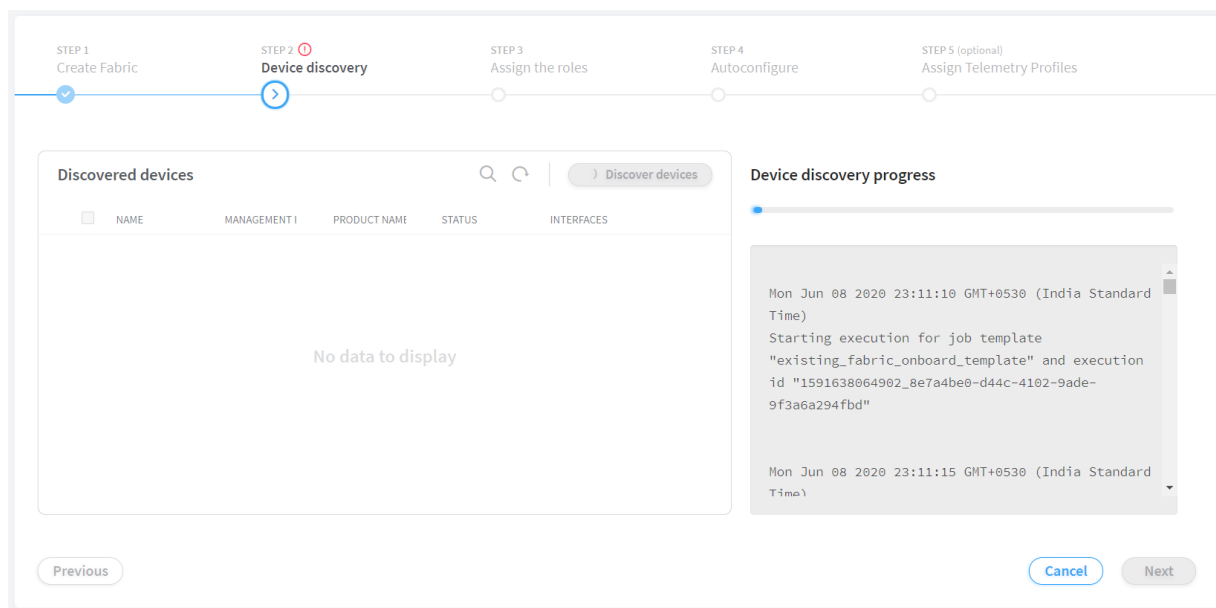
[Image Management | 54](#)

[Terminating Ongoing Fabric Jobs | 113](#)

Discover a Device

Device discovery is initiated as soon as you click **Next** on the Fabrics page. For more information on creating a new fabric, see "[Create a Fabric](#)" on page 7.

Figure 5: Device Discovery Page



- If you have followed the steps provided in the **Provisioning Option - New Fabric** (greenfield) section of the "[Create a Fabric](#)" on page 7 topic, clicking **Next** on the Fabrics page initiates the following fabric onboarding tasks:
 1. Based on the management subnet information that you provide, the DHCP configuration file (dnsmasq) is generated.
 2. After the devices are allotted IP addresses, the Dynamic Host Configuration Protocol (DHCP) lease file is generated with the device IP address and MAC address information.

Devices corresponding to the serial numbers listed in the Device Info section of the input YAML file are discovered.

The following base configuration is pushed to the discovered devices.

```
system {
  host-name "<serial-number>"
  root-authentication {
    encrypted-password "<encrypted-password>";
  }
  services {
    ssh {
      root-login allow;
    }
    telnet;
    netconf {
      ssh;
    }
  }
}
protocols {
  lldp {
    interface all;
  }
}
```

3. The devices are discovered and all configured interfaces available on the discovered devices are onboarded.
 4. The discovered devices obtain neighboring device information by using Link Layer Discovery Protocol (LLDP). Only devices that are part of the fabric are added.
 5. The node profiles available in the input YAML file are associated with multiple products and hardware. If the discovered device product name is associated with any listed product or hardware, the corresponding node profile is associated with that device.
 6. The DHCP IP is set as a static IP on the management interface.
 7. The input YAML supplemental configuration file is applied to the device.
- If you have followed the steps provided in the **Provisioning Option - Existing Fabric** (brownfield) section of the ["Create a Fabric" on page 7](#) topic, clicking **Next** on the Fabrics page initiates the following fabric onboarding tasks:

1. If you have entered a management subnet value, all reachable devices are discovered with a ping sweep.
If /32 is provided in the management subnet, only /32 hosts are discovered.
2. The devices are discovered and all configured interfaces available on the discovered devices are onboarded.
3. The discovered devices obtain neighboring device information by using Link Layer Discovery Protocol (LLDP). Only devices that are part of the fabric are added.
4. The node profiles available in the input YAML file are associated with multiple products and hardware. If the discovered device product name is associated with any listed product or hardware, the corresponding node profile is associated with that device.

The **Device discovery progress** bar on the Discovered devices page displays the progress of the device discovery job. See [Figure 6 on page 23](#).

Figure 6: Device Discovery Progress

Device discovery progress



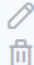
The devices that are discovered are listed in the Discovered Devices table and are in **Active** state. However, if a device image was upgraded during the initial zero-touch-provisioning, the device is in **Changed** state.

You can add a discovered device to the fabric by following these steps:

Select the device you want to add by selecting the check box next to the device name.

You can select more than one device.

Figure 7: Discovered Devices Table

	NAME	MANAGEMENT IP	PRODUCT NAME	STATUS	INTERFACES	
<input checked="" type="checkbox"/>	DC2-Spine2	10.XX.XX.XX	qfx10002-36q	ONBOARDED	41	...
<input type="checkbox"/>	DC2-Leaf3	10.XX.XX.XX	qfx5100-24q...	ONBOARDED	33	
<input type="checkbox"/>	DC2-Leaf2	10.XX.XX.XX	qfx5100-24q...	ONBOARDED	32	...
<input type="checkbox"/>	DC2-PNF	10.XX.XX.XX	srx5400	DISCOVERED	0	...
<input type="checkbox"/>	DC2-Leaf4	10.1.2.26	qfx5100-24q	ONBOARDED	33	

1 item selected | [Select all](#) | [Deselect all](#)

Click **Next** to assign roles.

Alternatively, you could proceed to assign roles by clicking **Next**, without selecting any device from the Discovered devices table. If you have not selected any device, all devices that are discovered will be added.

The Assign the roles page is displayed. For more information on assigning roles to devices, see ["Assign a Role to a Device" on page 25](#).

RELATED DOCUMENTATION

[Create a Fabric | 7](#)

[Assign a Role to a Device | 25](#)

[View Node Profile Information | 109](#)

[Delete a Fabric | 30](#)

[Image Management | 54](#)

Assign a Role to a Device

After you have completed the steps provided in the "[Discover a Device](#)" on [page 21](#) topic, you can assign roles to the devices from the Assign the Roles page.

Figure 8: Assign the Roles Page

The screenshot shows the 'Assign the Roles' page in a network management interface. The page is divided into five steps: STEP 1 Create Fabric, STEP 2 Device discovery, STEP 3 Assign the roles (current step), STEP 4 Autoconfigure, and STEP 5 (optional) Assign Telemetry Profiles. Below the progress bar is a table titled 'Assign to devices' with columns for NAME, MANAGEMENT IP, NODE PROFILE, ROLE, ROUTING ROLES, and AUTOCONFIGURE. The table lists several devices with their respective roles and routing roles. At the bottom, there are 'Previous', 'Cancel', and 'Autoconfigure' buttons.

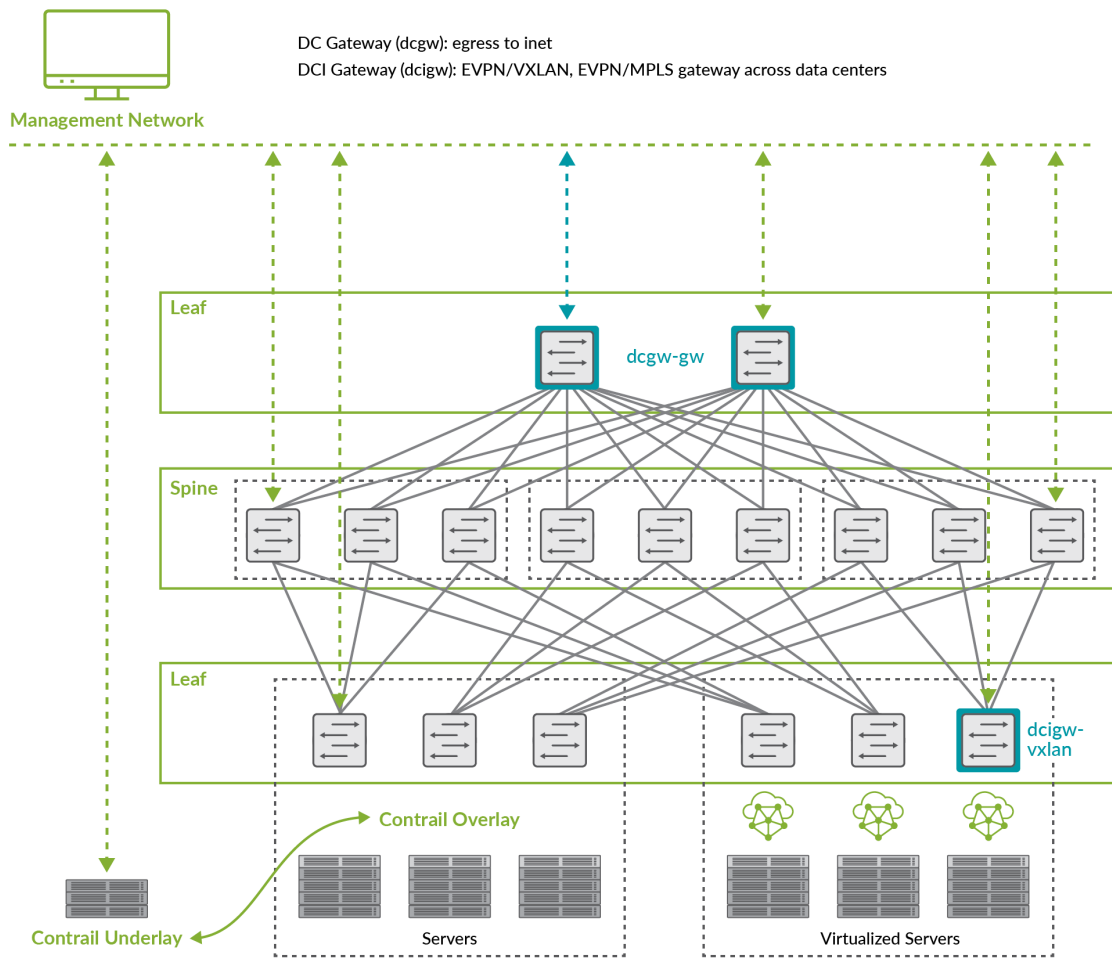
NAME	MANAGEMENT IP	NODE PROFILE	ROLE	ROUTING ROLES	AUTOCONFIGURE
DC2-Spine2	10.XX.XX.XX	juniper-qfx10k	-	-	False
DC2-Leaf3	10.XX.XX.XX	juniper-qfx5k-lean	-	-	False
DC2-Leaf2	10.XX.XX.XX	juniper-qfx5k-lean	-	-	False
DC2-PNF	10.XX.XX.XX	juniper-srx	-	-	False
DC2-Leaf4	10.XX.XX.XX	juniper-qfx5k-lean	-	-	False
DC2-Leaf1	10.XX.XX.XX	juniper-qfx5k-lean	-	-	False

Contrail Networking uses tags to identify functions that various devices in a DC fabric can provide. Contrail Networking uses the following roles to tag devices:

Physical roles A physical role determines whether a device can act as a leaf, spine, or physical network function (PNF).

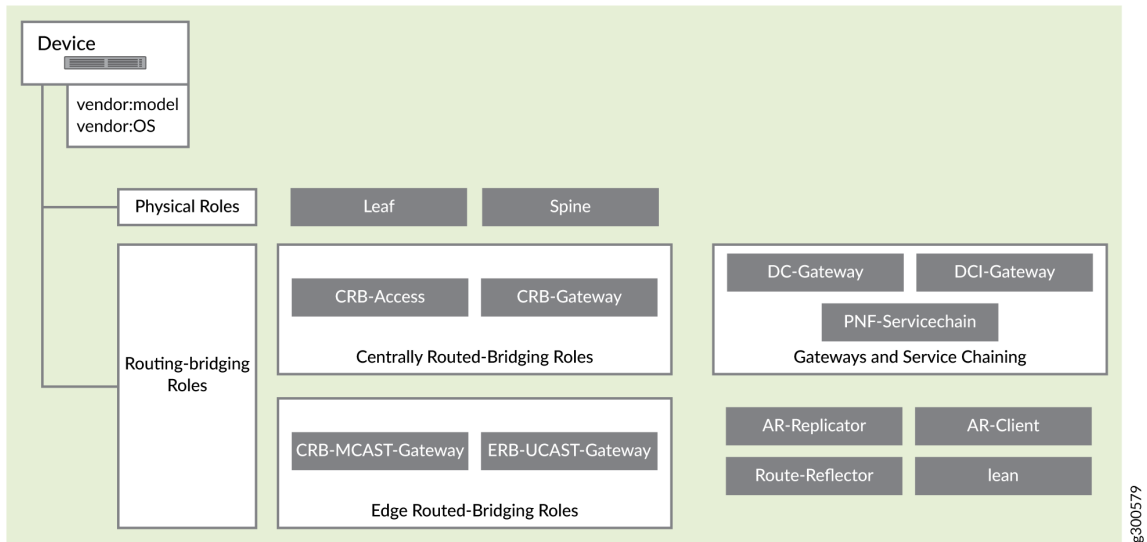
Routing-bridging roles These roles are specific to a set of capabilities that a device can deliver in a data center fabric with EVPN-VXLAN.

Figure 9: Sample Data Center Topology



For more information on supported hardware platforms, associated node profiles and roles, see ["Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles"](#) on page 185.

Figure 10: Supported Physical Roles and Routing-Bridging Roles



Steps to Assign Roles to Devices

You can assign roles to a device, or assign roles to multiple devices of the same node profile.

1. Follow these steps to assign a role to a device.
 - a. From the Assign to devices table, select the device you want to assign a role to by selecting the check box next to the device name.
 - b. Click the **Assign** icon at the end of the row to assign roles. See [Figure 11 on page 27](#).

Figure 11: Assign a Role to a Device

Assign to devices						
NAME	MANAGEMENT IP	NODE PROFILE	ROLE	ROUTING ROLES	AUTOCONFIGURE	
<input type="checkbox"/> DC2-Spine2	10.XX.XX.XX	juniper-qfx10k	-	-	False	
<input checked="" type="checkbox"/> DC2-Leaf3	10.XX.XX.XX	juniper-qfx5k-lean	-	-	False	Assign Role

Follow these steps to assign roles to multiple devices of the same node profile.

- a. From the assign to devices table, select the devices you want to assign a role to by selecting the check box next to the device name.
- b. Click **Assign roles to devices of same node profile** as shown in [Figure 12 on page 28](#).

Figure 12: Assign Role to Multiple Devices

<input type="checkbox"/>	NAME	MANAGEMENT IP	NODE PROFILE	ROLE	ROUTING ROLES	AUTOCONFIGURE	
<input type="checkbox"/>	DC2-Spine2	10.XX.XX.XX	juniper-qfx10k	-	-	False	...
<input checked="" type="checkbox"/>	DC2-Leaf3	10.XX.XX.XX	juniper-qfx5k-lean	-	-	False	...
<input checked="" type="checkbox"/>	DC2-Leaf2	10.XX.XX.XX	juniper-qfx5k-lean	-	-	False	...

The Assign role to devices pop-up is displayed.

2. You can now assign physical roles and routing-bridging roles.

- a. Select a physical role from the Physical Role list.

For example, select **spine** role.

- b. Select a routing-bridging role from the Routing Bridging Roles list.

For example, select **CRB-Gateway** role. You can select more than one routing-bridging role.

For more information on supported hardware platforms, associated node profiles and roles, see ["Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles"](#) on page 185.

- c. Click **Assign** to confirm selection.

3. Click **Autoconfigure** to initiate the auto-configuration job.

The Autoconfigure page is displayed.

The **Autoconfigure progress** bar on the Discovered devices page displays the progress of the auto-configuration job. Once the auto-configuration job is completed, click **Next**. The Assign Telemetry Profiles page is displayed.

For more information on assigning telemetry profiles, see ["Assign Telemetry Profiles"](#) on page 29.

RELATED DOCUMENTATION

[Create a Fabric | 7](#)

[Discover a Device | 21](#)

[View Node Profile Information | 109](#)

[Delete a Fabric | 30](#)

[Image Management | 54](#)

[Assign Telemetry Profiles | 29](#)

[Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles | 185](#)

Assign Telemetry Profiles

After you have successfully assigned roles to devices and run the auto-configuration job, you can assign telemetry profiles. This is an optional step when you perform the device onboarding task. You can also assign telemetry profiles after device onboarding. For more information on assigning telemetry profiles after device onboarding, see [Contrail Insights Flows in Contrail Command](#).

You can assign a telemetry profile to a device. Each telemetry profile is linked to an sFlow profile and you can have only one sFlow profile per telemetry profile. You can either link a telemetry profile to an existing sFlow profile or create a new sFlow profile while creating the telemetry profile. Three default sFlow profiles and telemetry profiles are predefined in the system when you bring up the cluster. You cannot edit or delete these default profiles.


The default telemetry profiles are

- sflow-access-interfaces—Indicates that sFlow is enabled on all the access interfaces on the device.
- sflow-fabric-interfaces—Indicates that sFlow is enabled on all the fabric interfaces.
- sflow-all-interfaces—Indicates that sFlow is enabled on all the interfaces on the device that has an sFlow profile attached to it.

You can assign a telemetry profile to more than one device. Follow these steps to assign a telemetry profile to device(s).

1. To assign a telemetry profile to a device,
 - a. Select the device you want to assign a telemetry profile to by selecting the check box next to the device name.
 - b. Click the **Assign** icon at the end of the row to assign telemetry profile. See [Figure 13 on page 29](#).

Figure 13: Assign a Telemetry Profile to a Device

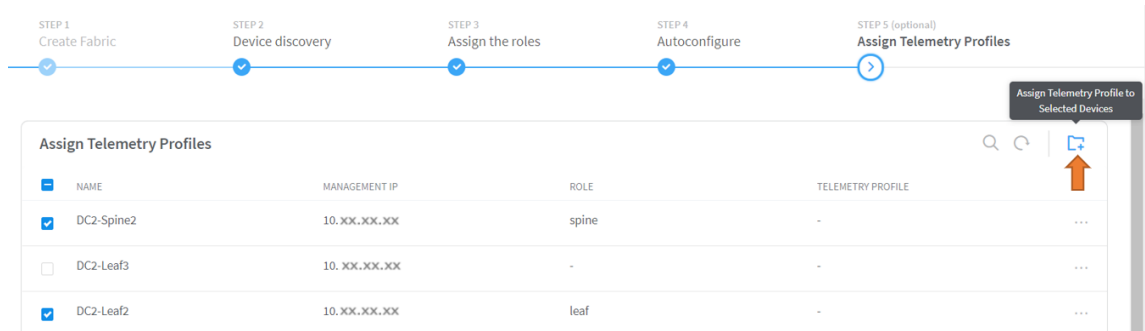


To assign a telemetry profile to more than one device,

- a. Select the devices you want to assign a telemetry profile to by selecting the check box next to the device name.

- b. Click Assign Telemetry Profile to Selected Devices as show in [Figure 14 on page 30](#).

Figure 14: Assign Telemetry Profile to Multiple Devices



The Assign Telemetry Profile to Device pop-up is displayed.

2. Select the telemetry profile from the Telemetry Profile list.
3. Click **Ok** to confirm selection.
4. Click **Finish** to exit the Create Fabric wizard.

The onboarding job is now complete.

RELATED DOCUMENTATION

[Create a Fabric | 7](#)

[Discover a Device | 21](#)

[View Node Profile Information | 109](#)

[Delete a Fabric | 30](#)

[Image Management | 54](#)

[Assign a Role to a Device | 25](#)

Delete a Fabric

You can delete a fabric by using the Contrail Command UI.

Follow these steps to delete a fabric:

1. Click **Fabrics**.

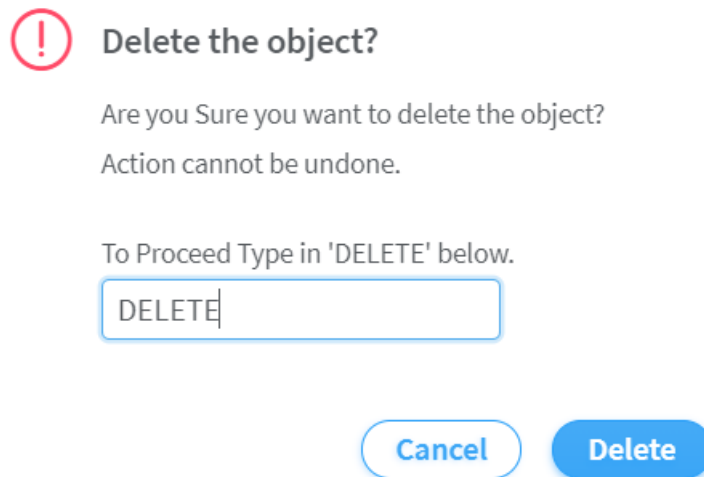
The Fabrics page is displayed.

2. Select the fabric you want removed by selecting the check box next to the name of fabric.

NOTE: Contrail Networking Release 5.0.1 does not support bulk deletion of fabric.

3. Click the **Delete** icon at the end of the row.
The **Delete the object?** pop-up is displayed.
4. Enter **DELETE** in the field as shown in [Figure 15 on page 31](#).

Figure 15: Delete Confirmation



The image shows a confirmation dialog box with a red warning icon. The title is "Delete the object?". Below the title, it asks "Are you Sure you want to delete the object?" and states "Action cannot be undone." There is a text input field with the word "DELETE" entered. Below the input field are two buttons: "Cancel" and "Delete".

! Delete the object?

Are you Sure you want to delete the object?
Action cannot be undone.

To Proceed Type in 'DELETE' below.

Cancel Delete

Finally, click **Delete** to delete the fabric.

RELATED DOCUMENTATION

[Create a Fabric | 7](#)

[Discover a Device | 21](#)

[Assign a Role to a Device | 25](#)

[View Node Profile Information | 109](#)

[Image Management | 54](#)

Provision Fabric Devices Using End-to-End ZTP

From Contrail Networking Release 5.1, you can provision fabric devices using Zero Touch Provisioning (ZTP).

ZTP allows you to provision new Juniper Networks devices in your network automatically, with minimal manual intervention.

This topic provides steps to provision fabric devices using ZTP and configure underlay network via Contrail Command UI.

NOTE: You must complete [Installing Contrail Command](#) before proceeding.


An All-in-One (AIO) Contrail cluster is a single Contrail Networking server with multiple VMs that supply controller, orchestrator, and, compute nodes.

AIO cluster node responds to DHCP requests to zeroize devices. Hence, AIO cluster node must be in the same subnet as of the management subnet.

NOTE: The minimum required version of Junos OS for QFX5000 and QFX10000 Series devices is 18.1R3-S5 or higher. The minimum required version of Junos OS for the MX devices is 18.4R2-S3. Also, all MX Series Routers must be zeroized for ZTP to work.

The following example illustrates the process for provisioning fabric devices using end-to-end ZTP.

Sample Topology



Prerequisites

- Contrail Command server:
 - Install CentOS 7.6.
- AIO Contrail cluster node:
 - Install CentOS 7.6.
 - Configure *eno1* port with the static IP **10.87.x.1/27**.

```

HWADDR=0c:xx:xx:xx:xx:4a
NM_CONTROLLED=no
BOOTPROTO=none
DEVICE=eno1
ONBOOT=yes
IPADDR=10.87.x.1
NETMASK=255.255.255.224
GATEWAY=10.87.6.30

```

- Configure *ens2f1* port with the static IP **10.1.x.2/24**.

```
HWADDR=90:xx:xx:xx:xx:a1
NM_CONTROLLED=no
BOOTPROTO=none
DEVICE=ens2f1
ONBOOT=yes
IPADDR=10.1.x.2
NETMASK=255.255.255.0
GATEWAY=10.1.x.1
```

command_servers.yml example file:

```
---
command_servers:
  server1:
    ip: 10.87.x.2
    connection: ssh
    ssh_user: root
    ssh_pass: password
    sudo_pass: password
    ntpserver: x.x.x

    # Specify either container_path
    # or registry details and container_name
    container_registry: x.x.x:5010
    container_name: contrail-command
    container_tag: master-720
    config_dir: /etc/contrail

    # contrail command container configurations given here go to /etc/contrail/contrail.yml
    contrail_config:
      # Database configuration. MySQL/PostgreSQL supported
      database:
        # MySQL example
        host: localhost
        user: root
        password: password
        name: contrail_test
        type: postgres
        dialect: postgres
```

```
# Max Open Connections for DB Server
max_open_conn: 100
connection_retries: 10
retry_period: 3s

# Log Level
log_level: debug

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

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

# Enable GRPC or not
enable_grpc: false

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

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

notify_etcd: false

# Keystone configuration
```

```
keystone:
  local: true
  assignment:
    type: static
  data:
    domains:
      default: &default
        id: default
        name: default
    projects:
      admin: &admin
        id: admin
        name: admin
        domain: *default
      demo: &demo
        id: demo
        name: demo
        domain: *default
    users:
      admin:
        id: admin
        name: Admin
        domain: *default
        password: password
        email: admin@x.com
        roles:
          - id: admin
            name: Admin
            project: *admin
      bob:
        id: bob
        name: Bob
        domain: *default
        password: bob_password
        email: bob@x.com
        roles:
          - id: Member
            name: Member
            project: *demo
  store:
    type: memory
    expire: 3600
  insecure: true
```

```
    authurl: https://localhost:9091/keystone/v3

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

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

watcher:
  enabled: false
  storage: json

client:
  id: admin
  password: password
  project_id: admin
  domain_id: default
  schema_root: /
  endpoint: https://localhost:9091

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

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



```
agent:
  enabled: true
  backend: file
  watcher: polling
  log_level: debug
cache:
  enabled: true
  timeout: 10s
  # how long revision deleted event preserved.
  max_history: 100000
rdbms:
  enabled: true
```

To provision fabric devices using ZTP via Contrail Command UI:

1. Log in to Contrail Command UI as a *super user* using root user credentials.



Log in

Select Cluster

Username

root

Password

Log in

2. Install bootstrap server.

Click **Servers**.

- a. Click **Create**.
- b. Enter the required details.
- c. Click **Create**.

Choose Mode* Express Detailed Bulk Import (csv)

Select workload type this server will be used for Physical/Virtual Node Baremetal

Hostname* Management IP* Management Interface

Credentials

MAC Address

Disk Partition(s)

Network Interfaces

Name*	IP Address*	
<input type="text" value="eno1"/>	<input type="text" value="x.x.x.x"/>	▼ ▲ 🗑️
Name*	IP Address*	
<input type="text" value="ens2f1"/>	<input type="text" value="x.x.x.x"/>	▼ ▲ 🗑️

[+ Add](#)

- Port *eno1* is connected to management VLAN.
 - Port *ens2f1* is connected to QFX ToR.
3. Create cluster by entering the required details.
- a. Click **Cluster**.
 - b. Click **Add Cluster**.
 - c. Enter the required details including **Inventory, Cloud Manager, Infrastructure Networks, Overcloud**, etc.
Check **Enable ZTP** checkbox.
 - **Default Vrouter Gateway** is the QFX ToR IRB IP. The IP is used for provisioning the network.

CONTROLLER_NODES and *CONTROL_NODES* are a part of Contrail Networking Configuration.

 - *CONTROLLER_NODES* IP is a static IP configured on port *eno1*.
 - *CONTROL_NODES* IP is a static IP configured on port *ens2f1*.

STEP 1 Inventory

STEP 2 **Cloud Manager**

STEP 3 Infrastructure Networks

STEP 4 (optional) Overcloud

STEP 5 (optional) Undercloud Nodes

STEP 6 (optional) Jumphost Nodes

STEP 7 Control Nodes

STEP 8 Orchestrator Nodes

STEP 9 (optional) Compute Nodes

STEP 10 (optional) Contrail Service Nodes

STEP 11 (optional) Appformix Nodes

Choose Provisioning Manager*

RHOSP Manager Contrail Cloud Manager

Cluster Name*

Container Registry* opencontrainightly Insecure Container Registry Username*

Container Registry Password* Contrail Version* latest

Provisioner Type Ansible

Domain Suffix local NTP Server Default Vrouter Gateway

Encapsulation Priority MPLSoGRE,MPLSoUDP,...

Enable ZTP ⓘ

▶ Contrail Configuration

High availability mode

Available servers

Search servers Add all <

HOSTNAME	IP ADDRESS	DISK PARTITION
Add servers to your inventory		

Assigned Control nodes

Search servers Remove all

HOSTNAME	IP ADDRESS	DISK PARTITION
5a12s1-node1		
Roles*		
contrail_config_node ×		
contrail_config_database_node ×		
contrail_analytics_node ×		
contrail_analytics_alarm_node ×		
contrail_analytics_snmp_node ×		
contrail_analytics_database_node ×		

NOTE: Set **enable_swift** to **yes** if the cluster will be used for any image management tasks on the fabric devices. Otherwise, set **enable_swift** to **no**.

- **enable_ironic** is used for life cycle management of Bare Metal Servers (BMS).

- **enable_swift** is used to provision Swift containers (object storage). All the images used during different fabric related tasks are stored in these containers.
- **enable_haproxy** is used when OpenStack controllers are set up in high availability (HA) mode.

Orchestrator type*
 Show Advanced

Container Registry

Openstack Release

Control & Data Network Virtual IP address Management Network Virtual IP address

Customize configuration ⓘ

Kolla Globals

Key	Value	
<input type="text" value="enable_ironic"/>	<input type="text" value="no"/>	▼ ▲ 🗑
<input type="text" value="enable_swift"/>	<input type="text" value="no"/>	▼ ▲ 🗑
<input type="text" value="enable_haproxy"/>	<input type="text" value="no"/>	▼ ▲ 🗑

+ Add

Kolla Passwords
+ Add

Available servers

Add all <

HOSTNAME	IP ADDRESS	DISK PARTITION
Add servers to your inventory		

Previous

Assigned Openstack nodes

Remove all

HOSTNAME	IP ADDRESS	DISK PARTITION
▼ 5a12s1-node1		🗑

Roles*

Next

Available servers

Add all <

HOSTNAME	IP ADDRESS	DISK PARTITION
Add servers to your inventory		

Assigned Service nodes

Remove all

HOSTNAME	IP ADDRESS	DISK PARTITION
▼ 5a12s1-node1		🗑

Default Vrouter Gateway*

Cluster overview

Display name	AIO
Container registry	-repo:5010 (insecure)
Contrail version	master-550
Provisioner type	ansible
Domain Suffix	local
NTP server	ntp.juniper.net
Default Vrouter Gateway	
Encapsulation priority	VXLAN,MPLSoUDP,MPLSoGRE
Enable ZTP	true
▶ Contrail configuration	
High availability mode	false
Orchestrator	openstack
Openstack release	queens
Openstack internal virtual IP	-
Openstack external virtual IP	-
Openstack registry	default
▶ Kolla globals	
Kolla passwords	-

Nodes overview



All cluster nodes	Control nodes	Compute nodes	Openstack nodes	Service nodes
NAME	TYPE	IP ADDRESS	NODE PROFILE	ROLES
5c10s7-node4	physical/virtual node			Control node; Compute n...

Previous

Provision

d. Click **Create**.

4. After creating the cluster, log in to the cluster using root user credentials.



Log in

Select Cluster

AIO-98fe2e93-4a0c-4602-b735-85e2697a2d92

Username

root

Password

Log in

5. Run fabric ZTP workflow to onboard the fabric devices

- a. Click **Fabrics**.
- b. Click **Create**.
- c. Click **New Fabric**.
- d. Click **Provision**.
- e. Enter the required details.

Table 4: Required Fields for creating Fabric

Field	Details
Overlay ASN (iBGP)	iBGP ASN pool for Contrail Networking overlay network. List of the ASN pools that can be used to configure the iBGP peers for the IP fabric
Underlay ASNs (eBGP)	eBGP ASN pool for fabric underlay network. List of the ASN pools that can be used to configure the eBGP peers for the IP fabric
Management subnet	List of the management network subnets for the fabric

Table 4: Required Fields for creating Fabric (Continued)

Field	Details
Fabric subnet	List of subnet prefixes that can be used for the P2P networks between fabric devices
Loopback subnet	List of the subnet prefixes that can be allocated to fabric device loopback IPs

Sample device_info.yml file

```

supplemental_day_0_cfg:
  - name: "cfg1"
    cfg: |
      set system ntp server 167.99.20.98
device_to_ztp:
  - serial_number: "DK588"
    supplemental_day_0_cfg: "cfg1"
    hostname: '5a12-qfx5'
  - serial_number: "VF3717350117"
    hostname: '5a12-qfx9'
  - serial_number: "11675330144"
  - serial_number: "74656088411"

```

NOTE: The YAML file lists the devices used for ZTP during a greenfield onboarding of devices. Contrail Networking Release 1907 introduces the ability to configure hostnames to the devices being onboarded. If the hostnames attribute is not specified, the device serial number is used as the hostname by default.

STEP 1 **Create Fabric** STEP 2 Device discovery STEP 3 Assign the roles STEP 4 Autoconfigure

Name*

Device credentials*

Overlay ASN (iBGP)*

Device Info*
 Download Template: [Ⓞ \(*.yaml\)](#)
[Ⓞ Upload .yaml or .yml](#)

Node profiles*

Management subnets

CIDR* Gateway*

f. Assign the roles to the fabric devices.

- DK588 as Spine with CRB-Gateway and Route-Reflector roles.
- WS3XXXX0049 as Leaf with CRB-Access role.

STEP 1 Create Fabric STEP 2 Device discovery STEP 3 **Assign the roles** STEP 4 Autoconfigure

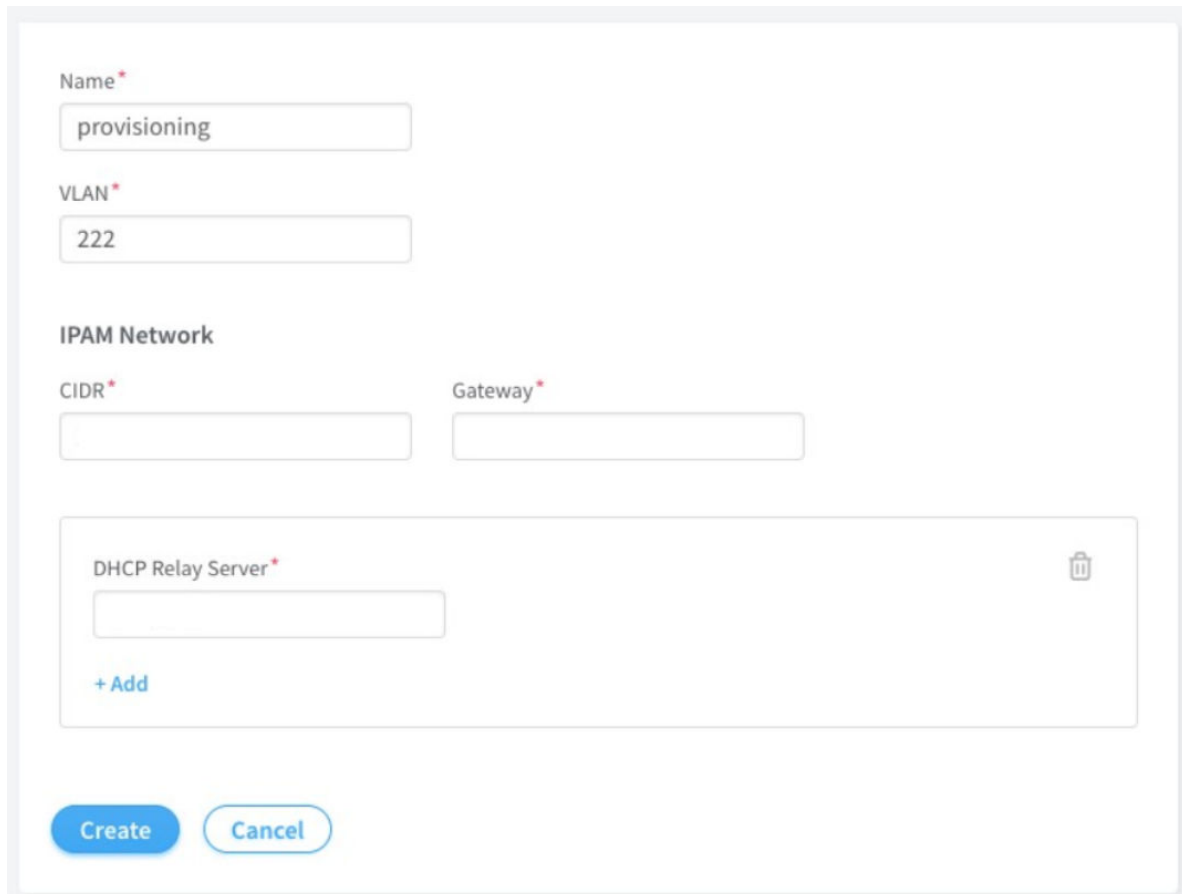
Assign to devices 🔍 ↻

<input type="checkbox"/>	NAME	MANAGEMENT IP	ROLE	ROUTING ROLES	AUTOCONFIGURE	
<input type="checkbox"/>	DK588		spine	CRB-GatewayRoute-Reflector	False	
<input type="checkbox"/>	WS3 70049		leaf	CRB-Access	False	...

No items selected

To configure underlay network via Contrail Command UI:

1. Create provisioning infrastructure network.
 - a. Click **Networks**.
 - b. Create a network by entering the required details.



The screenshot shows a web form for creating a network. The form is titled "Name*" and has a text input field containing "provisioning". Below this is a "VLAN*" field with the value "222". The "IPAM Network" section contains two fields: "CIDR*" and "Gateway*", both of which are empty. Below these is a "DHCP Relay Server*" section with a text input field and a trash icon. A "+ Add" link is located below the input field. At the bottom of the form are two buttons: "Create" (a solid blue button) and "Cancel" (a white button with a blue border).

2. Import server topology.
 - a. Click **Servers**.
 - b. Click **Import**.
 - c. Upload the **server topology** file.

Import Server

To import a Server, please upload a file (*.json or *.yaml) from your computer

Download Template: [⬇️ \(*.json\)](#) [⬇️ \(*.yaml\)](#)

Drag a file here, or [browse](#)

📎 server_01.yaml

Cancel

Import

Sample server topology yaml file:

```
nodes:
  - name: 5a12s1
    type: baremetal
    ports:
      - name: ens2f1
        mac_address: 90:xx:xx:xx:xx:a1
        switch_name: WS37XXX049
        port_name: xe-0/0/46
        switch_id: 3c:61:04:63:0e:80
```

Table 5: Required Fields for server topology yaml file

Field	Details
name	Name of the infrastructure BMS node
type	Type of the infrastructure BMS node. It must be "baremetal"

Table 5: Required Fields for server topology yaml file (Continued)

Field	Details
ports	List of the ports of BMS node connected to the TOR switch
name	Name of the BMS port
switch_name	TOR switch name
port_name	TOR port name

3. Import server node profile.

You must create server node profile for the Contrail Networking Controller server.

- a. Click **Servers**.
- b. Click **Node Profiles**.
- c. Click **Import**.
- d. Upload the **server node profile** file.

Table 6: Required fields for Server Node Profile

Field	Details
kind	Resource type
name	Name of a resource
fq_name	Fully Qualified name of a resource
parent_type	Node profile parent resource type. It must be "global-system-config"
node_profile_vendor	Node Profile vendor name

Table 6: Required fields for Server Node Profile (Continued)

Field	Details
node_profile_type	Node profile type. It must be "end-system" for servers
hardware_refs	List of references to the hardware models supported by the node profile
card_refs	List of references to the interface cards

Sample server node profile json file:

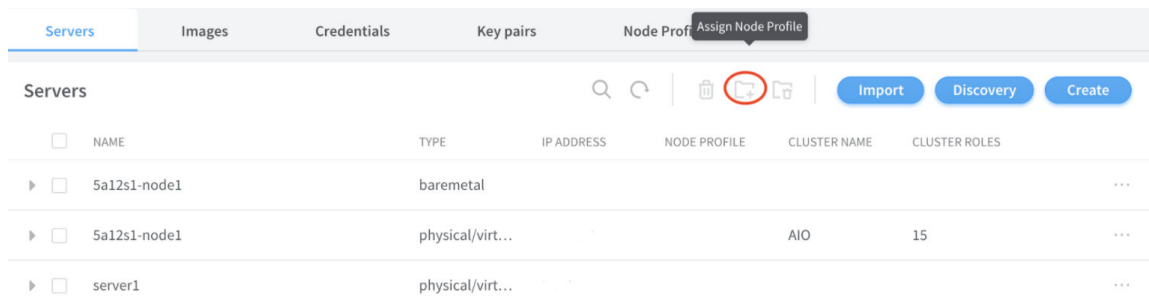
```
{
  "resources": [
    {
      "kind": "card",
      "data": {
        "name": "dell-bms-card",
        "fq_name": ["dell-bms", "dell-bms-card"],
        "interface_map": {
          "port_info": [{"name": "ens2f1", "labels": ["provisioning"]}]}
      }
    },
    {
      "kind": "hardware",
      "data": {
        "name": "dell-bms",
        "fq_name": ["dell-bms"],
        "card_refs": [{"to": ["dell-bms", "dell-bms-card"]}]}
    },
    {
      "kind": "node_profile",
      "data": {
        "hardware_refs": [{"to": ["dell-bms"]}]}],
        "parent_type": "global-system-config",
        "name": "Dell_BMS_01",
        "fq_name": ["default-global-system-config", "Dell_BMS_01"],
        "node_profile_vendor": "Dell",
      }
    }
  ]
}
```

```

    "node_profile_type": "end-system"
  }
}
]
}

```

4. Assign node profile to the server.
 - a. Click **Servers**.
 - b. Select the required server from the list.
 - c. Click **Assign Node Profile**.



Once the above procedure is completed, change the default route from *management* port to the *access* port.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2003	The minimum required version of Junos OS for the MX devices is 18.4R2-S3. Also, all MX Series Routers must be zeroized for ZTP to work.
1907	Contrail Networking Release 1907 introduces the ability to configure hostnames to the devices being onboarded

RELATED DOCUMENTATION

[Installing Contrail Command](#)

[Terminating Ongoing Fabric Jobs](#) | 113

3

CHAPTER

Fabric Configuration

- Image Management | 54
- Onboard Brownfield Devices | 57
- Onboard Greenfield Devices | 67
- Device Import | 78
- Create Virtual Network | 82
- Create Logical Routers | 90
- Create Network Policy | 92
- Create Network IPAM | 94
- Reconfigure Roles | 96
- Managing Custom Roles | 99
- View Node Profile Information | 109
- Monitoring Fabric Jobs | 110
- Terminating Ongoing Fabric Jobs | 113
- Adding a Leaf or Spine Device to an Existing Fabric Using ZTP | 115
- Grouping Fabric Devices and Roles Using Device Functional Groups | 118
- Creating Layer 3 PNF Service Chains for Inter-LR Traffic | 121
- Creating VNF Service Chains for Inter-LR Traffic | 128
- Retaining the AS Path Attribute in a Service Chain | 164
- Assisted Replication of Broadcast, Unknown Unicast, and Multicast Traffic | 165
- Running Generic Device Operations Commands In Contrail Command | 168

Adding DHCP Server Information for Virtual Networks and Logical Routers | 173

Return Material Authorization | 179

Approaches to Enable External Connectivity for Overlay Networks | 183

Contrail Networking Supported Hardware Platforms and Associated Roles And
Node Profiles | 185

Image Management

IN THIS SECTION

- [Upload a New Device Image | 54](#)

This topic provides instructions to upload a new device image to the Contrail Networking fabric.

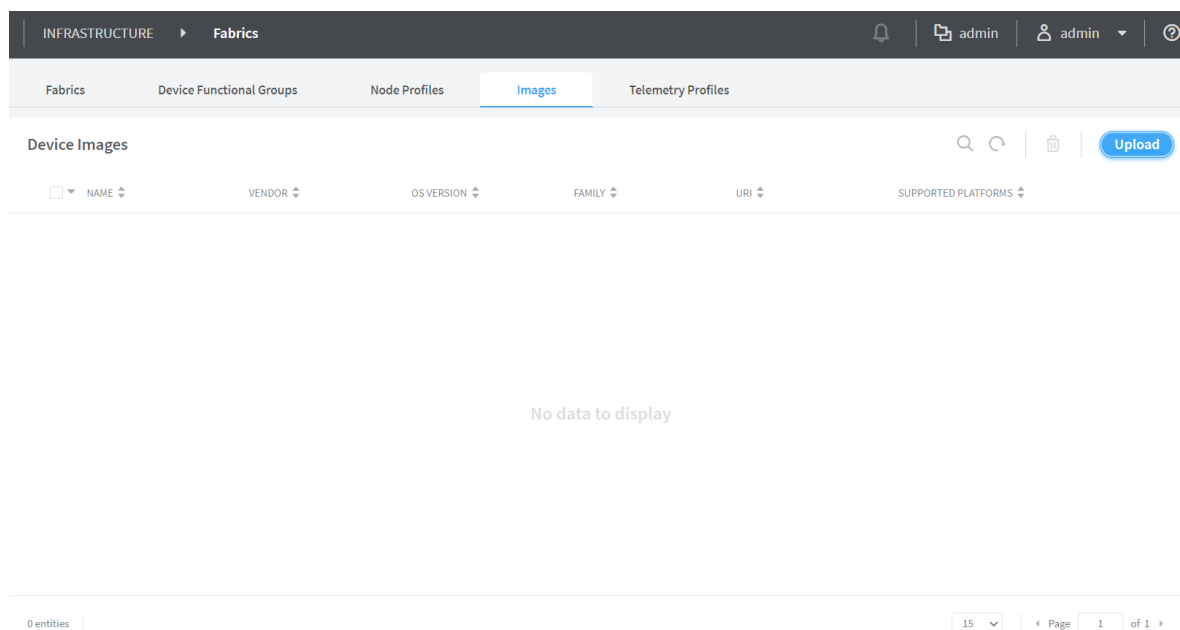
Upload a New Device Image

Follow these steps to upload a new device image:

1. Click **Infrastructure>Fabrics>Images**.

The Device Images page is displayed. See [Figure 16 on page 54](#).

Figure 16: Device Images



2. Click **Upload**.

The Upload Image pop-up is displayed. See [Figure 17 on page 55](#).

Figure 17: Upload Image

Upload Image

Device Image Tags Permissions

Name *

Pick a File *

Drag file here or [browse](#)

Vendor Name *

Device Family *

Supported Platforms *

Os Version *

Image MD5

Image SHA1

3. Enter the following information given in [Table 7 on page 55](#).

Table 7: Upload Image Details

Field	Action
Name	Enter a name for the device image.

Table 7: Upload Image Details (*Continued*)

Field	Action
Pick a file	Click Upload and navigate to the local directory and select the device image file. Click Open to confirm selection.
Vendor Name	Enter name of the vendor.
Device Family	Select the device family from the list.
Supported platforms	Select the hardware platforms that are compatible with the image file, from the list.
OS version	Enter the OS version.
Image MD5	(Optional) Enter MD5 checksum value.
Image SHA1	(Optional) Enter SHA1 checksum value.

NOTE: The images that you upload can not have the same vendor name, device family, supported platforms, or OS version. The Contrail Command UI will not allow you to upload two image files with the same field information.

4. Click **Upload** to begin uploading the device image file.

You are redirected to the Device Images page. When the image upload is complete, the device image is listed in Device Images page.

RELATED DOCUMENTATION

[Create a Fabric | 7](#)

[Discover a Device | 21](#)

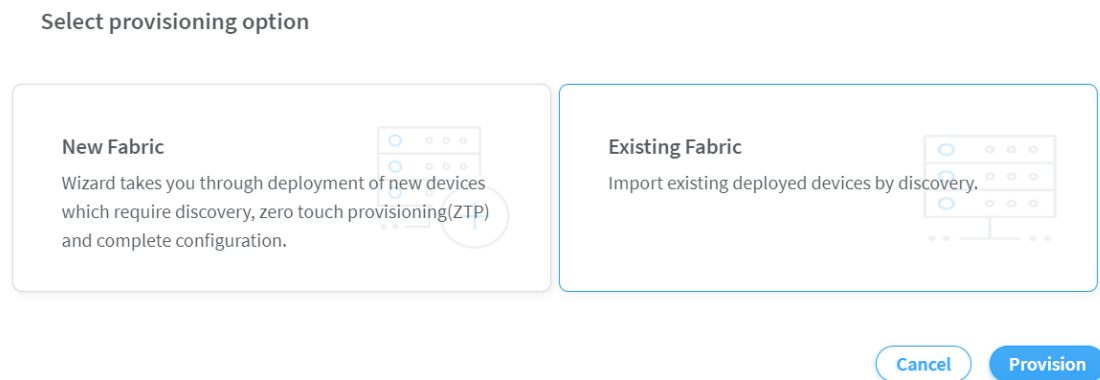
[Assign a Role to a Device | 25](#)

Onboard Brownfield Devices

Follow these steps to onboard brownfield devices from the Contrail Command user interface (UI):

1. Click **Infrastructure>Fabrics**.
The Fabrics page is displayed.
2. Click **Create**.
You are prompted to select a provisioning option.
3. Click **Existing Fabric** to import existing (brownfield) devices by discovery. See [Figure 18 on page 57](#).

Figure 18: Select Existing Fabric



4. Click **Provision**.
The Create Fabric page is displayed.
5. Enter the information as given in [Table 8 on page 57](#).

Table 8: Provision Existing Fabric

Field	Action
Name	Enter a name for the fabric.

Table 8: Provision Existing Fabric *(Continued)*

Field	Action
Overlay ASN (iBGP)	<p>Enter autonomous system (AS) number in the range of 1-65,535.</p> <p>If you enable 4 Byte ASN in Global Config, you can enter 4-byte AS number in the range of 1-4,294,967,295.</p>
Node profiles	<p>Add node profiles.</p> <p>You can add more than one node profile.</p> <p>All preloaded node profiles are added to the fabric by default. You can remove a node profile by clicking X on the node profile. For more information, see "View Node Profile Information" on page 109.</p> <p>For more information on supported hardware platforms, associated node profiles and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185.</p>
Disable VLAN-VN Uniqueness Check	<p>Select this check box when you are using the enterprise style of configuration but want to disable the requirement that every VLAN ID must have a 1:1 mapping with a VNI. Enterprise style of configuration is enabled by selecting the VLAN-ID Fabric-Wide Significance check box.</p>

Table 8: Provision Existing Fabric (Continued)

Field	Action
VLAN-ID Fabric Wide Significance	<p>Select the check box to enable enterprise style of configuration for the CRB-Access role on QFX devices. De-select the check box to enable service provider style of configuration for the CRB-Access role. The check box is selected by default since enterprise style is the default setting.</p> <p>Once configured you can modify the enterprise style setting to service provider style of configuration. However, you cannot modify the service provider style to enterprise style of configuration without having to recreate the fabric.</p> <p>The service provider style of configuration allows for customization of Ethernet-based services at the logical interface level. Each logical interface is bound to a unique VLAN ID. With the enterprise style of configuration, logical interfaces are placed into Layer 2 mode by specifying ethernet-switching as the interface family. The ethernet-switching family can be configured only on a single logical unit, unit 0. For more information on enterprise and service provider type of configurations, see Flexible Ethernet Services Encapsulation.</p> <p>NOTE: Contrail Networking Release 1909 supports QFX10002-60C device running Junos OS Release 19.1R2 and later. QFX10002-60C device works only if enterprise style of configuration is enabled. To enable enterprise style of configuration, select the VLAN-ID Fabric Wide Significance check box when onboarding the QFX10002-60C device. For more information on enterprise style of configuration, see "Configuring EVPN VXLAN Fabric with Multitenant Networking Services" on page 289.</p> <p>For more information on supported hardware platforms and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185.</p>

Table 8: Provision Existing Fabric (Continued)

Field	Action
Device credentials	Enter the device credentials to access the fabric devices for discovery. If your fabric devices have different username and password combinations for device access, click the + Add option to add additional username and password credentials.
Management subnets	<p>Enter the following information to auto-assign management IP addresses to devices:</p> <p>CIDR—Enter the block of IP addresses that will be assigned as management IP addresses. The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Gateway—Enter gateway address for the devices in the management subnet that connect to the fabric.</p>
Loopback subnets	<p>Enter loopback subnet (lo0) address.</p> <p>The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.</p>
Underlay ASNs (eBGP)	<p>Enter autonomous system (AS) number in the range of 1-65,535.</p> <p>If you enable 4 Byte ASN in Global Config, you can enter 4-byte AS number in the range of 1-4,294,967,295.</p> <ul style="list-style-type: none"> • Enter minimum value in ASN From field. • Enter maximum value in ASN To field.

Table 8: Provision Existing Fabric *(Continued)*

Field	Action
Fabric subnets	<p>Enter fabric CIDR address. The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.</p>
LR Loopback subnets	<p>Enter an IP subnet to be assigned as loopback interface (lo0) addresses used in Logical Routers (LR). The LR loopback interface IP address is required for eBGP peering to external or unmanaged devices.</p> <p>The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p>
Loopback subnets (CIDR)	<p>Enter loopback address.</p> <p>Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.</p> <p>If you assign the AR-Replicator and AR-Client roles to enable assisted replication on the QFX10000 devices in a datacenter, you must enter loopback address. For more information, see "Assign a Role to a Device" on page 25.</p>
PNF Servicechain subnets	<p>Enter the IP subnet for allocating IP addresses in the PNF Servicechain subnets field to establish EBGP session between PNF device and SPINE switch.</p> <p>This is an optional field that should be left blank when you are not creating service chains.</p>

Table 8: Provision Existing Fabric *(Continued)*

Field	Action
Advanced interface filters	<p>Create an interface filter to filter the interfaces to include in the fabric. By default, all interfaces identified as participating in Contrail are imported into the fabric during the fabric provisioning process. If an interface filter is set, the fabric provisioning process includes the interfaces that are participating in Contrail and that match the interface filter in the fabric.</p> <p>To create an interface filter, choose the operation as regex and enter the filter characters in the Expression field. The Expression field supports all characters - including metacharacters - allowed in Python regex filters. For example, you can enter <code>^xe</code> in the Expression field to filter out all 10Gbps xe interfaces from the fabric.</p>
Import configured interfaces	<p>Choose this option if configured interfaces need to be imported into the fabric in addition to runtime interfaces. With some exceptions, a configured interface is generally an interface that has been configured in the Junos OS software.</p> <p>A runtime interface is generally an interface that has not been configured in Junos OS. You can confirm which interfaces are configured interfaces by entering the <code>show interfaces</code> command at the configuration mode prompt(<code>#</code>) in Junos. Only runtime interfaces are imported into the fabric by default.</p>


6. Click **Next**.

The Device discovery page is displayed.

The **Device discovery progress** bar on the Device discovery page displays the progress of the device discovery job.

Figure 19: Device Discovery Progress Bar

Device discovery progress



The list of devices discovered are listed in the Discovered devices page.

7. Select the device you want to add by selecting the check box next to the device name.
You can select more than one device.
8. Click **Next** to assign roles.
Assign the Roles page is displayed.
9. From the assign to devices table, select the device you want to assign a role to by selecting the check box next to the device name.
Click the **Assign** icon at the end of the row to assign roles. The Assign role to devices pop-up is displayed.
10. You can now assign physical roles and routing-bridging roles.
 - a. Select a physical role from the Physical Role list.
 - b. Select a routing-bridging role from the Routing Bridging Roles list.

Assigning Roles for Spine Devices:

- Select **spine** from the Physical Role list.
- Select **CRB-Gateway** from the Routing Bridging Roles list.

Assigning Roles for Leaf Devices:

- Select **leaf** from the Physical Role list.
- Select **CRB-Access** from the Routing Bridging Roles list.

Assigning Roles for PNF Devices:

- Select **PNF** from the Physical Role list.
- Select **CRB-Access** and **PNF-Servicechain** from the Routing Bridging Roles list.

NOTE: The number of PNF instances you can create depends on the subnet mask of the pnf-servicechain-subnet that you provided during fabric onboarding. You can create multiple /29 subnets from the pnf-servicechain-subnet.

For example, if a /24 subnet is provided for the pnf-servicechain-subnet, then, you can create $2^5 = 32(29-24=5)$ subnets out of it. Each PNF uses a pair of /29 subnets. Thus, for a /24 subnet, you can have a maximum of 16 PNFs.

Assigning Roles for VNF Devices:

- Select **VNF** from the Physical Role list.
- Select **CRB-Access** from the Routing Bridging Roles list.

NOTE: **ERB-UCAST-Gateway** routing bridging role is also supported.

NOTE: When you configure a QFX series device as a data center gateway, ensure that you assign DC-Gateway role to the spine device.

To assign a DC-Gateway role to a spine device,

- Select **spine** from the Physical Role list.
- Select **DC-Gateway** from the Routing Bridging Role list.

Click **Assign** to confirm selection.

11. Click **Autoconfigure** to initiate the auto-configuration job.

The Autoconfigure page is displayed.

The **Autoconfigure progress** bar on the Discovered devices page displays the progress of the auto-configuration job. Once the auto-configuration job is completed, click **Next**. The Assign Telemetry Profiles page is displayed.

Starting with Contrail Networking Release 2008, you can apply MTU, admin state, flow control, LACP force up, interface type attributes to physical interfaces; and MTU to logical interfaces. These attributes are applied to physical and logical interfaces after you **Autoconfigure** the devices.

To apply these attributes to interfaces:

- Navigate to **Infrastructure > Fabric**.
- Select the desired fabric from the list.
- Select the desired fabric device from the list.
- Click **Physical Interfaces > Create**.
- Enter the required details.

Figure 20: Create Physical Interface

Create Physical Interface

Interface Permissions

Name*

Type

MTU ⓘ

Description

Admin State On

Flow Control ⓘ

- f. Click **Create**.
- g. Click **Logical Interfaces > Create**.
- h. Enter the required details.

Create Logical Interface

Interface
Permissions

Name* ?

Connected Physical Interface*

MTU ?

Description

Cancel
Create

i. Click **Create**.

- 12.** (Optional) Assign telemetry profiles. For more information, see "[Assign Telemetry Profiles](#)" on page 29.

PNF service chain and VNF service chain does not use telemetry profiles.

- 13.** Click **Finish** to exit the Create Fabric wizard.

The onboarding job is now complete.

NOTE: After the devices are onboarded, if you edit the fabric topology by adding new spine or leaf devices or by adding new links between devices, you *must* onboard the edited devices again. If you do not onboard the devices after edits to the initial configuration, underlay formation for the edited devices fails. You can choose to onboard individual devices by clicking the **Onboard**

button for the selected device in the **Fabric Devices** tab of the **Infrastructure > Fabrics > Fabric_Name** page.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2008	Starting with Contrail Networking Release 2008, you can apply MTU, admin state, flow control, LACP force up, interface type attributes to physical interfaces; and MTU to logical interfaces.

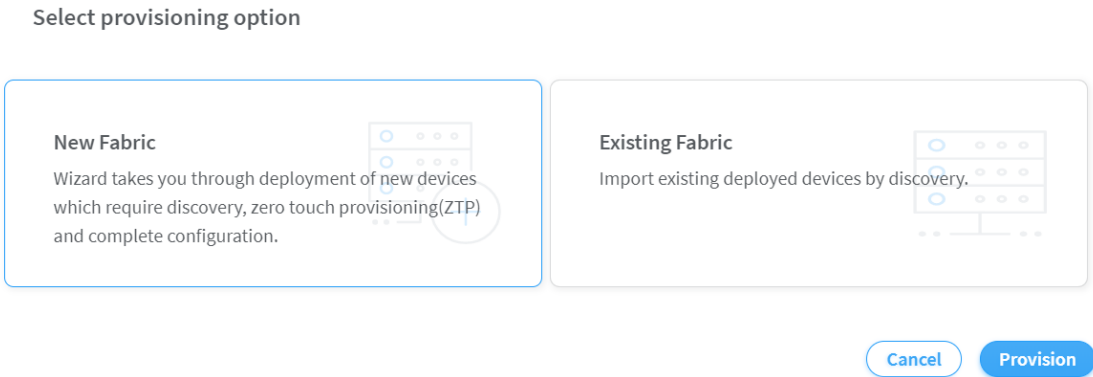
Onboard Greenfield Devices

You can use zero-touch-provisioning (ZTP) to deploy greenfield devices.

Follow these steps to deploy greenfield devices by using the Contrail Command UI.

1. Click **Infrastructure>Fabrics**.
The Fabrics page is displayed.
2. Click **Create**.
You are prompted to select a provisioning option.
3. Click **New Fabric** to deploy new (greenfield) devices by discovery. See [Figure 21 on page 67](#).

Figure 21: Select Existing Fabric



4. Click **Provision**.

The Create Fabric page is displayed.

5. Enter the information given in [Table 9 on page 68](#) if you have selected New Fabric as the provisioning option.

Table 9: Provisioning Option - New Fabric

Field	Action
Name	<p>Enter a name for the fabric.</p> <p>The name identifies the fabric on all fabric configuration and monitoring pages.</p>
Device credentials	<p>Enter root user password.</p> <p>The password entered in this field becomes the root password to access every device in the fabric.</p>
Overlay ASN (iBGP)	<p>Enter autonomous system (AS) number in the range of 1-65,535.</p> <p>If you enable 4 Byte ASN in Global Config, you can enter 4-byte AS number in the range of 1-4,294,967,295.</p>
Device Info	<p>Upload YAML file.</p> <p>This YAML file contains the serial numbers of each device in the fabric for device discovery. Click browse and navigate to the local directory and select the YAML file. Click Open to confirm.</p> <p>Alternatively, you can drag and drop the .yaml or .yml file in the Device Info box.</p> <p>To create this YAML file, click (* .yml) in the Template field, download the file, modify the file to include the serial numbers and hostnames for your fabric devices, and save the file.</p> <p>For a sample YAML file, see "No Link Title" on page 72.</p>

Table 9: Provisioning Option - New Fabric (Continued)

Field	Action
Node profiles	<p>Add node profiles.</p> <p>You can add more than one node profile.</p> <p>All preloaded node profiles are added to the fabric by default. You can remove a node profile by clicking X on the node profile. For more information, see "View Node Profile Information" on page 109.</p> <p>For more information on supported hardware platforms, associated node profiles and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185.</p>
Upgrade devices during the process?	<p>Select the Upgrade devices during the process? check box as given in Figure 22 on page 73 to enable the OS Version list.</p> <p>Starting with Contrail Networking Release 1907, you can upgrade a device during the ZTP process.</p>
OS Version	<p>Select the OS version you want to upgrade the device to, from the OS Version list.</p> <p>The OS Version list is enabled when you select the Upgrade devices during the process? check box.</p> <p>NOTE: The options in the OS Version list are the OS versions of the images that you uploaded.</p>
Disable VLAN-VN Uniqueness Check	<p>Select this check box when you are using the enterprise style of configuration but want to disable the requirement that every VLAN ID must have a 1:1 mapping with a VNI. Enterprise style of configuration is enabled by selecting the VLAN-ID Fabric-Wide Significance check box.</p>

Table 9: Provisioning Option - New Fabric (Continued)

Field	Action
VLAN-ID Fabric Wide Significance	<p>Select the VLAN-ID Fabric Wide Significance check box to enable enterprise style of configuration for the CRB-Access role on QFX devices. Deselect the check box to enable service provider style of configuration for the CRB-Access role. The check box is selected by default since enterprise style is the default setting.</p> <p>Once configured you can modify the enterprise style setting to service provider style of configuration. However, you cannot modify the service provider style to enterprise style of configuration without having to recreate the fabric.</p> <p>NOTE: Contrail Networking Release 1909 supports QFX10002-60C devices running Junos OS Release 19.1R2 and later. QFX10002-60C device works only if enterprise style of configuration is enabled. To enable enterprise style of configuration, select the VLAN-ID Fabric Wide Significance check box when onboarding the QFX10002-60C device. For more information on enterprise style of configuration, see "Configuring EVPN VXLAN Fabric with Multitenant Networking Services" on page 289.</p> <p>For more information on supported hardware platforms and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185.</p>
Underlay ASNs (eBGP)	<p>Enter autonomous system (AS) number in the range of 1-65,535.</p> <p>If you enable 4 Byte ASN in Global Config, you can enter 4-byte AS number in the range of 1-4,294,967,295.</p> <ul style="list-style-type: none"> • Enter minimum value in ASN From field. • Enter maximum value in ASN To field.

Table 9: Provisioning Option - New Fabric (Continued)

Field	Action
Management subnets	<p>Enter the following information to auto-assign management IP addresses to devices:</p> <p>CIDR—Enter the block of IP addresses that will be assigned as management IP addresses. The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Gateway—Enter gateway address for the devices in the management subnet that connect to the fabric.</p>
Fabric subnets (CIDR)	<p>Enter fabric CIDR address. The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.</p>
Loopback subnets (CIDR)	<p>Enter loopback subnet (lo0) address.</p> <p>The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.</p>
LR Loopback subnets	<p>Enter an IP subnet to be assigned as loopback interface (lo0) addresses used in Logical Routers (LR). The LR loopback interface IP address is required for eBGP peering to external or unmanaged devices.</p> <p>The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p>
PNF Servicechain subnets	<p>Enter the IP subnet for allocating IP addresses in the PNF Servicechain subnets field to establish EBGP session between PNF device and SPINE switch. This is an optional field that should be left blank when you are not creating service chains.</p>

Table 9: Provisioning Option - New Fabric (Continued)

Field	Action
Advanced interface filters	<p>Create an interface filter to filter the interfaces to include in the fabric. By default, all interfaces identified as participating in Contrail are imported into the fabric during the fabric provisioning process. If an interface filter is set, the fabric provisioning process includes the interfaces that are participating in Contrail and that match the interface filter in the fabric.</p> <p>To create an interface filter, choose the operation as regex and enter the filter characters in the Expression field. The Expression field supports all characters - including metacharacters - allowed in Python regex filters. For example, you can enter <code>^xe</code> in the Expression field to filter out all 10Gbps xe interfaces from the fabric.</p>
Import configured interfaces	<p>Choose this option if configured interfaces need to be imported into the fabric in addition to runtime interfaces. With some exceptions, a configured interface is generally an interface that has been configured in the Junos OS software.</p> <p>A runtime interface is generally an interface that has not been configured in Junos OS. You can confirm which interfaces are configured interfaces by entering the <code>show interfaces</code> command at the configuration mode prompt(<code>#</code>) in Junos. Only runtime interfaces are imported into the fabric by default.</p>

Sample YAML File Snippet

```

supplemental_day_0_cfg:
  - name: 'cfg1'
    cfg: |
      set system ntp server 167.XX.XX.XX
device_to_ztp:
  - serial_number: 'serial number'
    supplemental_day_0_cfg: 'cfg1'
    hostname: '<host name>'
    device_functional_group: 'dfg1'
  - serial_number: 'serial number'

```

```

supplemental_day_0_cfg: 'cfg1'
- serial_number: 'serial number'
- serial_number: 'serial number'

```

where,

supplemental_day_0_cfg is the additional configuration that is pushed on to the device during ZTP.

serial_number is the serial number of the device that is added to the fabric.

hostname is the device host name. If host name is not set, the serial number of the device is set as the device host name by default.

Figure 22: Deploy Greenfield Devices

6. Click Next.

The Device discovery page is displayed.

The **Device discovery progress** bar on the Device discovery page displays the progress of the device discovery job.

Figure 23: Device Discovery Progress Bar

Device discovery progress



The list of devices discovered are listed in the Discovered devices page.

7. Select the device you want to add by selecting the check box next to the device name.
You can select more than one device.
8. Click **Next** to assign roles.
Assign the Roles page is displayed.
9. From the assign to devices table, select the device you want to assign a role to by selecting the check box next to the device name.
Click the **Assign** icon at the end of the row to assign roles. The Assign role to devices pop-up is displayed.
10. You can now assign physical roles and routing-bridging roles.
 - a. Select a physical role from the Physical Role list.
 - b. Select a routing-bridging role from the Routing Bridging Roles list.

Assigning Roles for Spine Devices:

- Select **spine** from the Physical Role list.
- Select **CRB-Gateway** from the Routing Bridging Roles list.

Assigning Roles for Leaf Devices:

- Select **leaf** from the Physical Role list.
- Select **CRB-Access** from the Routing Bridging Roles list.

Assigning Roles for PNF Devices:

- Select **PNF** from the Physical Role list.
- Select **CRB-Access** and **PNF-Servicechain** from the Routing Bridging Roles list.

NOTE: The number of PNF instances you can create depends on the subnet mask of the pnf-servicechain-subnet that you provided during fabric onboarding. You can create multiple /29 subnets from the pnf-servicechain-subnet.

For example, if a /24 subnet is provided for the pnf-servicechain-subnet, then, you can create $2^5 = 32(29-24=5)$ subnets out of it. Each PNF uses a pair of /29 subnets. Thus, for a /24 subnet, you can have a maximum of 16 PNFs.

Assigning Roles for VNF Devices:

- Select **VNF** from the Physical Role list.
- Select **CRB-Access** from the Routing Bridging Roles list.

NOTE: **ERB-UCAST-Gateway** routing bridging role is also supported.

Click **Assign** to confirm selection.

11. Click **Autoconfigure** to initiate the auto-configuration job.

The Autoconfigure page is displayed.

The **Autoconfigure progress** bar on the Discovered devices page displays the progress of the auto-configuration job. Once the auto-configuration job is completed, click **Next**. The Assign Telemetry Profiles page is displayed.

Starting with Contrail Networking Release 2008, you can apply MTU, admin state, flow control, LACP force up, interface type attributes to physical interfaces; and MTU to logical interfaces. These attributes are applied to physical and logical interfaces after you **Autoconfigure** the devices.

To apply these attributes to interfaces:

- a. Navigate to **Infrastructure > Fabric**.
- b. Select the desired fabric from the list.
- c. Select the desired fabric device from the list.
- d. Click **Physical Interfaces > Create**.
- e. Enter the required details.

Figure 24: Create Physical Interface

Create Physical Interface

Interface Permissions

Name*

Type

MTU ⓘ

Description

Admin State On

Flow Control ⓘ

- f. Click **Create**.
- g. Click **Logical Interfaces > Create**.
- h. Enter the required details.

Create Logical Interface

Interface Permissions

Name* ⓘ

Connected Physical Interface*

MTU ⓘ

Description

i. Click **Create**.

12. (Optional) Assign telemetry profiles. For more information, see ["Assign Telemetry Profiles"](#) on page 29.

PNF service chain and VNF service chain does not use telemetry profiles.

13. Click **Finish** to exit the Create Fabric wizard.

The onboarding job is now complete.

NOTE: After the devices are onboarded, if you edit the fabric topology by adding new spine or leaf devices or by adding new links between devices, you *must* onboard the edited devices again. If you do not onboard the devices after edits to the initial configuration, underlay formation for the edited devices fails. You can choose to onboard individual devices by clicking the **Onboard**

button for the selected device in the **Fabric Devices** tab of the **Infrastructure > Fabrics > Fabric_Name** page.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2008	Starting with Contrail Networking Release 2008, you can apply MTU, admin state, flow control, LACP force up, interface type attributes to physical interfaces; and MTU to logical interfaces.
1909	Contrail Networking Release 1909 supports QFX10002-60C devices running Junos OS Release 19.1R2 and later.
1908	Select the VLAN-ID Fabric Wide Significance check box to enable enterprise style of configuration for the CRB-Access role on QFX devices. Deselect the check box to enable service provider style of configuration for the CRB-Access role.
1907	Starting with Contrail Networking Release 1907, you can upgrade a device during the ZTP process.

Device Import

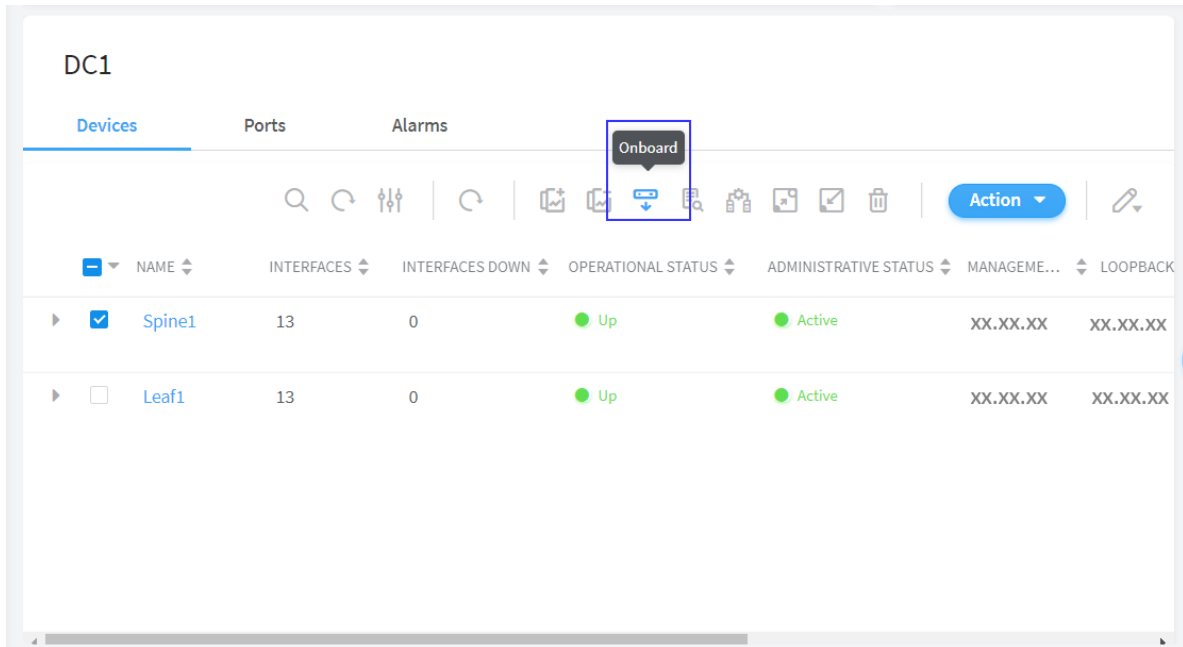
You can import devices or onboard devices to an existing fabric by using the Contrail Command user interface (UI).

Contrail Networking Release 2011 also supports topology discovery when you run the onboard device job.

Follow these steps to import devices by using the Contrail Command UI.

1. Navigate to **Infrastructure>Fabrics**.
The Fabrics page is displayed.
2. Click the name of the fabric to view the list of fabric devices.
3. Select the device you want to onboard by selecting the check box next to the name of the device.
You can select more than one device at a time.
4. Click the **Onboard** icon as show in [Figure 25 on page 79](#).


Figure 25: Onboard Fabric Device



The Onboard devices pop-up is displayed.

5. From the Onboard devices pop-up, you can:

Figure 26: Onboard Devices Pop-up



- a. Select and import certain types of interfaces to the device(s).

To import a particular type of interface, select the check box next to the type of interface. Uncheck all check boxes to import all types of interfaces.

If a particular interface is not displayed as a option, you can select regex while adding advanced interface filters as described in the following step.

b. Add advanced interface filters.

You can create an interface filter to filter interfaces to be included in the fabric. By default, all interfaces identified as participating in Contrail are imported into the fabric during the fabric provisioning process. If an interface filter is set, the fabric provisioning process includes the interfaces that are participating in Contrail and that match the interface filter in the fabric.

To add advanced interface filters, click **+Add**. The Operation and Expression fields are enabled.

Select **regex** from the Operation list. Enter filter characters in the Expression field.

The Expression field supports all characters, including metacharacters, that are allowed in Python regex filters. For example, you can enter `^xe` in the Expression field to filter out all 10Gbps xe interfaces from the fabric.

c. Import configured interfaces.

Select the Import Configured Interfaces check box only if configured interfaces need to be imported into the fabric in addition to runtime interfaces.

With some exceptions, a configured interface is an interface that has been configured in the Junos OS software. A runtime interface is an interface that has not been configured in Junos OS. Only runtime interfaces are imported into the fabric by default.

6. Click **Onboard** to onboard device(s).

The device is now onboarded.

Alternatively, to cancel device onboard, click **Cancel**.

With Contrail Networking Release 2011, topology discovery job is also initiated when you run the onboard devices job.

Follow these steps to view physical interfaces connected to the device you just onboarded. You can view fabric links from the results of the topology discovery job as well.

1. Navigate to **Infrastructure>Fabrics**.

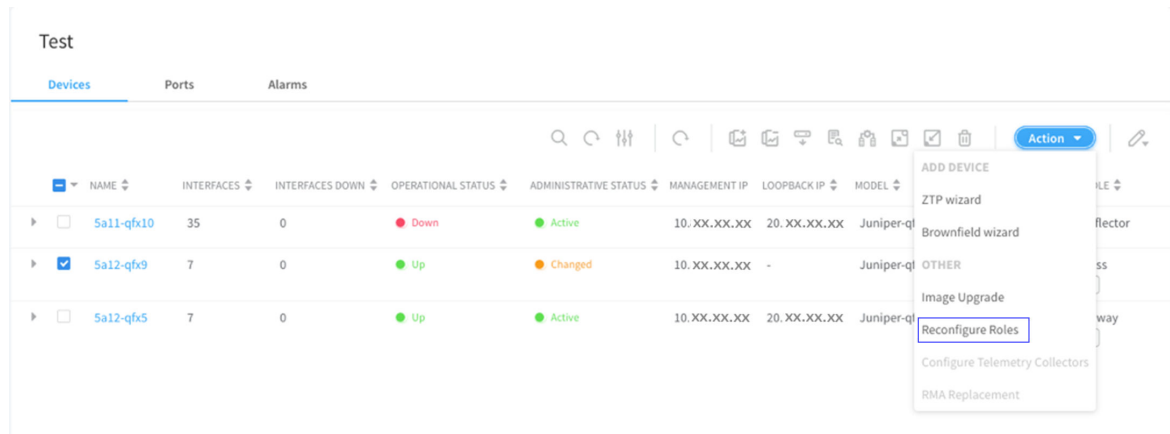
The Fabrics page is displayed.

2. Click the name of the fabric to view the list of fabric devices.

3. Select the device that was onboarded by selecting the check box next to the name of the device.

4. Click **Action>Reconfigure Roles** as shown in [Figure 27 on page 81](#).

Figure 27: Reconfigure Roles



5. Click **Autoconfigure** to reinitiate the auto-configuration job.

The Autoconfigure page is displayed.

The Autoconfigure progress bar on the Discovered devices page displays the progress of the auto-configuration job. Once the auto-configuration job is completed, click **Next**. The Assign Telemetry Profiles page is displayed.

6. (Optional) Assign telemetry profiles. For more information, see ["Assign Telemetry Profiles" on page 29](#).
7. Click **Finish** to exit the wizard.
8. From the list of devices that is displayed, click the name of the device that you just onboarded.

The device overview is displayed.

9. Click the **Physical Interfaces** tab to view the physical interfaces of the selected device.

Figure 28: Topology Discovery

▶	<input type="checkbox"/>	et-0/0/30	-	External	rvpg-test2	1	● ENABLED	...
▶	<input type="checkbox"/>	et-0/0/33	-	Fabric	-	2	● ENABLED	...
▶	<input type="checkbox"/>	et-0/0/34	-	Access	rvpg-test	1	● ENABLED	...
▶	<input type="checkbox"/>	et-0/0/35	-	Access	rvpg-test2	1	● ENABLED	...
▶	<input type="checkbox"/>	lo0	-	-	-	2	● ENABLED	...

Interfaces that are connected to other devices in the fabric are marked as **Fabric** as seen in [Figure 28 on page 81](#). Interfaces that are connected to external devices are marked as **External**. Interfaces that are not initially connected to a fabric device or external device but belong to a Virtual Port Group are marked as **Access**. Interfaces without a type are default interfaces.

10. (Optional) To view logical interfaces and hardware inventory of the device you just onboarded, click the **Logical Interfaces** tab and **Hardware Inventory** tab respectively.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2011	Contrail Networking Release 2011 also supports topology discovery when you run the onboard device job.

Create Virtual Network

A virtual network is a collection of endpoints, such as virtual machine instances, that can communicate with each other. You can also connect virtual networks to your on-premises network. A virtual network in a EVPN VXLAN data center corresponds to a bridge domain for one tenant in a multi-tenant data center fabric.

Follow these steps to create a virtual network from the Contrail Command user interface (UI).

1. Navigate to **Overlay>Virtual Networks**.

The All Networks page is displayed.

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

The Create Virtual Network page is displayed.

3. Enter a name for the network in the **Name** field.

4. Select VN Fabric Type.

Select **Routed** to enable routed virtual network functionality. A routed virtual network represents a layer 3 subnet between the fabric (border gateway) and the third-party physical network device. For more information, see ["Using Static, eBGP, PIM, and OSPF Protocols to Connect to Third-Party Network Devices" on page 258](#).

Select **Switched** (default option) for tenant virtual network on leaf, bare metal server, or vRouter.

5. Select network policies from the **Network Policies** list. You can select more than one network policy.

Network policies provide connectivity between virtual networks by allowing or denying specified traffic. They define the access control lists to virtual networks. To create a new network policy, navigate to **Overlay>Network Policies**.

For more information on creating network policies, see "[Create Network Policy](#)" on page 92.

NOTE: You can attach a network policy to the virtual network after you have created the virtual network.

6. Select any one of the following preferred allocation mode.
 - Flat subnet only
 - Flat subnet preferred
 - (Default) User defined subnet only
 - User defined subnet preferred

An allocation mode indicates how you choose a subnet. You select **Flat subnet only** or **Flat subnet preferred** allocation mode when the subnet is shared by multiple virtual networks. However, you select **(Default) User defined subnet only** or **User defined subnet preferred** allocation mode when you want to define a subnet range.

7. Enter subnet information as given in [Table 10 on page 83](#).

Table 10: Subnet Information

Field	Action
Network IPAM	Select the IP address management method that controls IP address allocation, DNS, and DHCP for the subnet.
CIDR	Enter the overlay subnet CIDR.
Allocation Pools	Enter a list of ranges of IP addresses for vRouter-specific allocation.
Gateway	Enter the gateway IP address of the overlay subnet. This field is disabled by default. To configure this field, uncheck Auto Gateway.
Service Address	Specify the user configured IP address for DNS Service instead of the default system allocated one.

Table 10: Subnet Information (Continued)

Field	Action
Auto Gateway	This check box is enabled by default and gateway address is allocated by the system. When this box is unchecked, gateway address is user configurable.
DHCP	Select this check box if you want Contrail to provide DHCP service.
DNS	Select this check box if you want the vRouter agent to provide DNS service.

8. Enter host route information.

Host routes are a list of prefixes and next hops that are passed to the virtual machine through DHCP.

a. Route Prefix—Enter a full CIDR value with an IP address and a subnet mask. For example, 10.0.0.0/24.

b. Next Hop—Enter next hop address.

9. Enter floating IP pool information.

A floating IP address is an IP address (typically public) that can be dynamically assigned to a running virtual instance. You can configure floating IP address pools in project networks, then allocate floating IP addresses from the pool to virtual machine instances in other virtual networks.

a. Pool Name—Enter pool name.

b. Projects—Select project from the list.

10. Enter fat flows information. See [Table 11 on page 84](#).

You can apply fat flows to all VMIs under the configured VN. Fat flows help reduce the number of flows that are handled by Contrail.

Table 11: Configure Fat Flow

Field	Action
Protocol	Select the application protocol.

Table 11: Configure Fat Flow (Continued)

Field	Action	
Port	<p>Enter a value between 0 through 65,535. Enter 0 to ignore both source and destination port numbers.</p> <p>NOTE: If you select ICMP as the protocol, the Port field is not enabled.</p>	
Ignore Address	<p>Configure fat flows to support aggregation of multiple flows into a single flow by ignoring source and destination ports or IP addresses. If you select Destination, only the Prefix Aggregation Source fields are enabled. If you select Source, only the Prefix Aggregation Destination fields are enabled. If you select the None (selected by default), both Prefix Aggregation Source and Prefix Aggregation Destination fields are enabled.</p>	
Prefix Aggregation Source	Source Subnet	<p>Enter the source IP address.</p> <p>Ensure that the source subnet of the flows match. For example, enter 10.1.0.0/24 to create fat flows with 10.1.0.0/24 as the subnet. The valid subnet mask range is /8 through /32.</p> <p>NOTE: For packets from the local virtual machine, source refers to the source IP of the packet. For packets from the physical interface, source refers to the destination IP of the packet.</p>
	Prefix	<p>Enter source subnet prefix length.</p> <p>The prefix length you enter is used to aggregate flows matching the source subnet. For example, when the source subnet is 10.1.0.0/16 and prefix length is 24, the flows matching the source subnet is aggregated to 10.1.x.0/24 flows. The valid the prefix length range is /(subnet mask of the source subnet) through /32.</p>

Table 11: Configure Fat Flow (Continued)

Field		Action
Prefix Aggregation Destination	Destination Subnet	<p>Enter the destination IP address.</p> <p>Ensure that the destination subnet of the flows match. Enter 10.1.0.0/24 to create fat flows with 10.1.0.0/24 as the subnet. The valid subnet mask range is /8 through /32.</p> <p>NOTE: For packets from the local virtual machine, destination refers to the destination IP of the packet. For packets from the physical interface, destination refers to the source IP of the packet.</p>
	Prefix	<p>Enter the destination subnet prefix length.</p> <p>The prefix length you enter is used to aggregate flows matching the destination subnet. For example, when the source subnet is 10.1.0.0/16 and prefix length is 24, the flows matching the source subnet is aggregated to 10.1.x.0/24 flows. The valid prefix length range is /(subnet mask of the destination subnet) through /32.</p>

11. Enter routing policy and bridge domain information as given below.

a. Select routing policy from the **Routing Policies** list.

To create a routing policy, navigate to **Overlay>Routing>Routing Policy**.

b. Define a list of route target prefixes.

Enter an IP address in the ASN field and Target in the range 0 through 65,535, or ASN in the range 1 through 65,535 and Target in the range 1 through 4,294,967,295 if 4-byte ASN is disabled. If 4-byte ASN is enabled, enter ASN in the range 1 through 4,294,967,295 and Target in the range 0 through 65,535.

c. Define export route targets.

You can advertise the matched routes from the local virtual routing and forwarding (VRF) table to the MPLS routing table.

Enter an IP address in the ASN field and Target in the range 0 through 65,535, or ASN in the range 1 through 65,535 and Target in the range 1 through 4,294,967,295 if 4-byte ASN is disabled. If 4-byte ASN is enabled, enter ASN in the range 1 through 4,294,967,295 and Target in the range 0 through 65,535.

d. Define import route targets.

Import the matched routes from the MPLS routing table and to the local virtual routing and forwarding (VRF) table.

Enter an IP address in the ASN field and Target in the range 0 through 65,535, or ASN in the range 1 through 65,535 and Target in the range 1 through 4,294,967,295 if 4-byte ASN is disabled. If 4-byte ASN is enabled, enter ASN in the range 1 through 4,294,967,295 and Target in the range 0 through 65,535.

- e. Enter bridge domain information. See [Table 12 on page 87](#).

A bridge domain is a set of logical interfaces that share the same flooding or broadcast characteristics.

Table 12: Bridge Domains

Field	Action
Name	Enter a name for the Layer 2 or Layer 3 bridge domain.
I-SID	Enter a Service Identifier in the range from 1 through 16777215.
MAC Learning	Enable or disable MAC learning. MAC learning is the process of obtaining the MAC addresses of all the nodes in a virtual network. It is enabled by default.
MAC Limit	Configure the maximum number of MAC addresses that can be learned.
MAC Move Limit	Configure the maximum number of times a MAC address move occurs in the MAC move time window. A MAC move is when a MAC address appears on a different physical interface or within a different unit of the same physical interface.
Time Window (secs)	Configure the period of time over which the MAC address move occurs. The default period is 10 seconds.
Aging Time (secs)	Configure the MAC table aging time, the maximum time that an entry can remain in the Ethernet Switching table before it is removed. The default time period is 300 seconds.

12. Enter advanced configuration information as given in [Table 13 on page 88](#).

Table 13: Advanced Configuration

Field	Action
Admin State	Select the administrative state of the virtual network.
Reverse Path Forwarding	Enable or disable Reverse Path Forwarding (RPF) check for the virtual network.
Shared	Select to share the virtual network with all tenants.
External	Select the check box to make the virtual networks reachable externally.
Allow Transit	Select to enable the transitive property for route imports.
Mirroring	Select to mark the virtual network as a mirror destination network.
Flood Unknown Unicast	Select to flood the network with packets with unknown unicast MAC address. By default, the packets are dropped.
Multiple Service Chains	Select to allow multiple service chains within two networks in a cluster.
IP Fabric Forwarding	Select to enable fabric based forwarding.
Forwarding Mode	Select the packet forwarding mode for the virtual network.
Extend to Physical Router(s)	Select the physical router to which you want to extend the logical router. The physical router provides routing capability to the logical router.
Static Route(s)	Select the static routes to be added to this virtual network.

Table 13: Advanced Configuration (*Continued*)

Field	Action
QoS	Select the QoS to be used for this forwarding class.
Security Logging Object(s)	Select the security logging object configuration for specifying session logging criteria.
ECMP Hashing Fields	<p>Configure one or more ECMP hashing fields.</p> <p>When configured all traffic destined to that VN will be subject to the customized hash field selection during forwarding over ECMP paths by vRouters.</p>
PBB Encapsulation	Select to enable Provider Backbone Bridging (PBB) EVPN tunneling on the network.
PBB ETree	<p>Select to enable PBB ETree mode on the virtual network which allows L2 communication between two end points connected to the vRouters.</p> <p>When the check box is deselected, end point communication happens through an L3 gateway provisioned in the remote PE site.</p>
Layer2 Control Word	Select to enable adding control word to the Layer 2 encapsulation.
SNAT	Select to provide connectivity to the underlay network by port mapping.
MAC Learning	<p>Enable or disable MAC learning.</p> <p>MAC learning is the process of obtaining the MAC addresses of all the nodes in a virtual network. It is enabled by default.</p>
Provider Network	<p>Select the provider network.</p> <p>The provider network specifies VLAN tag and the physical network name.</p>

Table 13: Advanced Configuration (Continued)

Field	Action
IGMP enable	Enable or disable IGMP.
Multicast Policies	Select the multicast policies. To create a policy, navigate to Overlay>Multicast Policies .
Max Flows	Enter the maximum number of flows permitted on each virtual machine interface of the virtual network.

13. Click Create.

The All Networks page is displayed. The virtual network that you created is displayed on this page.

Create Logical Routers

A logical router replicates the functions of a physical router. It connects multiple virtual networks. A logical router performs a set of tasks that can be handled by a physical router, and contains multiple routing instances and routing tables.

Follow these steps to create a logical router (LR).

1. Navigate to **Overlay>Logical Routers and click **Create**.**

The Create Logical Routers page is displayed.

2. Enter the following information as given in [Table 14 on page 90](#).**Table 14: Create a Logical Router**

Field	Action
Name	Enter a name for the Logical Router.
Admin State	Select the administrative state that you want the device to be in when the router is activated. Up is selected by default.

Table 14: Create a Logical Router (Continued)

Field	Action
Logical Router Type	Select SNAT Routing or VXLAN Routing from the list.
Choose Fabric	Select the fabric that you are associating this logical router to.
Connected Networks	Select the networks that you want to connect this logical router to.
Extend to Physical Router	Select the physical router(s) to which you want to extend virtual networks or routed virtual networks to, from the Extend to Physical Router list. A physical router provides routing capability to the logical router.
Reconfigure Physical Routers	This link is enabled when you select a routed virtual network from the Connected networks list. Click Reconfigure Physical Router to reconfigure a physical router that you want to extend a virtual network to. For more information, refer to the Create Logical Routers section of the "Using Static, eBGP, PIM, and OSPF Protocols to Connect to Third-Party Network Devices" on page 258 topic.
Public Logical Router	(Optional) Select this check box if you want the logical router to function as a public logical router.
NAT	Select this check box to enable Network Address Translation (NAT). This check box is disabled by default.
VxLAN Network Identifier	Enter VXLAN network identifier in the range from 1 through 16,777,215. This field is disabled by default.
DHCP IP Address	Enter DHCP relay server IP address. You can add more than one IP address. To add another address, click +Add .

Table 14: Create a Logical Router (*Continued*)

Field	Action
Route Target(s)	<p>Click +Add to add route targets.</p> <p>Enter Autonomous System (AS) number in the ASN field.</p> <ul style="list-style-type: none"> • Enter ASN in the range of 1-4,294,967,295, when 4 Byte ASN is enabled in Global Config. • Enter ASN in the range of 1-65,535, when 4 Byte ASN is disabled. • You can also add suffix <i>L</i> or <i>l</i> (<i>lower-case L</i>) at the end of a value in the ASN field to assign an AS number in 4-byte range. Even if the value provided in the ASN field is in the range of 1-65,535, adding <i>L</i> or <i>l</i> (<i>lower-case L</i>) at the end of the value assigns the AS number in 4-byte range. If you assign the ASN field a value in the 4-byte range, you must enter a value in the range of 0-65,535 in the Target field. <p>Enter route target in the Target field.</p> <ul style="list-style-type: none"> • Enter route target in the range of 0-65,535, when 4 Byte ASN is enabled and ASN field is assigned a 4-byte value. • Enter route target in the range of 0-4,294,967,295, when the ASN field is assigned a 2-byte value.

3. Click **Create** to create the logical router.

The Logical Routers page is displayed.

NOTE: The router_interface object (Virtual Port) is created as part of the LR creation and VN extension to Spines workflow. While planning the IP address for spines, you must be aware that an extra one IP address is required for the router_interface object which gets created automatically.

Create Network Policy

A network policy is a set of access control rules that can be attached to virtual networks. A network policy determines what traffic that is allowed or denied on the network.

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

1. Navigate to **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 a name for the policy in the Policy Name field.

4. Enter the following information as given in [Table 15 on page 93](#) to define a policy rule.

You can define more than one rule for a policy.

Table 15: Define Policy Rule

Field	Action
Action	To allow traffic to pass through the network, select Pass . To deny traffic, select Deny .
Protocol	Select a protocol you want to associate with traffic. Any is selected by default.
Source Type	Select the source type for this policy rule.
Source	Select the traffic source based on the source type you have selected. For example, if you select CIDR as the Source Type, enter the source subnet in the Source field.
Source Port	Leave the default option, Any , as is.
Direction	Determine the direction of traffic flow that you want to apply this policy rule. You can select < or > .
Destination Type	Select the destination type for this policy rule.
Destination	Select the traffic destination based on the destination type you have selected. For example, if you select CIDR as the Destination Type, enter the destination subnet in the Destination field.
Destination Ports	Leave the default option, Any , as is.

Table 15: Define Policy Rule *(Continued)*

Field	Action
Advanced Options	Select this check box to view more options that you can configure for this policy rule.
Services	Select the network services you want to apply to this policy rule.
QoS	Select the QoS you want to apply to this policy rule.
Log	Select this check box to log traffic pattern.
Mirror	Select this check box to mirror traffic pattern.

5. (Optional) Click **+Add** to add another policy rule.

6. Click **Create** to create the network policy.

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

(Optional) Attach a network policy to a virtual network.

1. Navigate to **Overlay>Virtual Networks**.

The All networks page is displayed.

2. To select the virtual network you want to add the policy to, select the check box next to the name of the virtual network. Then click the **Edit** icon at the end of the row.

The Edit Virtual Network page is displayed.

3. Select the network policy from the Network Policies list and click **Save**.

The policy is now added and the All networks page is displayed.

Create Network IPAM

A network IP Address Management (IPAM) enables you to manage DNS and DHCP services that assign IP addresses to hosts on a network.

Follow these steps to create a network IPAM by using the Contrail Command UI.

1. Navigate to **Overlay>Network IPAM** click **Create**.

The Create IP Address Management page is displayed.

2. Enter a name for the network IPAM in the **Name** field.

3. Select a subnet method to indicate how you choose a subnet.

Select **User Defined** option button when you want to define a subnet range. Select **Flat** option button when the subnet is shared by multiple virtual networks.

4. Follow these steps if you select **Flat** as the subnet method.

The Subnet(s) section is displayed when you select **Flat** as the subnet method.

- a. Enter valid IPv4 subnet or mask in the **CIDR** field

- b. Enter the gateway IP address in the **Gateway**.

The Gateway field is disabled by default. Clear the **Auto Gateway** check box to enable this field.

- c. **Auto Gateway** is selected by default.

Clear this check box to manually enter gateway IP address in the Gateway field.

- d. **DHCP** check box is selected by default.

Dynamic Host Configuration Protocol (DHCP) dynamically assigns IP addresses to hosts on a network.

- e. Click **+Add** in the Allocation Pool(s) section and add the following information.

An allocation pool is the subnet pool from the defined CIDR, from which Contrail Networking allocates IP addresses.

- **Start (Allocation Pool)**—Enter the starting IP address in the range of IP addresses that can be allocated.
- **End (Allocation Pool)**—Enter the ending IP address in the range of IP addresses that can be allocated.
- **vRouter Specific Pool**—This check box is selected by default. This is the pool from which vRouter allocates IP addresses to workloads.

5. Select a method to associate an IPAM to a DNS Server from the **DNS Method** list.

- Select **Default** when the DNS resolution for virtual machines are performed based on the name server configuration.
- Select **Tenant** to use tenant DNS servers.

If you select **Tenant** radio option, the Tenant DNS Server IPs section is enabled. Enter DNS server IP information in **DNS Server IP** field. You can add more IP addresses by clicking **+Add**.

- Select **Virtual DNS** to use virtual DNS servers to resolve DNS requests from virtual machines.

If you select **Virtual DNS** radio option, select virtual DNS information from the **Virtual DNS** field that is enabled.

- Select **None** for no DNS support.

6. Enter IPv4 address of NTP server in the **NTP Server IP** field.

7. Enter a domain name for the NTP server in the **Domain Name** field.

The Domain Name field is enabled only when you have selected Default, Tenant, or None as the DNS method.

8. Click **Create** to create the IPAM.

The IP Address Management page is displayed.

Reconfigure Roles

You can reconfigure fabric device roles after you have completed the fabric onboarding process. For more information on assigning roles to devices during the fabric onboarding process, see ["Assign a Role to a Device" on page 25](#).

Follow these steps to reconfigure a device role.

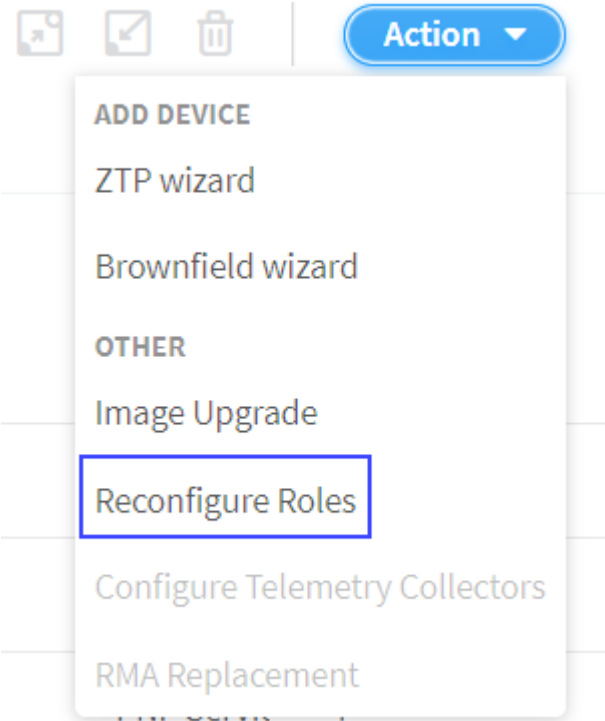
1. Navigate to **Infrastructure>Fabrics**.

The Fabrics page is displayed.

2. Click the name of the fabric to view the list of fabric devices.

3. Click **Action>Reconfigure Roles**. See [Figure 29 on page 97](#).

Figure 29: Reconfigure Roles



The Assign to devices page is displayed.

4. Follow these steps to reassign a role to a device.
 - a. From the Assign to devices table, select the device you want to assign a role to by selecting the check box next to the device name.
 - b. Click the **Assign** icon at the end of the row to assign roles. See [Figure 30 on page 97](#).

Figure 30: Assign a Role to a Device

The image shows a table titled 'Assign to devices'. The table has columns for NAME, MANAGEMENT IP, NODE PROFILE, ROLE, ROUTING ROLES, and AUTOCONFIGURE. Two rows are visible: 'DC2-Spine2' and 'DC2-Leaf3'. The 'DC2-Leaf3' row is selected, indicated by a blue checkmark in the first column. At the end of the 'DC2-Leaf3' row, there is an 'Assign Role' icon (a blue square with a white refresh symbol) which is highlighted with a red rectangular border.

NAME	MANAGEMENT IP	NODE PROFILE	ROLE	ROUTING ROLES	AUTOCONFIGURE
<input type="checkbox"/> DC2-Spine2	10.XX.XX.XX	juniper-qfx10k	-	-	False
<input checked="" type="checkbox"/> DC2-Leaf3	10.XX.XX.XX	juniper-qfx5k-lean	-	-	False

Follow these steps to assign roles to multiple devices of the same node profile.

- a. From the assign to devices table, select the devices you want to assign a role to by selecting the check box next to the device name.

- b. Click **Assign roles to devices of same node profile** as shown in [Figure 31 on page 98](#).

Figure 31: Assign Role to Multiple Devices

NAME	MANAGEMENT IP	NODE PROFILE	ROLE	ROUTING ROLES	AUTOCONFIGURE
<input type="checkbox"/> DC2-Spine2	10.XX.XX.XX	juniper-qb10k	-	-	False
<input checked="" type="checkbox"/> DC2-Leaf3	10.XX.XX.XX	juniper-qb5k-lean	-	-	False
<input checked="" type="checkbox"/> DC2-Leaf2	10.XX.XX.XX	juniper-qb5k-lean	-	-	False

The Assign role to devices pop-up is displayed.

5. You can now assign physical roles and routing-bridging roles.

- a. Select a physical role from the Physical Role list.

For example, select **spine** role.

- b. Select a routing-bridging role from the Routing Bridging Roles list.

For example, select **CRB-Gateway** role. You can select more than one routing-bridging role.

For more information on supported hardware platforms, associated node profiles and roles, see ["Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles"](#) on page 185.

- c. Click **Assign** to confirm selection.

6. Click **Autoconfigure** to reinitiate the auto-configuration job.

The Autoconfigure page is displayed.

The **Autoconfigure progress** bar on the Discovered devices page displays the progress of the auto-configuration job. Once the auto-configuration job is completed, click **Next**. The Assign Telemetry Profiles page is displayed.

For more information on assigning telemetry profiles, see ["Assign Telemetry Profiles"](#) on page 29.

RELATED DOCUMENTATION

[Assign a Role to a Device](#) | 25

[Assign Telemetry Profiles](#) | 29

Managing Custom Roles

IN THIS SECTION

- [Adding Custom Roles | 99](#)
- [Backup and Restore Custom Roles | 104](#)

These topics provide instructions to add custom roles, backup custom roles, and restore custom roles.

Adding Custom Roles

Follow these steps to add a custom overlay role to device(s) managed by Contrail Enterprise Multicloud (CEM).

1. Run the following command from the server running the device manager container.

```
# docker exec -it config_devicemgr_1 bash
```

2. Define the custom role object in `predef_payloads.json`.

```
(config-device-manager) # vi /opt/contrail/fabric_ansible_playbooks/conf/  
predef_payloads.json
```

3. Define custom overlay roles.

All roles are defined in the “data”->“object_type”: “overlay-role”, “objects”: [section. As an example, `motd-test` is defined as a custom overlay role.

```
"object_type": "overlay-role",  
  "objects": [  
    {  
      "fq_name": [  
        "default-global-system-config", "motd-test"  
      ],
```

```

    "name": "motd-test"
  },

```

4. After the roles are defined, map the custom role to devices and underlay roles.

All device profiles are listed in the “data”->“object_type”: “node-profile” section. Custom roles are added in the node_profile_roles -> role_mappings section.

As an example, add motd-test custom role to juniper-mx device that is already configured with leaf physical role.

```

{
  "fq_name": [
    "default-global-system-config", "juniper-mx"
  ],
  "name": "juniper-mx",
  "node_profile_vendor": "Juniper",
  "node_profile_device_family": "junos",
  "node_profile_hitless_upgrade": true,
  "node_profile_roles": {
    "role_mappings": [
      {
        "physical_role": "leaf",
        "rb_roles": ["CRB-Access", "CRB-Gateway", "DC-Gateway", "Route-Reflector",
"DCI-Gateway", "ERB-UCAST-Gateway", "DCI-Gateway", "CRB-MCAST-Gateway", "PNF-Servicechain",
"AR-Client", "motd-test"]
      },
      {
        "physical_role": "spine",
        "rb_roles": ["lean", "CRB-Access", "CRB-Gateway", "DC-Gateway", "Route-
Reflector", "CRB-MCAST-Gateway", "DCI-Gateway", "PNF-Servicechain", "AR-Client"]
      }
    ]
  }
}

```

5. Edit the all.yml file.

```

(config-device-manager) # vi /opt/contrail/fabric_ansible_playbooks/group_vars/all.yml

```

6. Create a directory, motd_test, and configure motd-test as leaf physical role in the feature_based_plugin_roles section of the all.yml file.

The Jinja template for each custom role and underlay role is stored in this directory.

```
feature_based_plugin_roles:
  motd-test@leaf:
    - motd_test
  CRB-Access@leaf:
    - overlay_storm_control
    - overlay_telemetry
```

NOTE: The name of the directory is `motd_test`. The name of the custom role is `motd-test`.

7. Add the `motd_test` directory to the configuration apply order in the `feature_apply_order` section.

```
feature_apply_order:
  - basic
  - underlay_ip_clos
  [...]
  - overlay_telemetry
  - motd_test
```

8. Add Jinja templates.

The device configuration templates are located here:

```
/opt/contrail/fabric_ansible_playbooks/config_templates/
```

Create `motd_test` directory.

```
(config-device-manager) # cd /opt/contrail/fabric_ansible_playbooks/config_templates/
(config-device-manager) # mkdir motd_test
(config-device-manager) # cd motd_test
```

9. Add custom role to a device.

The configuration file name depends on the device type. In this example, the file name for the devices are as follows:

- MX Series devices: `juniper_junos_motd_test.j2`
- QFX Series devices: `juniper_junos-qfx_motd_test.j2`

Create a configuration file, `juniper_junos_motd_test.j2`, to add a custom role for the MX series device.

```
(config-device-manager) # vi juniper_junos_motd_test.j2
```

10. Configure the file using the `set` command.

```
set groups {{cfg_group}} system login message MOTD_TEST
```

Using `{{cfg_group}}` allows you to separate custom group configuration from predefined CEM roles.

11. Exit the container and then restart it.

```
(config-device-manager) # exit  
# docker restart config_devicemgr_1
```

The new role is now seen in the Contrail Command user interface (UI). See [Figure 32 on page 103](#).


Figure 32: motd-test Routing-Bridging Role


Assign role to 1 devices

Physical Role* 

leaf 

Routing Bridging Roles 

DC-Gateway  

CRB-Gateway 

CRB-Gateway 

CRB-MCAST-Gateway

DC-Gateway 

DCI-Gateway

ERB-UCAST-Gateway

motd-test 

Route-Reflector

12. (Optional) After applying the role, log in to the device and confirm that the configuration is applied.

```

MX> show configuration | compare rollback 1
[edit groups]
  __contrail_overlay_networking__ { ... }
+  __contrail_motd_test__ {
+    system {
+      login {
+        message MOTD_TEST;
+      }
+    }
+  }
[edit]
- apply-groups [ re0 __contrail_basic__ __contrail_underlay_ip_clos__
__contrail_underlay_infra_bms_access__ __contrail_overlay_bgp__ __contrail_overlay_evpn__
__contrail_overlay_evpn_access__ __contrail_overlay_evpn_gateway__
__contrail_overlay_evpn_type5__ __contrail_overlay_dhcp_relay__
__contrail_overlay_security_group__ __contrail_overlay_lag__
__contrail_overlay_multi_homing__ __contrail_overlay_fip_snat__
__contrail_overlay_networking__ ];
+ apply-groups [ re0 __contrail_basic__ __contrail_underlay_ip_clos__
__contrail_underlay_infra_bms_access__ __contrail_overlay_bgp__ __contrail_overlay_evpn__
__contrail_overlay_evpn_access__ __contrail_overlay_evpn_gateway__
__contrail_overlay_evpn_type5__ __contrail_overlay_dhcp_relay__
__contrail_overlay_security_group__ __contrail_overlay_lag__
__contrail_overlay_multi_homing__ __contrail_overlay_fip_snat__
__contrail_overlay_networking__ __contrail_motd_test__ ];

```

Backup and Restore Custom Roles

IN THIS SECTION

- [Backup Custom Roles | 105](#)
- [Restore Custom Roles | 105](#)

These topics provide instructions to backup and restore custom roles.

Backup Custom Roles

Follow these steps to backup custom roles.

1. Backup `/opt/contrail/fabric_ansible_playbooks/conf/predef_payloads.json`.

```
# docker exec -it config_devicemgr_1 cat /opt/contrail/fabric_ansible_playbooks/conf/
predef_payloads.json > predef_payloads.json.bak
```

2. Verify the overlay roles and node profiles.

For example, when you back up custom roles during the upgrade process, the existing containers are removed and the custom roles are erased. You will not be able to restore these roles and configurations once it has been erased. Hence, ensure that you back up the correct file by verifying the roles and node profiles before the upgrade process.

3. Backup `all.yml`.

```
# docker exec -it config_devicemgr_1 cat /opt/contrail/fabric_ansible_playbooks/group_vars/
all.yml > all.yml.bak
```

4. Verify the `feature_based_plugin_roles` and `feature apply order` sections.
5. Backup custom roles.

```
# docker exec -it config_devicemgr_1 tar --exclude "overlay*" -czvf custom_roles.tar.gz /opt/
contrail/fabric_ansible_playbooks/config_templates/
# docker cp config_devicemgr_1:custom_roles.tar.gz
```

Ensure that you add `--exclude` before `"overlay*"` when you back up custom roles. If you do not add `--exclude`, both custom roles as well as predefined roles are backed up. This might cause any predefined roles that had bug fixes for the next release to be overwritten.

For example, assume that there is a fix in the Jinja template for `overlay_evpn` predefined role for Contrail Networking Release 2005. When you upgrade Contrail Networking Release 2003 to Contrail Networking Release 2005, and subsequently use the backup file (that was backed up not using `--exclude`), all fixes related to predefined roles for Contrail Networking Release 2005 will be overwritten when you use the config template from the backup file.

Restore Custom Roles

Follow these steps to restore custom roles.

1. Copy the new `/opt/contrail/fabric_ansible_playbooks/conf/predef_payloads.json` file.

```
# docker exec -it config_devicemgr_1 cat /opt/contrail/fabric_ansible_playbooks/conf/
predef_payloads.json > predef_payloads.json.new
```

2. Verify overlay roles and node profiles.

```
# diff -u predef_payloads.json.bak predef_payloads.json.new
```

NOTE: New roles might have been added when CEM was updated.

3. Restore the `predef_payloads.json` changes.

- a. Open `predef_payloads.json`.

```
docker exec -it config_devicemgr_1 vi /opt/contrail/fabric_ansible_playbooks/conf/
predef_payloads.json
```

- b. Verify that all roles are defined in the `"data"->"object_type": "overlay-role"`, `"objects": [` section, and `motd-test` is defined as a custom overlay role.

For more information, see the `predef_payloads.json.bak` file. For a file diff, refer to section two of the `predef_payloads.json.bak` file.

```
"object_type": "overlay-role",
  "objects": [
    {
      "fq_name": [
        "default-global-system-config", "motd-test"
      ],
      "name": "motd-test"
    },
  ],
```

- c. Verify that the custom role is mapped to devices and underlay roles.

All device profiles are listed in the `"data"->"object_type": "node-profile"` section. Custom roles are added in the `node_profile_roles -> role_mappings` section.

As an example, add `motd-test` custom role to `juniper-mx` device that is already configured with leaf physical role.

```
{
  "fq_name": [
    "default-global-system-config", "juniper-mx"
  ],
  "name": "juniper-mx",
  "node_profile_vendor": "Juniper",
  "node_profile_device_family": "junos",
  "node_profile_hitless_upgrade": true,
  "node_profile_roles": {
    "role_mappings": [
      {
        "physical_role": "leaf",
        "rb_roles": ["CRB-Access", "CRB-Gateway", "DC-Gateway", "Route-Reflector",
"DCI-Gateway", "ERB-UCAST-Gateway", "DCI-Gateway", "CRB-MCAST-Gateway", "PNF-
Servicechain", "AR-Client", "motd-test"]
      },
      {
        "physical_role": "spine",
        "rb_roles": ["lean", "CRB-Access", "CRB-Gateway", "DC-Gateway", "Route-
Reflector", "CRB-MCAST-Gateway", "DCI-Gateway", "PNF-Servicechain", "AR-Client"]
      }
    ]
  }
}
```

4. Backup all.yml.

```
# docker exec -it config_devicemgr_1 cat /opt/contrail/fabric_ansible_playbooks/group_vars/
all.yml > all.yml.new
```

5. Verify changes in `feature_based_plugin_roles` and `feature` apply order.

```
# diff -u all.yml.bak all.yml.new
```

NOTE: New roles might have been added when CEM was updated.

6. Restore all.yml.

For more information, see the all.yml.bak file. For a file diff, refer to section five of the all.yml.bak file.

a. Open all.yml.

```
docker exec -it config_devicemgr_1 vi /opt/contrail/fabric_ansible_playbooks/group_vars/
all.yml
```

b. Verify that feature_based_plugin_roles has all roles and role mappings. Ensure that motd-test role is also added.

```
feature_based_plugin_roles:
  motd-test@leaf:
    - motd_test
  CRB-Access@leaf:
    - overlay_storm_control
    - overlay_telemetry
```

c. Verify that the feature_apply_order describes the order of templates that are applied on devices. Ensure that motd_test is also added.

```
feature_apply_order:
  - basic
  - underlay_ip_clos
  [...]
  - overlay_telemetry
  - motd_test
```

7. Restore custom roles Jinja templates.

```
# docker cp custom_roles.tar.gz config_devicemgr_1:/
# docker exec -it config_devicemgr_1 tar xzvf custom_roles.tar.gz
```

8. Restart the container and verify roles.

```
# docker restart config_devicemgr_1
```

RELATED DOCUMENTATION

| *Upgrading Contrail Networking using Contrail Command*

View Node Profile Information

You can view basic device information, vendor information, vendor hardware information, supported routing bridging roles, supported physical roles, assigned devices, and node permission information of a node on the Node Profiles page of the Contrail Command UI.

Follow these steps to view node profiles:

1. Click **Infrastructure>Fabrics>Node Profiles**.

The Node Profiles page is displayed. See [Figure 33 on page 109](#).

Figure 33: Node Profiles

Node Profiles		
NAME	DEVICE FAMILY	VENDOR
▶ juniper-mx	junos	Juniper
▶ juniper-qfx10k	junos-qfx	Juniper
▶ juniper-qfx10k-lean	junos-qfx	Juniper
▶ juniper-qfx5k	junos-qfx	Juniper
▶ juniper-qfx5k-lean	junos-qfx	Juniper
▶ juniper-srx	junos	Juniper

2. Select the node profile you want to view by clicking the arrow next to the node profile name.

NOTE: The supported node profiles are juniper-mx, juniper-qfx10k, juniper-qfx10k-lean, juniper-qfx5k, juniper-qfx5k-lean, juniper-qfx5k-erb-only, juniper-qfx5120 and juniper-srx.

For more information on supported hardware platforms, associated node profiles and roles, see ["Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles"](#) on page 185.

The details and permissions of the node profile are displayed.

By default, all preloaded node profiles are available for devices in a fabric.

RELATED DOCUMENTATION

[Create a Fabric | 7](#)

[Discover a Device | 21](#)

[Assign a Role to a Device | 25](#)

[Delete a Fabric | 30](#)

[Image Management | 54](#)

[Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles | 185](#)

Monitoring Fabric Jobs

Starting with Contrail Networking Release 2008, you can view a detailed summary of all Contrail Command initiated jobs and transactions for the past three days in the **Monitoring > Operations** page in Contrail Command. The **Operations** page replaces the **Jobs** page in the UI. You can view fabric transaction type, completion status, the start and end times, as well as the execution IDs.

Click the plus (+) icon to expand the fabric transaction details. From release 2008, you can also view supplemental device configurations resulting from active and completed transactions. These supplemental configuration pushes are not directly initiated jobs, but device configurations that occur as a result of user-initiated fabric transactions. Examples of fabric transactions include fabric onboarding, roles assignment, hitless upgrade, fabric deleting, and so on. When you click to expand a transaction or a parent job, you can see the supplemental device configurations, henceforth referred to as child jobs, and their details as well. The transaction status is displayed as **In progress** for ongoing transactions, **Completed** for completed transactions, and **Warning** for transactions with failed child jobs.

Figure 34: Monitoring > Operations : Fabric Transactions and Jobs

JOB TYPE	STATUS	START	END	EXECUTION ID
Role Assignment	Warning	07/28/2020 11:18:14 PM	07/28/2020 11:21:19 PM	1595958488290_993679fe-024d-4887-a653-cd5179169702
Device Config Push for fab01 > 5a12-qfx14	Completed	07/28/2020 11:20:46 PM	07/28/2020 11:21:19 PM	1595958640733_a4140352-a74b-4124-a89b-5716137b637f
Device Config Push for fab01 > 5a12-qfx14	Error	07/28/2020 11:19:32 PM	07/28/2020 11:19:32 PM	1595958560693_74bfd852-22ed-4b2e-8690-b01aa4630e1e
Device Config Push for fab01 > 5a12-qfx14	Error	07/28/2020 11:18:38 PM	07/28/2020 11:19:17 PM	1595958510552_a86ed40f-43b7-4f48-ab8d-a3b269cf2b23
Role Assignment for fab01	Completed	07/28/2020 11:18:14 PM	07/28/2020 11:18:41 PM	1595958488290_993679fe-024d-4887-a653-cd5179169702
Role Assignment	Completed	07/28/2020 10:55:00 PM	07/28/2020 10:55:25 PM	1595957094641_8e1920cb-4fdf-41b0-9744-493b1fe69f68
Virtual Port Group 'vpg1' Update	Completed	07/28/2020 10:17:26 PM	07/28/2020 10:17:43 PM	req-7d056cd8-677b-4c47-8a83-c59774ef0eca
Role Assignment	Completed	07/28/2020 1:35:53 PM	07/28/2020 1:36:50 PM	1595923547599_1b6ef7653-e184-4d00-bc1a-10afa670876d
Virtual Port Group 'vpg1' Create	Completed	07/28/2020 1:33:41 PM	07/28/2020 1:34:07 PM	req-ea32b643-d3fe-4a1c-8208-269cdc882ce9
Logical Router 'LR1' Create	Warning	07/28/2020 1:32:20 PM	07/28/2020 1:33:07 PM	req-41bfbe3b-ba76-4007-85f0-53f5155c4137
Role Assignment	Completed	07/28/2020 1:11:51 PM	07/28/2020 1:12:52 PM	1595922105732_0e872a8b-d7c8-4ebd-97a4-cf167b195ac0
Fabric Onboarding	Completed	07/28/2020 1:06:33 PM	07/28/2020 1:10:40 PM	1595921787092_2de67256-8b20-45ff-b568-8d5d2f3255eb

Click a transaction to view more information including the actual configuration of the job. When you click a transaction, you are navigated to the **Job Details** page with two panes. The left pane lists all the child jobs that are created as a result of the fabric transaction and the actual parent job is listed at the bottom. You can view the name, status, start and end times, and execution IDs of all the jobs, similar to the previous main page. In ongoing jobs, the parent job status is completed only after all the child jobs have completed successfully. If a child job fails, the child job status is displayed as **Error** but the parent job status is displayed as **Completed**.

NOTE: For completed transactions with one or more failed child jobs, the main fabric transaction status is displayed as **Warning**, but the parent job status is displayed as **Completed**.

Click a job to view the full job logs including error messages and the actual configuration pushes of the job. The logs are displayed in the right pane in the **Logs** tab. Click the **Configuration** tab to view the configuration information of that job. You can also download the logs and configuration information using the download button. Additionally, you can search the logs and configuration details for specific keywords using the search option.

Figure 35: Job Details

The screenshot displays the 'Job Details' page in the Contrail Command interface. The breadcrumb navigation shows 'MONITORING > Operations > Job Details'. The page title is 'Role Assignment'. On the left, there is a table titled 'Jobs for Role Assignment' with columns for NAME, STATUS, START, END, and EXECUTION ID. The table contains four rows, all with a 'Completed' status. On the right, there is a 'Configuration' section with a code editor showing network configuration details, including BGP group settings and neighbor configurations. An 'Abort' button is visible in the top right corner of the configuration section.

NAME	STATUS	START	END	EXECUTION ID
Device Config Push...	Completed	07/28/2020 1:36:33 PM	07/28/2020 1:36:50 PM	1595923587857_fd...
Device Config Push...	Completed	07/28/2020 1:36:04 PM	07/28/2020 1:36:25 PM	1595923558131_32...
Device Config Push...	Completed	07/28/2020 1:36:03 PM	07/28/2020 1:36:22 PM	1595923557842_0c...
Role Assignment f...	Completed	07/28/2020 1:35:53 PM	07/28/2020 1:36:05 PM	1595923547599_1...

```

{"config_diffs":
[edit groups __contrail_overlay_bgp__ policy-options policy-state
+ route-type external]
[edit groups __contrail_overlay_bgp__ policy-options policy-state
+ route-type external]
[edit groups __contrail_overlay_bgp__ routing-options resolution]
rib bgp.rtarget.0 { ... }
+ rib bgp.l3vpn.0 {
+ resolution-ribs [ inet.3 inet.0 ];
+ }
[edit groups __contrail_overlay_bgp__ protocols bgp group _contrail
+ family inet-vpn {
+ unicast;
+ }
+ family inet6-vpn {
+ unicast;
+ }
+ }
[edit groups __contrail_overlay_bgp__ protocols bgp group _contrail
+ neighbor 192.168.100.50 {
+ peer-as 64512;
+ }
- neighbor 10.10.10.252 {
- peer-as 64512;
- }
[edit groups __contrail_overlay_bgp__ protocols bgp]

```

You can also terminate an ongoing parent job but *only* a parent job in this page. The **Abort** button on the top right of the page is enabled for ongoing jobs. For completed or failed jobs, the **Abort** button is greyed out.

The **Operations** page enables you to view transactions and jobs directly or indirectly initiated by Contrail Command. You can view transactions history of three days. This page is particularly useful for troubleshooting purposes, because you can see job logs configuration details and error messages and you can see exactly where the job failed and debug the errors.

Caveats

- You cannot rollback configurations. This page is used solely for tracking progress and debugging issues and you cannot undo an operation.
- The child jobs are labeled on a best-effort basis. Multiple supplemental configuration pushes as a result of a complex parent job can mislabel jobs. For example, an update job can be labeled as a create job.

Prior to release 2008, you can view detailed information, status, and logs of all active, failed, and completed fabric jobs for the previous 24 hours in the **Monitoring > Jobs** page in Contrail Command. You can also terminate ongoing jobs from the job status monitoring page. The **Jobs** page is available only in releases 1912 through 2005.

Navigate to the **Monitoring > Jobs** page to view fabric jobs history and status. Alternatively, click the bell icon on the menu bar on the top of any page to view a truncated list of the latest jobs. Click **See All** to

view the complete list of active and completed jobs. The jobs are displayed in a descending order of the latest job to the oldest. You can also use the search option to search for a particular job.

The job summary information provides information on the job type, progress, start and end times, and the execution ID. The job status indicates if a job is currently ongoing, completed, or failed. Click the job type to view additional information including the job percentage completion information in the progress bar and the complete job logs. The job logs also display information about error messages for failed jobs.

You can also terminate an ongoing job when you click the job type. The **Abort** button on the top right of the page is enabled for ongoing jobs. For completed or failed jobs, the **Abort** button is greyed out.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2008	Starting with Contrail Networking Release 2008, you can view a detailed summary of all Contrail Command initiated jobs and transactions for the past three days in the Monitoring > Operations page in Contrail Command. The Operations page replaces the Jobs page in the UI.
1912	Prior to release 2008, you can view detailed information, status, and logs of all active, failed, and completed fabric jobs for the previous 24 hours in the Monitoring > Jobs page in Contrail Command.

RELATED DOCUMENTATION

[Terminating Ongoing Fabric Jobs](#) | 113

Terminating Ongoing Fabric Jobs

In Contrail Networking Release 1910, you can use Contrail Command user interface (UI) to terminate an ongoing fabric job.

You can terminate an ongoing fabric job in the following workflows:

- Provisioning fabric devices using Zero-touch-provisioning. For more information, see "[Provision Fabric Devices Using End-to-End ZTP](#)" on page 32.
- Importing existing brownfield devices and deploying new greenfield devices. For more information, see "[Create a Fabric](#)" on page 7.

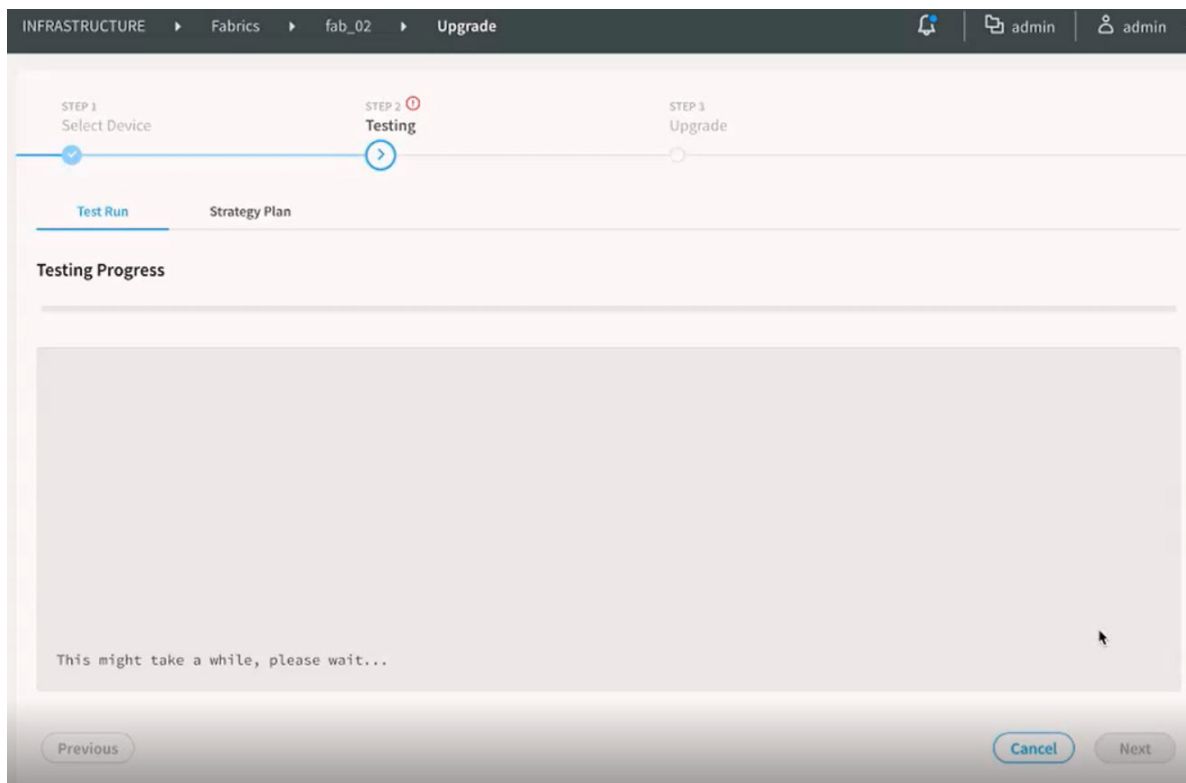
- Initiating auto-configuration job. For more information, see ["Assign a Role to a Device" on page 25](#).
- Device image upgrade job. For more information, see ["Performing Hitless Software Upgrade on Data Center Devices" on page 326](#).
- Activating or deactivating Maintenance Mode job. For more information, see ["Activating Maintenance Mode on Data Center Devices" on page 293](#).
- Discovery of servers. For more information, see ["Onboarding and Discovery of Bare Metal Servers" on page 404](#) and ["Fabric Discovery and ESXi Discovery by Using Contrail Command" on page 351](#).

In releases prior to release 1910, you could not terminate an ongoing fabric job.

For example, you can terminate an ongoing device image upgrade job while upgrading a device image in a fabric. The following steps provide instructions on how you can terminate the device image upgrade job:

1. [Figure 36 on page 114](#) displays that the image upgrade job is in progress.

Figure 36: Image Upgrade Job In Progress



2. To terminate this image upgrade job, click **Cancel**.
3. In the pop-up that is displayed, click **Abort and Exit**.

The image upgrade job is terminated.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1910	In Contrail Networking Release 1910, you can use Contrail Command user interface (UI) to terminate an ongoing fabric job.

RELATED DOCUMENTATION

| [Monitoring Fabric Jobs](#) | 110

Adding a Leaf or Spine Device to an Existing Fabric Using ZTP

Starting with Contrail Networking release 1911, you can expand an existing greenfield fabric deployment by adding new leaf or spine devices. The feature is especially useful when you do not add all the required devices to the fabric on Day One and want to add devices to the fabric at a later point. You can add new devices to a fabric by uploading a YAML file that contains the device information.

To add a device to a fabric:

1. Log into Contrail Command and navigate to Infrastructure > Fabrics > *Fabric Name*.
2. Click **Actions** > **ZTP Wizard**.

The Create Fabric page is displayed.

3. Click the **browse** link in the **Device Info** field and upload the YAML file that contains information about the device, such as the serial number of the device, that you want to add to the fabric.

Alternatively, you can drag and drop the .YAML or .yaml in the Device Info field. You can add multiple new devices at a time. To add multiple devices, specify the serial numbers of the devices and the configuration details for each device in the YAML file as shown in the sample below.

Sample YAML File

```
supplemental_day_0_cfg:
  - name: 'cfg1'
```

```

  cfg: |
    set system ntp server 167.xx.20.98
device_to_ztp:
- serial_number: '74035760356'
  supplemental_day_0_cfg: 'cfg1'
  hostname: 'test_device'
  device_functional_group: 'dfg1'
  loopback_ip: '10.xx.10.5'
  underlay_asn: '64015'
  mgmt_ip: '10.xx.13.42'
  to_ztp: false
- serial_number: '55674325815'
  supplemental_day_0_cfg: 'cfg1'
- serial_number: '11675330144'
- serial_number: '74656088411'

```

Table 16: Fields in the YAML file

Field	Description
serial_number	Specify the device serial number.
hostname	Enter a host name to identify the device. If a host name is not provided, the device serial number is used to identify the device.
device_functional_group	Enter a device functional group. A device functional group enables you to assign properties like OS version, and physical and routing-bridging roles to a user-defined group of devices instead of groups defined by node profiles. Using device functional groups enables configuring the devices in mixed-mode where devices in a single fabric can support different OS versions. Device functional groups are supported only for greenfield fabric onboarding.
loopback_ip	Enter an IP subnet to be assigned as loopback interface (lo0) addresses to fabric devices. The loopback interface IP addresses are required for iBGP peering in the overlay network. The field value must include a full CIDR with an IP address and a subnet mask.
underlay_asn	Enter the autonomous system number (ASN) that will be used to enable external BGP (eBGP) in the underlay network.

Table 16: Fields in the YAML file (Continued)

Field	Description
mgmt_ip	Enter the IP address to be assigned as management IP address for the devices.

Figure 37: Add a Device by Uploading a YAML File

You can see that the fields **Name**, **ASN Range**, **Fabric Subnet**, **Loopback Subnets**, and **PNF Servicechain Subnets** are grayed out.

4. Enter the root user password and click **Next** to proceed to the Device discovery page. Complete the steps in "[Discover a Device](#)" on page 21.
5. Click **Next** to assign roles and complete the steps in "[Assign a Role to a Device](#)" on page 25.
6. Click **Autoconfigure**.
7. Navigate to Infrastructure > Fabrics > **Fabric Name** > **Fabric Devices** page and verify that the new devices are listed.
8. Click **Action** > **Reconfigure Roles** to assign role to the new device.

NOTE: To add devices successfully using this method, you must ensure that ASN numbers, Fabric Subnet IPs, and Loopback Subnet IPs are available for the number of devices to be added.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1911	Starting with Contrail Networking release 1911, you can expand an existing greenfield fabric deployment by adding new leaf or spine devices.

RELATED DOCUMENTATION

[Discover a Device | 21](#)

[Assign a Role to a Device | 25](#)

Grouping Fabric Devices and Roles Using Device Functional Groups

Contrail Networking fabric management currently provides pre-defined node profiles to configure certain properties, such as supported routing-bridging roles, for a specified class of devices. Node profiles are defined on a per-vendor-family basis. Contrail Networking Release 1911 enables you to assign properties like OS version, and physical and routing-bridging roles to a user-defined group of devices using device functional groups (DFGs) instead of a grouping defined by node profiles. This is particularly useful in mixed mode where devices in a single fabric support multiple OS versions. These properties are applied while provisioning fabric devices using Zero Touch Provisioning (ZTP) or during device Return Material Authorization (RMA).

Contrail Command contains a set of predefined device functional groups. You can view existing groups in the **Device Functional Groups** tab of the **Infrastructure > Fabrics** page.

For the list of predefined device functional groups, see [Table 17 on page 119](#).

Table 17: List of Predefined Device Functional Groups

Device Functional Group	Description	OS Version	Routing-Bridging Roles
L2-Server-Leaf	Provides layer 2 servers connectivity with ingress replication for multicast in the spine.	18.4R2	CRB-Access
L3-Server-Leaf	Provides layer 3 servers connectivity.	19.1R3	ERB-UCAST-Gateway
L3-Storage-Leaf	Provides layer 3 connectivity to storage arrays.	18.4R2	ERB-UCAST-Gateway
L3-Server-Leaf-with-Optimized-Multicast	Provides layer 3 servers connectivity with optimized multicast traffic.	18.4R2	ERB-UCAST-Gateway, AR-Client
Centrally-Routed-Border-Spine	Provides layer 3 routing for layer 2 server leafs and route reflector and ingress replication. Provides DCGW service, DCI GW service, and connectivity to firewalls.	18.4R2	Route-Reflector, CRB-Gateway, DC-Gateway, DCI-Gateway, PNF-Servicechain
Centrally Routed-Border-Spine-With-Optimized-Multicast	Provides layer 3 routing and gateway services for layer 2 server leafs. Provides route reflector and assisted replication services.	18.4R2	Route-Reflector, AR-Replicator, CRB-Gateway, DC-Gateway, DCI-Gateway, PNF-Servicechain
Border-Spine-in-Edge-Routed	Provides layer 3 gateway service and route reflector service.	18.4R2	Route-Reflector, DC-Gateway, DCI-Gateway, PNF-Servicechain, ERB-UCAST-Gateway, CRB-MCAST-Gateway
Border-Leaf-in-Edge-Routed	Provides layer 3 gateway service and route reflector service.	18.4R2	Route-Reflector, DC-Gateway, DCI-Gateway, PNF-Servicechain, ERB-UCAST-Gateway, CRB-MCAST-Gateway

Table 17: List of Predefined Device Functional Groups (Continued)

Device Functional Group	Description	OS Version	Routing-Bridging Roles
Lean-Spine-with-Route-Reflector	Spine only acting as Route Reflector.	18.4R2	Route-Reflector, lean

For more information on supported hardware platforms and routing-bridging roles, see "[Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles](#)" on page 185.

You can also create custom device functional groups by clicking **Create** on the top right corner of the **Infrastructure > Fabrics > Device Functional Groups** page. Device functional groups are added in the `fabric_ztp.yml` file under **Device Info** used during fabric creation in the UI.

To group devices and assign properties using device functional groups, you must:

1. Create a new device functional group. Alternatively, you can use the predefined device functional groups.

To create a new device functional group.

- a. Click **Create** on the **Device Functional Groups** tab of the **Infrastructure > Fabrics** page. The **Create Device Functional Group** page appears.
- b. Enter the required information. You can select a physical role, multiple routing-bridging roles, and the associated devices. You can also specify the required OS version.
- c. Click **Create**. The newly created device functional group is listed in the **Device Functional Groups** tab.

2. the device functional group in the **Device Info** YAML file used during fabric creation.

To a device functional group.

- a. Click **Create** on the **Fabrics** tab of the **Infrastructure > Fabrics** page. The **Select Provisioning Option** page appears.
- b. Select the **New Fabric** option since device functional groups are supported only on greenfield deployments. The **Create Fabric** page appears.
- c. Edit the `fabric_ztp.yml` file under **Device Info** to add the device functional group. Add `device_functional_group: '<>'` to the `fabric_ztp.yml` YAML file. For a sample YAML file, see "[Create a Fabric](#)" on page 7.
- d. Enter the required information as per the steps provided in the "[Create a Fabric](#)" on page 7 topic and click **Next**. The **Device discovery** page is displayed.

- e. After you have completed the steps provided in the ["Discover a Device" on page 21](#) topic, click **Next**. The **Assign the Roles** page is displayed.
- f. The preassigned roles and device names from the previously defined device functional group is prepopulated and displayed. Click **Autoconfigure** to continue and complete the fabric creation process.

The device functional groups are used for image upgrade during ZTP, addition of new devices, and also during RMA.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1911	Contrail Networking Release 1911 enables you to assign properties like OS version, and physical and routing-bridging roles to a user-defined group of devices using device functional groups (DFGs) instead of a grouping defined by node profiles.

RELATED DOCUMENTATION

[Create a Fabric | 7](#)

[Assign a Role to a Device | 25](#)

[Return Material Authorization | 179](#)

Creating Layer 3 PNF Service Chains for Inter-LR Traffic

IN THIS SECTION

- [Onboard Fabric Devices | 123](#)
- [Configure Virtual Networks | 123](#)
- [Configure Virtual Port Groups | 124](#)
- [Configure Logical Routers | 124](#)

- [Configure PNF | 124](#)
- [View Service Appliance Sets and Service Appliances | 127](#)

Contrail Networking provides layer 3 physical network functions (PNF) support to create service chains for inter-LR (logical router) traffic. Contrail Networking automates configuration of QFX and SRX devices to allow movement of inter-LR traffic between bare metal servers through layer 3 PNF.

Figure 38: Example Topology

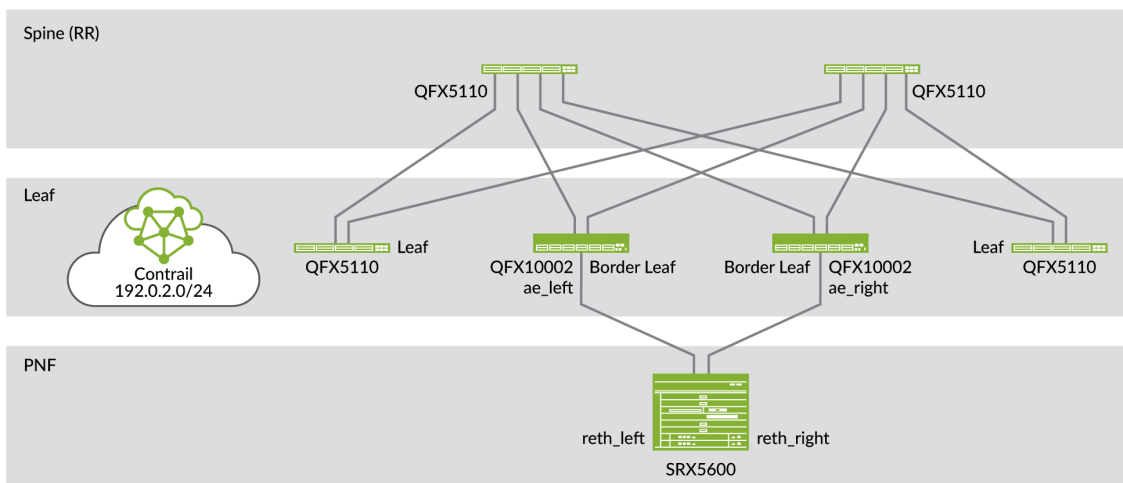



Figure 38 on page 122 shows an example topology of how a PNF device (SRX5600) is used to allow inter-LR traffic to pass through a service chain. You can use the SRX device as a layer 3 PNF device after you have configured the device during device onboarding. The PNF device is connected to border leaf or spine devices.

Getting Started

The general workflow to create a PNF service chain is as follows:



These topics provide instructions to create a PNF service chain.

Onboard Fabric Devices

Follow the steps provided in the ["Onboard Brownfield Devices" on page 57](#) topic to onboard brownfield fabric devices and assign roles to the devices.

While onboarding devices, ensure that you enter the IP subnet in the **PNF Servicechain subnets** field to establish EBGP session between PNF device and Spine switch.

See [Table 18 on page 123](#) for an example configuration of a centrally-routed bridging (CRB) architecture that includes PNF functionality. The SRX device uses the physical role, **pnf**, and routing-bridging role, **PNF-Servicechain**. The border leaf device uses **PNF-Servicechain** routing-bridging role.

Table 18: Assign Roles to Devices

Device	Physical Role	Routing-Bridging Role
Spine devices	spine	CRB-Gateway, Route-Reflector, CRB-MCAST-Gateway
Border leaf	leaf	PNF-Servicechain
Leaf devices	leaf	CRB-Access
SRX Device	pnf	PNF-Servicechain

Configure Virtual Networks

Follow the steps provided in the ["Create Virtual Network" on page 82](#) topic to create virtual networks.

After you have created the virtual networks, you create a network policy. For more information on creating a network policy and attaching the network policy to the virtual network, see ["Create Network Policy" on page 92](#).

Configure Virtual Port Groups

Follow the steps provided in the ["Configuring Virtual Port Groups" on page 238](#) topic to configure virtual port groups. A virtual port group defines leaf device interfaces attached to end hosts

Ensure that you assign the virtual port group to the virtual network that you created.

For example, when you create two virtual networks, VN-A and VN-B, you will have to create one virtual port group for VN-A and another for VN-B.

Configure Logical Routers

Follow the steps provided in the ["Create Logical Routers" on page 90](#) topic to configure logical routers.

While creating logical router, ensure that you

- Select **VXLAN Routing** as the Logical Router Type.
- Select the virtual network(s) from the Connected Networks list.
- Select the physical routers (the spine devices) to which you want to extend the logical router.

Configure PNF

Configuring PNF includes the following:

- Creating a PNF Service Template to define the physical connectivity of the PNF to the fabric.
- Creating a PNF Service Instance to define the interconnection of the two logical routers.

Follow these steps to create PNF service template and PNF service instance by using the Contrail Command UI.

1. Navigate to **Services>Deployments**.

The VNF Service Instances page is displayed.

2. Click the **PNF** tab.

The PNF Service Instances page is displayed.

3. Click **Create** and select **Instance (with Template)** from the list.

The Create PNF Service Instance page is displayed.

4. Enter the following information in the PNF Service Template pane.

Table 19: Enter PNF Service Template Information

Field	Action
Name	Enter a name for the PNF template.
PNF Device	Select the PNF device you want to use for this service chain.
PNF Left Interface	Select the left interface of the PNF device.
PNF Left Fabric	Select the fabric connected to the left interface of the PNF device.
PNF Left Attachment Points	Select the physical router attached to the left interface of the PNF device from the Physical Router list. Select the left interface of the physical router from the Left Interface list.
PNF Right Interface	Select the right interface of the PNF device.
PNF Right Fabric	Select the fabric connected to the right interface of the PNF device.
PNF Right Attachment Points	Select the physical router attached to the right interface of the PNF device from the Physical Router list. Select the right interface of the physical router from the Right Interface list.

5. Click **Next** to confirm.

The PNF Service Instance pane is displayed.

After you create the PNF service template, you can use the PNF service template to enable the PNF service instance.

6. Enter the following information in the PNF Service Instance Pane.

Table 20: Enter PNF Service Instance Information

Field	Action
Name	Enter a name for the PNF service instance.

Table 20: Enter PNF Service Instance Information (*Continued*)

Field	Action
Service Template	The PNF service template is selected by default.
PNF eBGP ASN	Enter the PNF eBGP AS number.
(Optional) Configure Static RP	<p>Select Configure Static RP check box to configure static rendezvous point (RP).</p> <p>The RP IP Address field is enabled. The RP is the router that receives multicast traffic.</p> <p>This field is required only when sending multicast traffic through the PNF service chain.</p>
(Optional) RP IP Address	<p>Enter the RP IP address of the router that receives multicast traffic.</p> <p>This field is required only when sending multicast traffic through the PNF service chain.</p>
Left Tenant Logical Router	Select the left tenant logical router. This interface is where the service chain starts.
PNF Left BGP Peer ASN	Displays the BGP AS number of the border leaf that the PNF device is connected to.
Left Service VLAN	<p>Enter left service VLAN ID.</p> <p>The VLAN ID must be unique.</p>
Right Tenant Logical Router	Select the right tenant logical router. This interface is where the service chain ends.
PNF Right BGP Peer ASN	Displays the BGP AS number of the border leaf that the PNF device is connected to.

Table 20: Enter PNF Service Instance Information (*Continued*)

Field	Action
Right Service VLAN	Enter right service VLAN ID. The VLAN ID must be unique.

- Click **Finish** to complete configuration.

The PNF Service Instances page is displayed. For a sample resulting configuration, see [Figure 39 on page 127](#).

Figure 39: Resulting Configuration

STATUS	SERVICE INSTANCE	SERVICE TEMPLATE	PNF eBGP ASN	LEFT LOGICAL ROUTER	LEFT SERVICE VLAN	RIGHT LOGICAL ROUTER	RIGHT SERVICE VLAN
Active	PNF-Instance-Test	PNF-Template-Test-ter	65112	PNF-LR-1	1111	PNF-LR-2	2222

Details		Permissions	
Instance Name	PNF-Instance-Test	Owner	ecd4c44227c440ab97701ca1d3d39ff4
Display Name	PNF-Instance-Test	Owner permissions	Read, Write, Refer
UUID	bba3dac5-94ea-4acb-966f-03ea10515f85	Global permissions	-
Template	PNF-Template-Test-template	Share	-
Port Tuples	PNF-Instance-Test;		
Status	Active		
Left Logical Router	PNF-LR-1		
Right Logical Router	PNF-LR-2		
PNF eBGP ASN	65112		
Left Service VLAN	1111		
Right Service VLAN	2222		

View Service Appliance Sets and Service Appliances

(Optional) Follow these steps to view Service Appliance Sets and Service Appliances by using the Contrail Command UI:

- Click **Services > Appliances**.

The Appliances page is displayed.

- Click **Service Appliance Sets** tab to view the list of available service appliance sets.
- Click **Service Appliance** tab to view the list of available service appliances.

Alternatively, you can also navigate to the **Monitoring>Operations** page to verify the status of the job.

RELATED DOCUMENTATION

[Fabric Overview | 4](#)

[Create a Fabric | 7](#)

Creating VNF Service Chains for Inter-LR Traffic

IN THIS SECTION

- [Onboard Brownfield Devices | 132](#)
- [Create Virtual Network | 142](#)
- [Configuring Virtual Port Groups | 150](#)
- [Create Logical Routers | 158](#)
- [Configure the Internal Virtual Networks | 160](#)
- [Create the Service Virtual Machine | 161](#)
- [Create VNF Service Template | 161](#)
- [Create VNF Service Instance | 162](#)
- [Create the Network Policy | 163](#)

Contrail Networking Release 1912 extends the service chaining functionality to bare metal servers (BMS). In earlier releases, Contrail Networking supports traffic flow between a virtual machine in one virtual network and a virtual machine in another virtual network. However, traffic flow between a virtual machine and BMS through a service chain was not supported. With Release 1912, Contrail Networking supports the movement of inter-LR traffic by using virtual network functions (VNF). This EVPN-based VXLAN (Ethernet VPN-based Virtual Extensible LAN) service chain supports bidirectional traffic flow through a service virtual machine.

VNF service chaining uses EVPN with VXLAN to enable traffic flow between:

- Two bare metal servers.

Figure 40: Traffic Flow Between Two Bare Metal Servers




Figure 40 on page 129 shows traffic flowing between two bare metal servers. Each bare metal server is connected to a logical router (virtual routing engine). These logical routers are configured to send traffic from the bare metal server in one virtual network to the bare metal server in the other virtual network, through the service virtual machine.

- A bare metal server and a virtual machine.

Figure 41: Traffic Flow Between a Bare Metal Server and a Virtual Machine




Figure 41 on page 129 shows traffic flowing between a bare metal server and a virtual machine. The bare metal server and the virtual machine are connected to logical routers. These logical routers are configured to send traffic from the bare metal server in one virtual network to the virtual machine in the other virtual network, through the service virtual machine.

- A virtual machine and a bare metal server.

Figure 42: Traffic Flow Between a Virtual Machine and a Bare Metal Server





Figure 42 on page 130 shows traffic flowing between a virtual machine and a bare metal server. The virtual machine and the bare metal server are connected to logical routers. These logical routers are configured to send traffic from the virtual machine in one virtual network to the bare metal server in the other virtual network, through the service virtual machine.

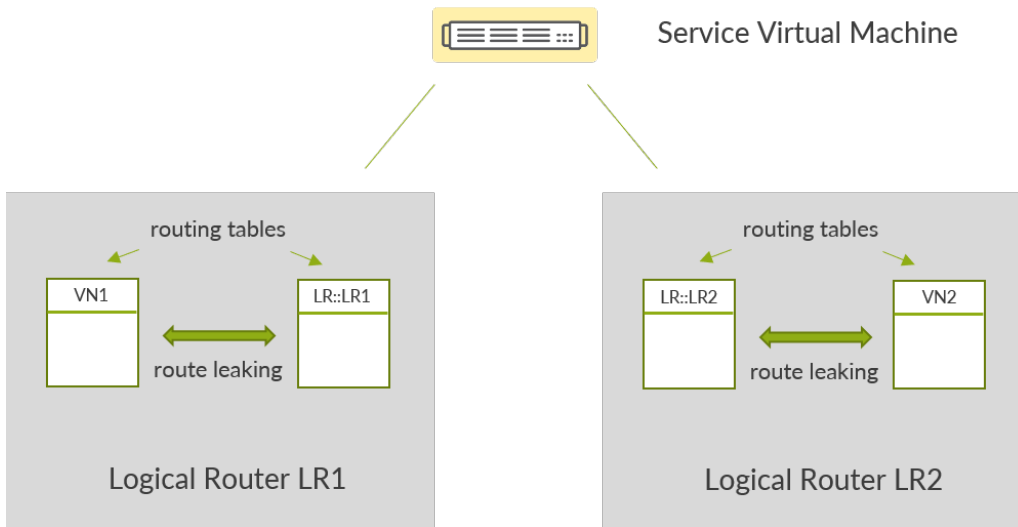
Figure 43 on page 130 shows how the service virtual machine, acting as a VNF, conceptually connects to the virtual networks.

Figure 43: VNF Connectivity



The VNF does not connect to VN1 and VN2 directly. Instead, the VNF connects to virtual networks (labelled LR::LR1 and LR::LR2) that are internally generated by Contrail Networking. These internal virtual networks learn of routes in VN1 and VN2 through route leaking, as shown in Figure 44 on page 131.

Figure 44: Route Leaking Between Internally-Generated and Explicitly-Created Virtual Networks



Contrail Networking creates LR::LR1 when you associate Logical Router LR1 with Virtual Network VN1, and LR::LR2 when you associate Logical Router LR2 with Virtual Network VN2. If you're not working with VNFs, then you can safely ignore these internally-generated virtual networks. If you're working with VNFs, as you are in this topic, then you must configure each of these internally-generated virtual networks with a subnet and associate these networks with the VNF. We'll show you how to do this later.

Routes are learned through route leaking and re-origination. This works as follows:

1. Routes to endpoints in VN1 are leaked to LR::LR1, and routes to endpoints in VN2 are leaked to LR::LR2.
2. Contrail Networking then installs LR::LR1 routes into LR::LR2, and LR::LR2 routes into LR::LR1. Prior to installing these routes, Contrail Networking re-originate the routes so that the service virtual machine is the next hop. This means that traffic going from LR::LR1 to LR::LR2 and from LR::LR2 to LR::LR1 will be routed to the service virtual machine.
3. The re-originated routes are then leaked from LR::LR1 to VN1, and from LR::LR2 to VN2.
4. Additionally, Contrail Networking configures the routing tables in the vRouter (on the server where the service virtual machine resides) so that it too has routes to VN1 and VN2.

The end result is that packets in one virtual network destined for the other virtual network are sent to the service virtual machine for processing.

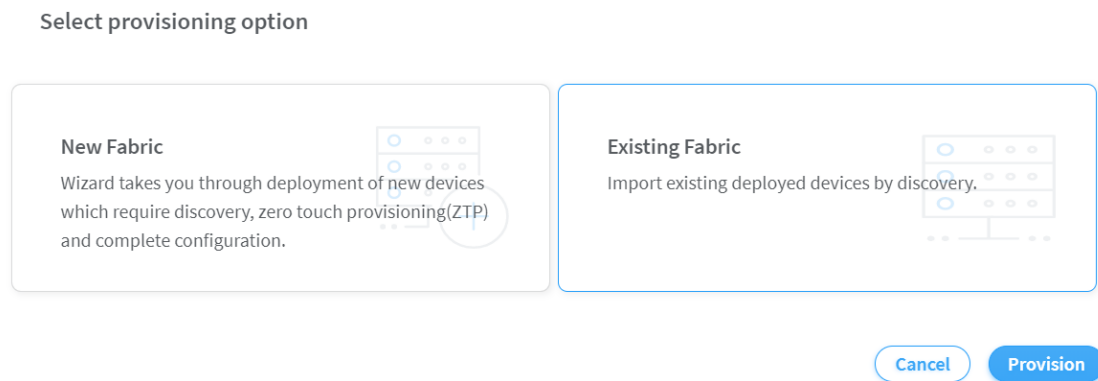
These topics provide instructions to create an EVPN-based VXLAN service chain.

Onboard Brownfield Devices

Follow these steps to onboard brownfield devices from the Contrail Command user interface (UI):

1. Click **Infrastructure**>**Fabrics**.
The Fabrics page is displayed.
2. Click **Create**.
You are prompted to select a provisioning option.
3. Click **Existing Fabric** to import existing (brownfield) devices by discovery. See [Figure 45 on page 132](#).

Figure 45: Select Existing Fabric



4. Click **Provision**.
The Create Fabric page is displayed.
5. Enter the information as given in [Table 21 on page 132](#).

Table 21: Provision Existing Fabric

Field	Action
Name	Enter a name for the fabric.
Overlay ASN (iBGP)	Enter autonomous system (AS) number in the range of 1-65,535. If you enable 4 Byte ASN in Global Config , you can enter 4-byte AS number in the range of 1-4,294,967,295.

Table 21: Provision Existing Fabric *(Continued)*

Field	Action
Node profiles	<p>Add node profiles.</p> <p>You can add more than one node profile.</p> <p>All preloaded node profiles are added to the fabric by default. You can remove a node profile by clicking X on the node profile. For more information, see "View Node Profile Information" on page 109.</p> <p>For more information on supported hardware platforms, associated node profiles and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185.</p>
Disable VLAN-VN Uniqueness Check	<p>Select this check box when you are using the enterprise style of configuration but want to disable the requirement that every VLAN ID must have a 1:1 mapping with a VNI. Enterprise style of configuration is enabled by selecting the VLAN-ID Fabric-Wide Significance check box.</p>

Table 21: Provision Existing Fabric (Continued)

Field	Action
VLAN-ID Fabric Wide Significance	<p>Select the check box to enable enterprise style of configuration for the CRB-Access role on QFX devices. De-select the check box to enable service provider style of configuration for the CRB-Access role. The check box is selected by default since enterprise style is the default setting.</p> <p>Once configured you can modify the enterprise style setting to service provider style of configuration. However, you cannot modify the service provider style to enterprise style of configuration without having to recreate the fabric.</p> <p>The service provider style of configuration allows for customization of Ethernet-based services at the logical interface level. Each logical interface is bound to a unique VLAN ID. With the enterprise style of configuration, logical interfaces are placed into Layer 2 mode by specifying ethernet-switching as the interface family. The ethernet-switching family can be configured only on a single logical unit, unit 0. For more information on enterprise and service provider type of configurations, see Flexible Ethernet Services Encapsulation.</p> <p>NOTE: Contrail Networking Release 1909 supports QFX10002-60C device running Junos OS Release 19.1R2 and later. QFX10002-60C device works only if enterprise style of configuration is enabled. To enable enterprise style of configuration, select the VLAN-ID Fabric Wide Significance check box when onboarding the QFX10002-60C device. For more information on enterprise style of configuration, see "Configuring EVPN VXLAN Fabric with Multitenant Networking Services" on page 289.</p> <p>For more information on supported hardware platforms and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185.</p>

Table 21: Provision Existing Fabric (Continued)

Field	Action
Device credentials	Enter the device credentials to access the fabric devices for discovery. If your fabric devices have different username and password combinations for device access, click the + Add option to add additional username and password credentials.
Management subnets	<p>Enter the following information to auto-assign management IP addresses to devices:</p> <p>CIDR—Enter the block of IP addresses that will be assigned as management IP addresses. The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Gateway—Enter gateway address for the devices in the management subnet that connect to the fabric.</p>
Loopback subnets	<p>Enter loopback subnet (lo0) address.</p> <p>The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.</p>
Underlay ASNs (eBGP)	<p>Enter autonomous system (AS) number in the range of 1-65,535.</p> <p>If you enable 4 Byte ASN in Global Config, you can enter 4-byte AS number in the range of 1-4,294,967,295.</p> <ul style="list-style-type: none"> • Enter minimum value in ASN From field. • Enter maximum value in ASN To field.

Table 21: Provision Existing Fabric *(Continued)*

Field	Action
Fabric subnets	<p>Enter fabric CIDR address. The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p> <p>Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.</p>
LR Loopback subnets	<p>Enter an IP subnet to be assigned as loopback interface (lo0) addresses used in Logical Routers (LR). The LR loopback interface IP address is required for eBGP peering to external or unmanaged devices.</p> <p>The field value must include a CIDR with an IP address and a subnet mask. For example, 192.0.20/24.</p>
Loopback subnets (CIDR)	<p>Enter loopback address.</p> <p>Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.</p> <p>If you assign the AR-Replicator and AR-Client roles to enable assisted replication on the QFX10000 devices in a datacenter, you must enter loopback address. For more information, see "Assign a Role to a Device" on page 25.</p>
PNF Servicechain subnets	<p>Enter the IP subnet for allocating IP addresses in the PNF Servicechain subnets field to establish EBGP session between PNF device and SPINE switch.</p> <p>This is an optional field that should be left blank when you are not creating service chains.</p>

Table 21: Provision Existing Fabric (*Continued*)

Field	Action
Advanced interface filters	<p>Create an interface filter to filter the interfaces to include in the fabric. By default, all interfaces identified as participating in Contrail are imported into the fabric during the fabric provisioning process. If an interface filter is set, the fabric provisioning process includes the interfaces that are participating in Contrail and that match the interface filter in the fabric.</p> <p>To create an interface filter, choose the operation as regex and enter the filter characters in the Expression field. The Expression field supports all characters - including metacharacters - allowed in Python regex filters. For example, you can enter <code>^xe</code> in the Expression field to filter out all 10Gbps xe interfaces from the fabric.</p>
Import configured interfaces	<p>Choose this option if configured interfaces need to be imported into the fabric in addition to runtime interfaces. With some exceptions, a configured interface is generally an interface that has been configured in the Junos OS software.</p> <p>A runtime interface is generally an interface that has not been configured in Junos OS. You can confirm which interfaces are configured interfaces by entering the <code>show interfaces</code> command at the configuration mode prompt(<code>#</code>) in Junos. Only runtime interfaces are imported into the fabric by default.</p>


6. Click **Next**.

The Device discovery page is displayed.

The **Device discovery progress** bar on the Device discovery page displays the progress of the device discovery job.

Figure 46: Device Discovery Progress Bar

Device discovery progress



The list of devices discovered are listed in the Discovered devices page.

7. Select the device you want to add by selecting the check box next to the device name.
You can select more than one device.
8. Click **Next** to assign roles.
Assign the Roles page is displayed.
9. From the assign to devices table, select the device you want to assign a role to by selecting the check box next to the device name.
Click the **Assign** icon at the end of the row to assign roles. The Assign role to devices pop-up is displayed.
10. You can now assign physical roles and routing-bridging roles.
 - a. Select a physical role from the Physical Role list.
 - b. Select a routing-bridging role from the Routing Bridging Roles list.

Assigning Roles for Spine Devices:

- Select **spine** from the Physical Role list.
- Select **CRB-Gateway** from the Routing Bridging Roles list.

Assigning Roles for Leaf Devices:

- Select **leaf** from the Physical Role list.
- Select **CRB-Access** from the Routing Bridging Roles list.

Assigning Roles for PNF Devices:

- Select **PNF** from the Physical Role list.
- Select **CRB-Access** and **PNF-Servicechain** from the Routing Bridging Roles list.

NOTE: The number of PNF instances you can create depends on the subnet mask of the pnf-servicechain-subnet that you provided during fabric onboarding. You can create multiple /29 subnets from the pnf-servicechain-subnet.

For example, if a /24 subnet is provided for the pnf-servicechain-subnet, then, you can create $2^5 = 32(29-24=5)$ subnets out of it. Each PNF uses a pair of /29 subnets. Thus, for a /24 subnet, you can have a maximum of 16 PNFs.

Assigning Roles for VNF Devices:

- Select **VNF** from the Physical Role list.
- Select **CRB-Access** from the Routing Bridging Roles list.

NOTE: **ERB-UCAST-Gateway** routing bridging role is also supported.

NOTE: When you configure a QFX series device as a data center gateway, ensure that you assign DC-Gateway role to the spine device.

To assign a DC-Gateway role to a spine device,

- Select **spine** from the Physical Role list.
- Select **DC-Gateway** from the Routing Bridging Role list.

Click **Assign** to confirm selection.

11. Click **Autoconfigure** to initiate the auto-configuration job.

The Autoconfigure page is displayed.

The **Autoconfigure progress** bar on the Discovered devices page displays the progress of the auto-configuration job. Once the auto-configuration job is completed, click **Next**. The Assign Telemetry Profiles page is displayed.

Starting with Contrail Networking Release 2008, you can apply MTU, admin state, flow control, LACP force up, interface type attributes to physical interfaces; and MTU to logical interfaces. These attributes are applied to physical and logical interfaces after you **Autoconfigure** the devices.

To apply these attributes to interfaces:

- Navigate to **Infrastructure > Fabric**.
- Select the desired fabric from the list.
- Select the desired fabric device from the list.
- Click **Physical Interfaces > Create**.
- Enter the required details.

Figure 47: Create Physical Interface

Create Physical Interface

Interface Permissions

Name ^{*}

Type

None ▾

MTU ⓘ

Description

Admin State On

Flow Control ⓘ

- f. Click **Create**.
- g. Click **Logical Interfaces > Create**.
- h. Enter the required details.

Create Logical Interface

Interface Permissions

Name* ⓘ

Connected Physical Interface*

MTU ⓘ

Description

i. Click **Create**.

12. (Optional) Assign telemetry profiles. For more information, see "[Assign Telemetry Profiles](#)" on page 29.

PNF service chain and VNF service chain does not use telemetry profiles.

13. Click **Finish** to exit the Create Fabric wizard.

The onboarding job is now complete.

NOTE: After the devices are onboarded, if you edit the fabric topology by adding new spine or leaf devices or by adding new links between devices, you *must* onboard the edited devices again. If you do not onboard the devices after edits to the initial configuration, underlay formation for the edited devices fails. You can choose to onboard individual devices by clicking the **Onboard**

button for the selected device in the **Fabric Devices** tab of the **Infrastructure > Fabrics > Fabric_Name** page.

Create Virtual Network

A virtual network is a collection of endpoints, such as virtual machine instances, that can communicate with each other. You can also connect virtual networks to your on-premises network. A virtual network in a EVPN VXLAN data center corresponds to a bridge domain for one tenant in a multi-tenant data center fabric.

Follow these steps to create a virtual network from the Contrail Command user interface (UI).

1. Navigate to **Overlay>Virtual Networks**.

The All Networks page is displayed.

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

The Create Virtual Network page is displayed.

3. Enter a name for the network in the **Name** field.

4. Select VN Fabric Type.

Select **Routed** to enable routed virtual network functionality. A routed virtual network represents a layer 3 subnet between the fabric (border gateway) and the third-party physical network device.

For more information, see ["Using Static, eBGP, PIM, and OSPF Protocols to Connect to Third-Party Network Devices" on page 258](#).

Select **Switched** (default option) for tenant virtual network on leaf, bare metal server, or vRouter.

5. Select network policies from the **Network Policies** list. You can select more than one network policy.

Network policies provide connectivity between virtual networks by allowing or denying specified traffic. They define the access control lists to virtual networks. To create a new network policy, navigate to **Overlay>Network Policies**.

For more information on creating network policies, see ["Create Network Policy" on page 92](#).

NOTE: You can attach a network policy to the virtual network after you have created the virtual network.

6. Select any one of the following preferred allocation mode.
 - Flat subnet only

- Flat subnet preferred
- (Default) User defined subnet only
- User defined subnet preferred

An allocation mode indicates how you choose a subnet. You select **Flat subnet only** or **Flat subnet preferred** allocation mode when the subnet is shared by multiple virtual networks. However, you select **(Default) User defined subnet only** or **User defined subnet preferred** allocation mode when you want to define a subnet range.

7. Enter subnet information as given in [Table 22 on page 143](#).

Table 22: Subnet Information

Field	Action
Network IPAM	Select the IP address management method that controls IP address allocation, DNS, and DHCP for the subnet.
CIDR	Enter the overlay subnet CIDR.
Allocation Pools	Enter a list of ranges of IP addresses for vRouter-specific allocation.
Gateway	Enter the gateway IP address of the overlay subnet. This field is disabled by default. To configure this field, uncheck Auto Gateway.
Service Address	Specify the user configured IP address for DNS Service instead of the default system allocated one.
Auto Gateway	This check box is enabled by default and gateway address is allocated by the system. When this box is unchecked, gateway address is user configurable.
DHCP	Select this check box if you want Contrail to provide DHCP service.
DNS	Select this check box if you want the vRouter agent to provide DNS service.

8. Enter host route information.

Host routes are a list of prefixes and next hops that are passed to the virtual machine through DHCP.

- a. **Route Prefix**—Enter a full CIDR value with an IP address and a subnet mask. For example, 10.0.0.0/24.
 - b. **Next Hop**—Enter next hop address.
9. Enter floating IP pool information.
- A floating IP address is an IP address (typically public) that can be dynamically assigned to a running virtual instance. You can configure floating IP address pools in project networks, then allocate floating IP addresses from the pool to virtual machine instances in other virtual networks.
- a. **Pool Name**—Enter pool name.
 - b. **Projects**—Select project from the list.
10. Enter fat flows information. See [Table 23 on page 144](#).
- You can apply fat flows to all VMIs under the configured VN. Fat flows help reduce the number of flows that are handled by Contrail.

Table 23: Configure Fat Flow

Field	Action
Protocol	Select the application protocol.
Port	Enter a value between 0 through 65,535. Enter 0 to ignore both source and destination port numbers. NOTE: If you select ICMP as the protocol, the Port field is not enabled.
Ignore Address	Configure fat flows to support aggregation of multiple flows into a single flow by ignoring source and destination ports or IP addresses. If you select Destination, only the Prefix Aggregation Source fields are enabled. If you select Source, only the Prefix Aggregation Destination fields are enabled. If you select the None (selected by default), both Prefix Aggregation Source and Prefix Aggregation Destination fields are enabled.

Table 23: Configure Fat Flow (Continued)

Field		Action
Prefix Aggregation Source	Source Subnet	<p>Enter the source IP address.</p> <p>Ensure that the source subnet of the flows match. For example, enter 10.1.0.0/24 to create fat flows with 10.1.0.0/24 as the subnet. The valid subnet mask range is /8 through /32.</p> <p>NOTE: For packets from the local virtual machine, source refers to the source IP of the packet. For packets from the physical interface, source refers to the destination IP of the packet.</p>
	Prefix	<p>Enter source subnet prefix length.</p> <p>The prefix length you enter is used to aggregate flows matching the source subnet. For example, when the source subnet is 10.1.0.0/16 and prefix length is 24, the flows matching the source subnet is aggregated to 10.1.x.0/24 flows. The valid the prefix length range is /(subnet mask of the source subnet) through /32.</p>
Prefix Aggregation Destination	Destination Subnet	<p>Enter the destination IP address.</p> <p>Ensure that the destination subnet of the flows match. Enter 10.1.0.0/24 to create fat flows with 10.1.0.0/24 as the subnet. The valid subnet mask range is /8 through /32.</p> <p>NOTE: For packets from the local virtual machine, destination refers to the destination IP of the packet. For packets from the physical interface, destination refers to the source IP of the packet.</p>
	Prefix	<p>Enter the destination subnet prefix length.</p> <p>The prefix length you enter is used to aggregate flows matching the destination subnet. For example, when the source subnet is 10.1.0.0/16 and prefix length is 24, the flows matching the source subnet is aggregated to 10.1.x.0/24 flows. The valid prefix length range is /(subnet mask of the destination subnet) through /32.</p>

11. Enter routing policy and bridge domain information as given below.

- a. Select routing policy from the **Routing Policies** list.

To create a routing policy, navigate to **Overlay>Routing>Routing Policy**.

b. Define a list of route target prefixes.

Enter an IP address in the ASN field and Target in the range 0 through 65,535, or ASN in the range 1 through 65,535 and Target in the range 1 through 4,294,967,295 if 4-byte ASN is disabled. If 4-byte ASN is enabled, enter ASN in the range 1 through 4,294,967,295 and Target in the range 0 through 65,535.

c. Define export route targets.

You can advertise the matched routes from the local virtual routing and forwarding (VRF) table to the MPLS routing table.

Enter an IP address in the ASN field and Target in the range 0 through 65,535, or ASN in the range 1 through 65,535 and Target in the range 1 through 4,294,967,295 if 4-byte ASN is disabled. If 4-byte ASN is enabled, enter ASN in the range 1 through 4,294,967,295 and Target in the range 0 through 65,535.

d. Define import route targets.

Import the matched routes from the MPLS routing table and to the local virtual routing and forwarding (VRF) table.

Enter an IP address in the ASN field and Target in the range 0 through 65,535, or ASN in the range 1 through 65,535 and Target in the range 1 through 4,294,967,295 if 4-byte ASN is disabled. If 4-byte ASN is enabled, enter ASN in the range 1 through 4,294,967,295 and Target in the range 0 through 65,535.

e. Enter bridge domain information. See [Table 24 on page 146](#).

A bridge domain is a set of logical interfaces that share the same flooding or broadcast characteristics.

Table 24: Bridge Domains

Field	Action
Name	Enter a name for the Layer 2 or Layer 3 bridge domain.
I-SID	Enter a Service Identifier in the range from 1 through 16777215.
MAC Learning	<p>Enable or disable MAC learning.</p> <p>MAC learning is the process of obtaining the MAC addresses of all the nodes in a virtual network. It is enabled by default.</p>

Table 24: Bridge Domains (Continued)

Field	Action
MAC Limit	Configure the maximum number of MAC addresses that can be learned.
MAC Move Limit	Configure the maximum number of times a MAC address move occurs in the MAC move time window. A MAC move is when a MAC address appears on a different physical interface or within a different unit of the same physical interface.
Time Window (secs)	Configure the period of time over which the MAC address move occurs. The default period is 10 seconds.
Aging Time (secs)	Configure the MAC table aging time, the maximum time that an entry can remain in the Ethernet Switching table before it is removed. The default time period is 300 seconds.

12. Enter advanced configuration information as given in [Table 25 on page 147](#).

Table 25: Advanced Configuration

Field	Action
Admin State	Select the administrative state of the virtual network.
Reverse Path Forwarding	Enable or disable Reverse Path Forwarding (RPF) check for the virtual network.
Shared	Select to share the virtual network with all tenants.
External	Select the check box to make the virtual networks reachable externally.
Allow Transit	Select to enable the transitive property for route imports.

Table 25: Advanced Configuration (*Continued*)

Field	Action
Mirroring	Select to mark the virtual network as a mirror destination network.
Flood Unknown Unicast	Select to flood the network with packets with unknown unicast MAC address. By default, the packets are dropped.
Multiple Service Chains	Select to allow multiple service chains within two networks in a cluster.
IP Fabric Forwarding	Select to enable fabric based forwarding.
Forwarding Mode	Select the packet forwarding mode for the virtual network.
Extend to Physical Router(s)	Select the physical router to which you want to extend the logical router. The physical router provides routing capability to the logical router.
Static Route(s)	Select the static routes to be added to this virtual network.
QoS	Select the QoS to be used for this forwarding class.
Security Logging Object(s)	Select the security logging object configuration for specifying session logging criteria.
ECMP Hashing Fields	Configure one or more ECMP hashing fields. When configured all traffic destined to that VN will be subject to the customized hash field selection during forwarding over ECMP paths by vRouters.
PBB Encapsulation	Select to enable Provider Backbone Bridging (PBB) EVPN tunneling on the network.

Table 25: Advanced Configuration (*Continued*)

Field	Action
PBB ETree	<p>Select to enable PBB ETree mode on the virtual network which allows L2 communication between two end points connected to the vRouters.</p> <p>When the check box is deselected, end point communication happens through an L3 gateway provisioned in the remote PE site.</p>
Layer2 Control Word	Select to enable adding control word to the Layer 2 encapsulation.
SNAT	Select to provide connectivity to the underlay network by port mapping.
MAC Learning	<p>Enable or disable MAC learning.</p> <p>MAC learning is the process of obtaining the MAC addresses of all the nodes in a virtual network. It is enabled by default.</p>
Provider Network	<p>Select the provider network.</p> <p>The provider network specifies VLAN tag and the physical network name.</p>
IGMP enable	Enable or disable IGMP.
Multicast Policies	<p>Select the multicast policies.</p> <p>To create a policy, navigate to Overlay>Multicast Policies.</p>
Max Flows	Enter the maximum number of flows permitted on each virtual machine interface of the virtual network.

13. Click Create.

The All Networks page is displayed. The virtual network that you created is displayed on this page.

Configuring Virtual Port Groups

This topic describes how to create virtual port groups (VPGs) from Contrail Command UI. Contrail Networking Release 2008 introduces a redesigned VPG-creation workflow. To create a VPG, perform the steps described in ["No Link Title" on page 150](#) if you are using release 2008 later and those described in ["No Link Title" on page 155](#) if you are using releases 2003 and 2005.

- **For release 2008:**

In Contrail Networking Release 2008, you can create a VPG without attaching VLANs. You have the ability to add VLANs after the VPG is created. In scaled setups, there can be a large number of VLANs, making it very hard to manage inside the create or edit Virtual Port Group pages. Release 2008 simplifies the assignment of VLANs by introducing a dedicated page for management. The VPG creation workflow comprises two steps with the first step being configuration of the VPG. Only when the configuration step is completed successfully can you assign the VLANs which is the second step.

To create virtual port groups in Contrail Command in release 2008:

1. Navigate to **Overlay > Virtual Port Group > Create Virtual Port Group**.

The **New Virtual Port Group** wizard is displayed.

2. Enter a name for the virtual port group in the **Virtual Port Group Name** field.
3. Select the fabric from the **Fabric Name** list.

The available physical interfaces on the devices in the selected fabric are listed.

4. From the **Available Physical Interface** box, select the physical interfaces to be included in the virtual port group by clicking the arrow next to each physical interface. The available physical interfaces are the interfaces available on TORs that are already onboarded.

The selected interfaces are displayed in the **Assigned Physical Interface** box.

If you select more than one interface on the same TOR as shown in [Figure 52 on page 156](#), a link aggregation group (LAG) is automatically created on the device.

5. Select a security group from the **Security Groups** list.

For enterprise style fabric configuration, attach a security group to the virtual port group. The policies defined in the security group is assigned to all the ports in the virtual port group. For service provider style fabric configuration, you can attach a security group to every VLAN.

6. Assign a port profile to the virtual port group by selecting a port profile from the **Port Profile** list.

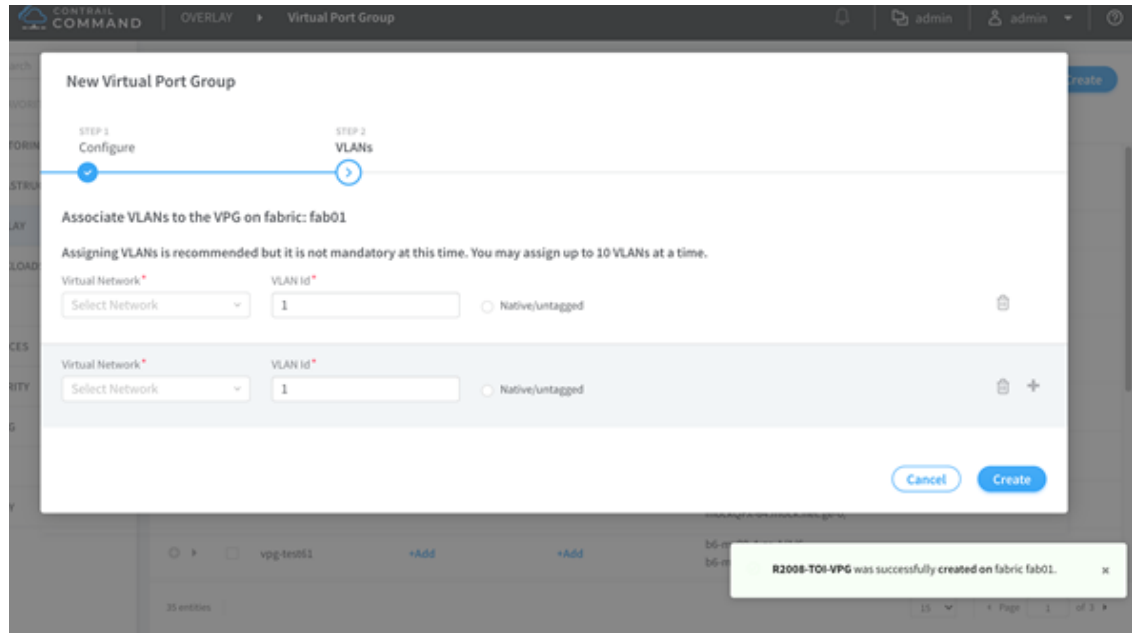
A port profile functions like a container that can support multiple port-related configurations, and allows you to apply those configurations by attaching them to the port profile.

7. Click **Next** to create the VPG. If VPG creation fails, an error message is displayed. If VPG creation is successful, you will be directed to the second step in the process, in which you can add the VLANs.
8. (Optional) You can assign VLANs in this step of the wizard. You can also add VLANs in the **Overlay > Virtual Port Group** page (see "10" on page 153). To add VLANs here, enter the information as shown in Table 26 on page 151.

Table 26: Enter VLAN Information

Field	Action
Virtual Network	Select the virtual network to which the virtual port group belongs.
VLAN ID	Enter the VLAN ID and network to which the VLAN is associated. If you enable the VLAN-ID Fabric-Wide Significance option when creating a fabric, you can associate one VLAN ID to only one virtual network. This ensures that the same VLAN ID is not associated with more than one virtual network within the same enterprise style fabric.
Native/untagged	Select this check box to allow a native/untagged virtual network (optional). You can assign only one native/untagged VLAN in a virtual port group.
Security Group	This field is available only in service provider style fabric configuration. Select a security group from the Security Groups list. You can attach a security group to each VLAN.

Figure 48: Assign VLANs



9. Click **Create**.

The newly created virtual port group is displayed in the Virtual Port Group page with details of the interfaces as shown in [Figure 49 on page 153](#).

Figure 49: Virtual Port Groups

NAME	VLANs	PHYSICAL INTERFACES	PORT PROFILE
PavelTest01	0 23	b6-mx80-4-re-0/0/0 b6-mx80-4-ge-1/2/8 + 1 more	
vpg-internal-0	1	b6-mx80-4-ae0	
vpg-test70	150 0 + 1 more +Add	mockQFX-64.mock.net-ge-0	
tst-tst	+Add	b6-mx80-4-ge-1/3/5 b6-mx80-4-ge-1/3/2 + 8 more	
tst-13454	+Add	b6-mx80-4-ge-1/3/3 b6-mx80-4-ge-1/2/9	
vpg-test64	+Add	mockQFX-64.mock.net-ge-0 mockQFX-64.mock.net-ge-0	
vpg-test62	+Add	mockQFX-64.mock.net-ge-0 mockQFX-64.mock.net-ge-0	
vpg-test61	+Add	b6-mx80-4-ge-1/1/6 b6-mx80-4-ge-1/0/10	

10. (Optional) To assign VLANs if not previously configured or to edit configured VLANs, perform one of the following steps.

- To edit or add only VLANs, click a VLAN or click **Add** next to the VPG name. The VLANs assignment page is displayed.
- To edit VPG information and/or edit VLANs, select a VPG and click the edit (pencil) icon. The **Edit VPG** page is displayed.

Edit the VPG information as required. Click **Save** to save the changes and remain on this page. Alternatively, click **Save and assign new VLANs** to save the changes and assign VLANs. The VLANs assignment page is displayed.

Figure 50: Edit VPG

Virtual Port Group name* Fabric name* Security Groups Port Profile

All changes made here will be committed to the Controller.

Available Physical Interface

DISPLAY NAME	PHYSICAL ROUTER
ge-1/2/1	b6-mx80-4
ge-1/1/11	b6-mx80-4
ge-1/1/1	b6-mx80-4
ge-1/2/3	b6-mx80-4

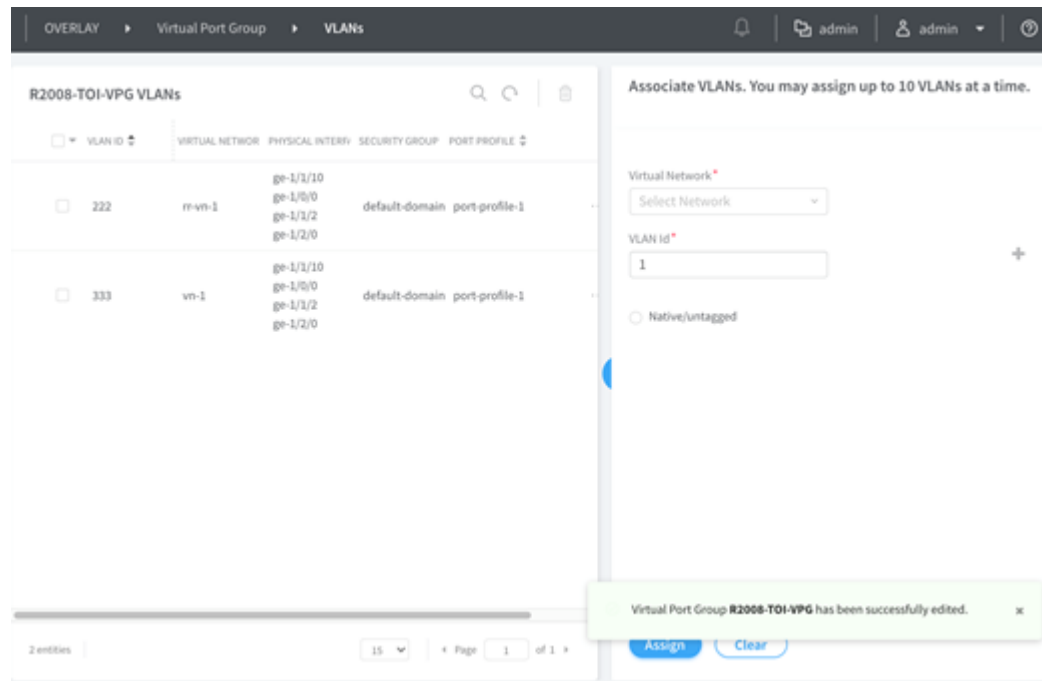
Assigned Physical Interface

DISPLAY NAME	PHYSICAL ROUTER
ge-1/1/10	b6-mx80-4
ge-1/2/0	b6-mx80-4
ge-1/0/0	b6-mx80-4
ge-1/1/2	b6-mx80-4

Buttons: Cancel, Save and assign new VLANs, Save

- The VLANs assignment page has two panels. The left panel lists all currently configured VLANs, if any. The right panel enables you to assign additional VLANs. Enter VLAN information and click **Assign** to attach the VLANs. The VLANs appear in the left panel. You can attach up to 10 VLANs at a time. You can also edit existing VLANs from this page. Successful and failed attempts at assigning and editing are indicated through success or error message pop-ups.

Figure 51: Edit VLANs



For better visibility, you can hide the right panel by clicking the blue expansion icon. You can also use this page to delete individual VLANs and bulk delete multiple VLANs.

- **For releases 2003 and 2005:**

To create virtual port groups in Contrail Command using releases 2003 and 2005:

1. Navigate to **Overlay > Virtual Port Group > Create Virtual Port Group**.

The Create Virtual Port Group page is displayed.

2. Enter a name for the virtual port group in the **Virtual Port Group Name** field.

3. Select virtual port group type.

With Contrail Networking Release 2003, you can create a routed virtual port group from the Contrail Command UI. Select the **Routed** option button to create a routed virtual port group. Select **Layer 2** option button to create a virtual port group.

4. Select the fabric from the **Fabric Name** list.


The available physical interfaces on the devices in the selected fabric are listed.

5. From the **Available Physical Interface** box, select the physical interfaces to be included in the virtual port group by clicking the arrow next to each physical interface. The available physical interfaces are the interfaces available on TORs that are already onboarded.

The selected interfaces are displayed in the **Assigned Physical Interface** box.

If you select more than one interface on the same TOR as shown in [Figure 52 on page 156](#), a link aggregation group (LAG) is automatically created on the device.

Figure 52: Select Interfaces on the Same TOR



- Assign a security group to the virtual port group by selecting a security group from the **Security Groups** list.

The policies defined in the security group is assigned to all the ports in the virtual port group.

- Select and assign a port profile from the **Port Profile** list.

A port profile functions like a container that can support multiple port-related configurations, and allows you to apply those configurations by attaching them to the port profile.

- Enter the following information as given in [Table 27 on page 156](#).

Table 27: Enter VLAN Information

Field	Action
Network	Select the virtual network to which the virtual port group belongs.
VLAN ID	Enter the VLAN ID and network to which the VLAN is associated. If you enable the VLAN-ID Fabric-Wide Significance option when creating a fabric, you can associate one VLAN ID to only one virtual network. This ensures that the same VLAN ID is not associated with more than one virtual network within the same enterprise style fabric.

Table 27: Enter VLAN Information (Continued)

Field	Action
Display Name	Enter the VLAN name. If the Auto Display Name field is selected, this field is autogenerated from the virtual port group name.
Auto Display Name	Select Auto Display Name if you want the VLAN name to be autogenerated from the virtual port group name.
Native/untagged	Select this check box to allow a native/untagged virtual network (optional). You can assign only one native/untagged VLAN in a virtual port group.

9. Click **Create**.

The newly created virtual port group is displayed on the Virtual Port Group page with details of the interfaces and the TORs as shown in [Figure 53 on page 157](#).

Figure 53: Virtual Port Groups

NAME	VLAN IDS	TOR PORT VLAN IDS	PHYSICAL INTERFACES	VIRTUAL NETWORK
vpg-internal-0		4094	ge-0/0/9:contrail-qfx5110-6	right_vn_1
vpg-test		4094	fxp0:contrail-srx5600-2 xe-0/0/32:2:bng-contrail-qfx-1... 1 more	left_vn_13

No items selected

You can delete a virtual port group by clicking the delete icon against the virtual port group. To delete a virtual port group, you must first remove the referenced VMI and the associated BMS instance from the virtual port group.

SEE ALSO

[Virtual Port Groups | 236](#)

Create Logical Routers

A logical router replicates the functions of a physical router. It connects multiple virtual networks. A logical router performs a set of tasks that can be handled by a physical router, and contains multiple routing instances and routing tables.

Follow these steps to create a logical router (LR).

1. Navigate to **Overlay > Logical Routers** and click **Create**.

The Create Logical Routers page is displayed.

2. Enter the following information as given in [Table 28 on page 158](#).

Table 28: Create a Logical Router

Field	Action
Name	Enter a name for the Logical Router.
Admin State	Select the administrative state that you want the device to be in when the router is activated. Up is selected by default.
Logical Router Type	Select SNAT Routing or VXLAN Routing from the list.
Choose Fabric	Select the fabric that you are associating this logical router to.
Connected Networks	Select the networks that you want to connect this logical router to.

Table 28: Create a Logical Router (*Continued*)

Field	Action
Extend to Physical Router	<p>Select the physical router(s) to which you want to extend virtual networks or routed virtual networks to, from the Extend to Physical Router list.</p> <p>A physical router provides routing capability to the logical router.</p>
Reconfigure Physical Routers	<p>This link is enabled when you select a routed virtual network from the Connected networks list. Click Reconfigure Physical Router to reconfigure a physical router that you want to extend a virtual network to.</p> <p>For more information, refer to the Create Logical Routers section of the "Using Static, eBGP, PIM, and OSPF Protocols to Connect to Third-Party Network Devices" on page 258 topic.</p>
Public Logical Router	(Optional) Select this check box if you want the logical router to function as a public logical router.
NAT	<p>Select this check box to enable Network Address Translation (NAT).</p> <p>This check box is disabled by default.</p>
VxLAN Network Identifier	<p>Enter VXLAN network identifier in the range from 1 through 16,777,215.</p> <p>This field is disabled by default.</p>
DHCP IP Address	<p>Enter DHCP relay server IP address.</p> <p>You can add more than one IP address. To add another address, click +Add.</p>

Table 28: Create a Logical Router (*Continued*)

Field	Action
Route Target(s)	<p>Click +Add to add route targets.</p> <p>Enter Autonomous System (AS) number in the ASN field.</p> <ul style="list-style-type: none"> • Enter ASN in the range of 1-4,294,967,295, when 4 Byte ASN is enabled in Global Config. • Enter ASN in the range of 1-65,535, when 4 Byte ASN is disabled. • You can also add suffix <i>L</i> or <i>l</i> (<i>lower-case L</i>) at the end of a value in the ASN field to assign an AS number in 4-byte range. Even if the value provided in the ASN field is in the range of 1-65,535, adding <i>L</i> or <i>l</i> (<i>lower-case L</i>) at the end of the value assigns the AS number in 4-byte range. If you assign the ASN field a value in the 4-byte range, you must enter a value in the range of 0-65,535 in the Target field. <p>Enter route target in the Target field.</p> <ul style="list-style-type: none"> • Enter route target in the range of 0-65,535, when 4 Byte ASN is enabled and ASN field is assigned a 4-byte value. • Enter route target in the range of 0-4,294,967,295, when the ASN field is assigned a 2-byte value.

3. Click **Create** to create the logical router.

The Logical Routers page is displayed.

NOTE: The router_interface object (Virtual Port) is created as part of the LR creation and VN extension to Spines workflow. While planning the IP address for spines, you must be aware that an extra one IP address is required for the router_interface object which gets created automatically.

Configure the Internal Virtual Networks

Use this procedure to configure the internal virtual networks.

When you connect a logical router to a virtual network, Contrail Networking automatically creates internal virtual networks. For logical routers named LR1 and LR2, the internal virtual networks are called LR::LR1 and LR::LR2 respectively. These networks attach to the service virtual machine.

1. Select **Overlay>Virtual Networks** to bring up the list of virtual networks.
2. Hover over the internal virtual network you want to configure (for example, LR::LR1) and click the Edit icon on far right of the row.

The Edit Virtual Network page appears.

3. In the Subnets section, click **+Add**.
4. Use the drop-down list to select the **Network IPAM** you want to use.
5. Specify the subnet in **CIDR** format (for example, 10.192.10.0/24).

The **Gateway** and **Service Address** are automatically filled in based on the subnet you configured. You are free to change these addresses although there's generally no need for you to do so.

6. Click **Save**.
7. Repeat these steps to configure the other internally-generated virtual network, but make sure to specify a different subnet.

Create the Service Virtual Machine

Use this procedure to create the service virtual machine, which is simply a compute workload.

1. Select **Workloads>Instances** to bring up the Instances page.
2. Click **Create** to create an instance.
3. Select **Virtual Machine** as the Server Type.
4. Specify the **Instance Name**.
5. Specify **Image** as the Boot Source.
6. Use the drop-down lists to select the **Image** and **Flavor**, which describe the image you want the VM to run and the compute specifications for the VM.

The drop-down lists are populated with the images and flavors you create through **Workloads>Images** and **Workloads>Flavors**.

7. Attach the VM to the internally-generated virtual networks LR::LR1 and LR::LR2 by using the arrows to move them from the **Available Networks** section to the **Allocated Networks** section.
8. Click **Create**.

Create VNF Service Template

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 a name for the 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 a service mode from the **Service Mode** list.

7. Select a service type from the **Service Type** list.

8. Add the left, right, and management interfaces in the Interface section.

- Select **left** as the interface type from the **Interface Type** list.
- Click **+ Add** and select **right** 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.

Create VNF Service Instance

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 a name for the service instance in the **Name** field.

4. Select the service template that you created from the **Service Template** list.

The **Interface Type** and **Virtual Network** fields are displayed for each interface.

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

- **left**—Select the left virtual network that you created.

- **right**—Select the right virtual network that you created.
6. Associate this service instance to the VNF you created earlier.
 - a. Expand the **Port Tuples** section.
 - b. Click **+Add**.
 - c. Use the drop-down list to specify the **Virtual Machine Interface** for the left and right interfaces.
 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 the Network Policy

Use this procedure to create the network policy that governs traffic going through the VNF.

1. Select **Overlay>Network Policies** to bring up the Network Policies page.
2. Click **Create**.
3. Provide a **Policy Name**.
4. In the Policy Rule(s) section, select **Network** as the **Source Type** and use the drop-down lists to specify the **Source** (for example, LR::LR1) and **Destination** (for example, LR::LR2) networks.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2008	Starting with Contrail Networking Release 2008, you can apply MTU, admin state, flow control, LACP force up, interface type attributes to physical interfaces; and MTU to logical interfaces.
2008	In Contrail Networking Release 2008, you can create a VPG without attaching VLANs. You have the ability to add VLANs after the VPG is created.
2003	With Contrail Networking Release 2003, you can create a routed virtual port group from the Contrail Command UI. Select the Routed option button to create a routed virtual port group.
1912	Contrail Networking Release 1912 extends the service chaining functionality to bare metal servers (BMS).

RELATED DOCUMENTATION

Creating Layer 3 PNF Service Chains for Inter-LR Traffic | 121

Retaining the AS Path Attribute in a Service Chain

Service chaining allows two virtual networks to communicate with each other using a service policy or network policy. The VNs communicate through services instances defined in the network policy. Service instances can be physical network functions or virtual network functions.

For data to traverse between VNs, the BGP route attributes are modified according to the network policy. One such BGP attribute, is the AS path attribute. AS path is a sequence of autonomous systems that network packets traverse. By default, the AS path is nullified while leaking routes from the source to the destination network in a service chain. Starting with Contrail Networking Release 2011, you can configure the AS path to be retained in the routes re-originated from the destination VN to the source VN in a service chain. You also have the ability to enable or disable the path retention for selected service chains.

NOTE: The AS path retention feature works only for virtual network functions.

You can enable or disable the **Retain AS Path** option while configuring the network policy. A network policy is unique to a service chain and configuring the knob at the policy level will apply that feature to that service chain and its component service instances defined by the policy. Even when service instances are shared between multiple service chains, the same service instance can behave differently in different service chains based on the **Retain AS Path** knob in the policy configuration.

To configure the AS Path attribute to be retained in the routes re-originated from the destination VN to the source VN in a service chain:

1. Navigate to **Overlay > Network Policies > Create Network Policy**.
The **Network Policy** tab of the **Create Network Policy** page is displayed.
2. Enter a name for the policy in the **Policy Name** field.
3. Edit the fields in the **Policy Rule(s)** section as per your policy requirement.
4. Click **Advanced Options** and edit the fields displayed as per your policy requirement.
5. Select the **Retain AS Path** check box if you want to retain the AS Path between the source and destination virtual networks. The check box is disabled, by default.
6. Click **Create** to create the network policy.

The **Network Policies** page is displayed. You can now attach the network policies to the required VNs.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2011	Starting with Contrail Networking Release 2011, you can configure the AS path to be retained in the routes re-originated from the destination VN to the source VN in a service chain. You also have the ability to enable or disable the path retention for selected service chains.

RELATED DOCUMENTATION

| [Creating VNF Service Chains for Inter-LR Traffic](#) | 128

Assisted Replication of Broadcast, Unknown Unicast, and Multicast Traffic

IN THIS SECTION

- [Benefits of Assisted Replication](#) | 167

Starting with Contrail Networking Release 1907, you can configure assisted replication on datacenter devices and assign the AR-Replicator and AR-Client roles to them.

Assisted replication or assisted ingress replication is a method to transport ingress broadcast, unknown unicast, and multicast (BUM) traffic in a more efficient way. In assisted replication, you configure a datacenter device as a dedicated replicator, which receives BUM traffic from the network virtualization edge (NVE) devices or provider edge (PE) devices.

An AR-Replicator is a network virtualization overlay (NVO) device or a provider edge device that replicates BUM traffic received through an overlay tunnel to other overlay tunnels and local attachment circuits. An AR-Client is a device that supports assisted replication and sends BUM traffic only to AR-Replicator.

You can designate powerful spines in the datacenter as replicators, which can receive the BUM traffic from ToRs and replicate them to the PEs in the network. To enable assisted replication, you can

configure the AR-Client role to MX Series, QFX10000 and QFX5000 devices as spine or leaf, and the AR-Replicator role to QFX10000 devices as spine or leaf.

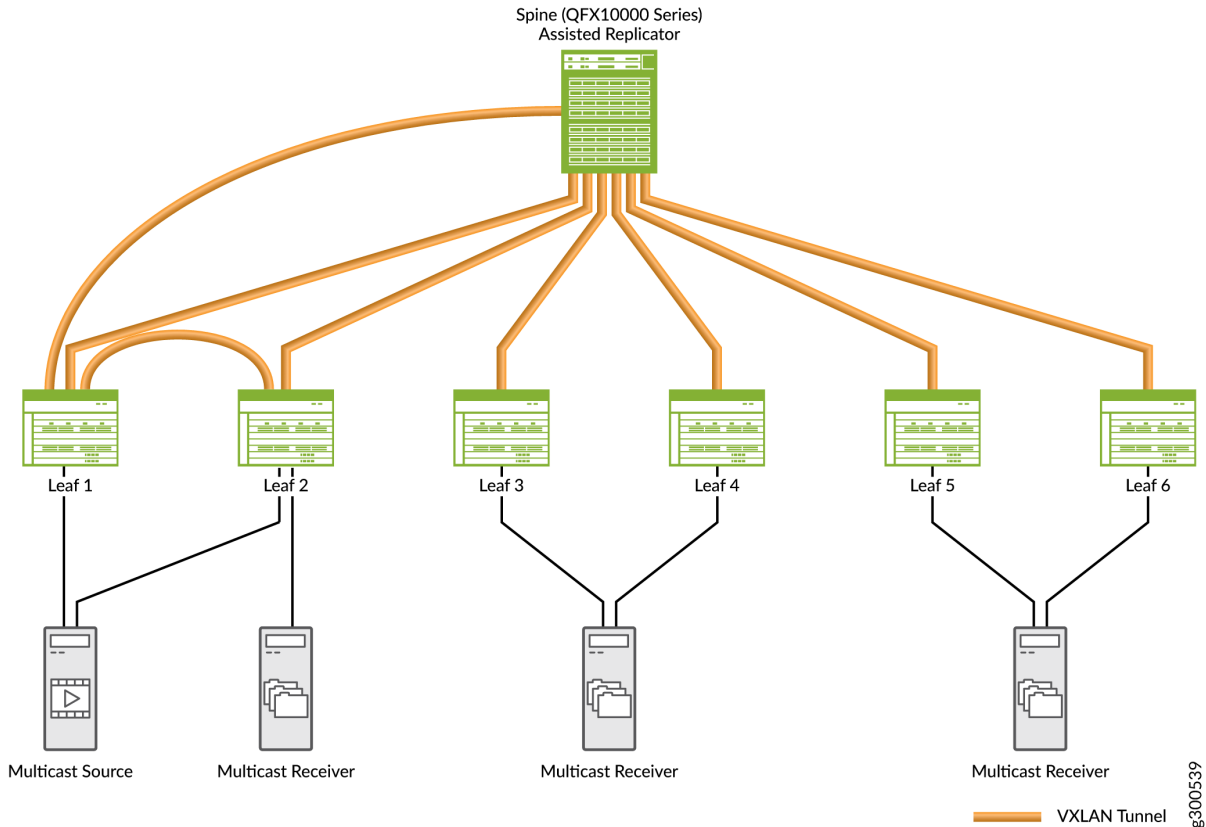
Assisted replication feature optimizes replication of ingress BUM traffic received from the CE interfaces by directing the BUM traffic towards a single EVPN core Replicator PE (a QFX10000 device) rather than sending it to all the PE devices for replication. Configuring a dedicated replicator for BUM traffic reduces the load on the PEs and improves the performance and efficiency of the network in transporting BUM traffic. This in turn leads to better utilization of the bandwidth.

NOTE: To configure AR-Client and AR-Replicator roles, the QFX10000 and QFX5000 series devices must be running Junos OS Release 18.4 R2 or later.

For a complete list of devices and the supported Junos OS and Contrail Networking releases, see ["Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles"](#) on page 185.

[Figure 54 on page 167](#) shows how Leaf 1 node sends traffic received from a multicast source to the assisted replicator through a VXLAN tunnel and how the assisted replicator replicates traffic to the leaf nodes. Here, the task of replicating BUM traffic to other PEs is shifted from the leaf to the spine, which functions as a dedicated replicator. Leaf 1, which receives the multicast traffic from the multicast source, replicates it to the spine as well as to the Leaf 2 through a VXLAN tunnel. The assisted replicator (spine) does not send the multicast traffic back to the leaf which the multicast source is connected to.

Figure 54: Traffic Path in Assisted Replication



Assisted replication is similar to conventional network segmentation. While the segments in conventional network segmentation could be in different regions, the segments in assisted replication exist in the same region.

You can configure assisted replication and assign roles to devices in a datacenter from the **Infrastructure > Fabrics > Create Fabric** page in Contrail Command. See ["Assign a Role to a Device"](#) on page 25 for step-by-step procedure to assign roles to a device.

Benefits of Assisted Replication

- Improves the performance and efficiency of the network by replicating BUM traffic towards a single EVPN core Replicator PE instead of sending to all PE devices.
- Reduces the load on the PEs, which in turn improves bandwidth utilization in the network.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1907	Starting with Contrail Networking Release 1907, you can configure assisted replication on datacenter devices and assign the AR-Replicator and AR-Client roles to them.

RELATED DOCUMENTATION

[Fabric Overview | 4](#)

[Create a Fabric | 7](#)

Running Generic Device Operations Commands In Contrail Command

Contrail Networking enables you to obtain device information, such as interface information, like input rate or output rate, or search for the name of an interface by providing its MAC address or IP address from the Contrail Command UI. You can run a specific generic device operations command on multiple devices at a time. A job template is defined for each generic device operations command. After you select the devices and specify the parameters defined in the job template, a job is created depending on the generic command you selected. The result of the job is then displayed for the selected device or devices.

You can select a maximum of 20 devices at a time and run a generic device operations command to view information about those devices.

You can run the following generic device operations commands:

- **Search using MAC or IP address**—Use this generic device operations command to identify the interface name if you know the IP address or the MAC address of an interface. This operation is useful to locate the interface by specifying the interface name and information such as name of the originating device and its loopback IP address.
- **Show MAC mobility**—Use this generic device operations command to display the current location of a MAC address within the fabric and its local and remote origin.
- **Show chassis information**—Use this generic device operations command to view a range of chassis-related information such as chassis environment, routing engine, chassis environment and so on.

- **Show operations information**—Use this generic device operations command to display a range of operational information such as BGP configuration information, EVPN configuration information, VLAN information and so on.
- **Show current or rollback configuration**—Use this generic device operations command to display the rollback configuration.
- **Show interface details**—Use this generic device operations command to display information about the physical or virtual interfaces on a fabric node.
 - **Show interfaces**—Use this generic device operations command to show a list of all runtime interfaces. You can use the filters to select the type of interface, such as physical or logical. You can also view particular types of interfaces using the `regex` filter.
 - **Show configured interfaces**—Use this generic device operations command to list all the configured interfaces. You can use the filters to select the type of interface, such as physical or logical. You can also view particular types of interfaces using the `regex` filter..
 - **Show interfaces by names**—Use this generic device operations command to check whether a particular type of interface is present in one or more of the devices selected. This operation is useful when you want to check which among the selected devices has an `xe-0/0/2` interface or an `100.0` interface. You can use the filters to select the type of interface, such as physical or logical. You can then enter the interface name you want to search for.

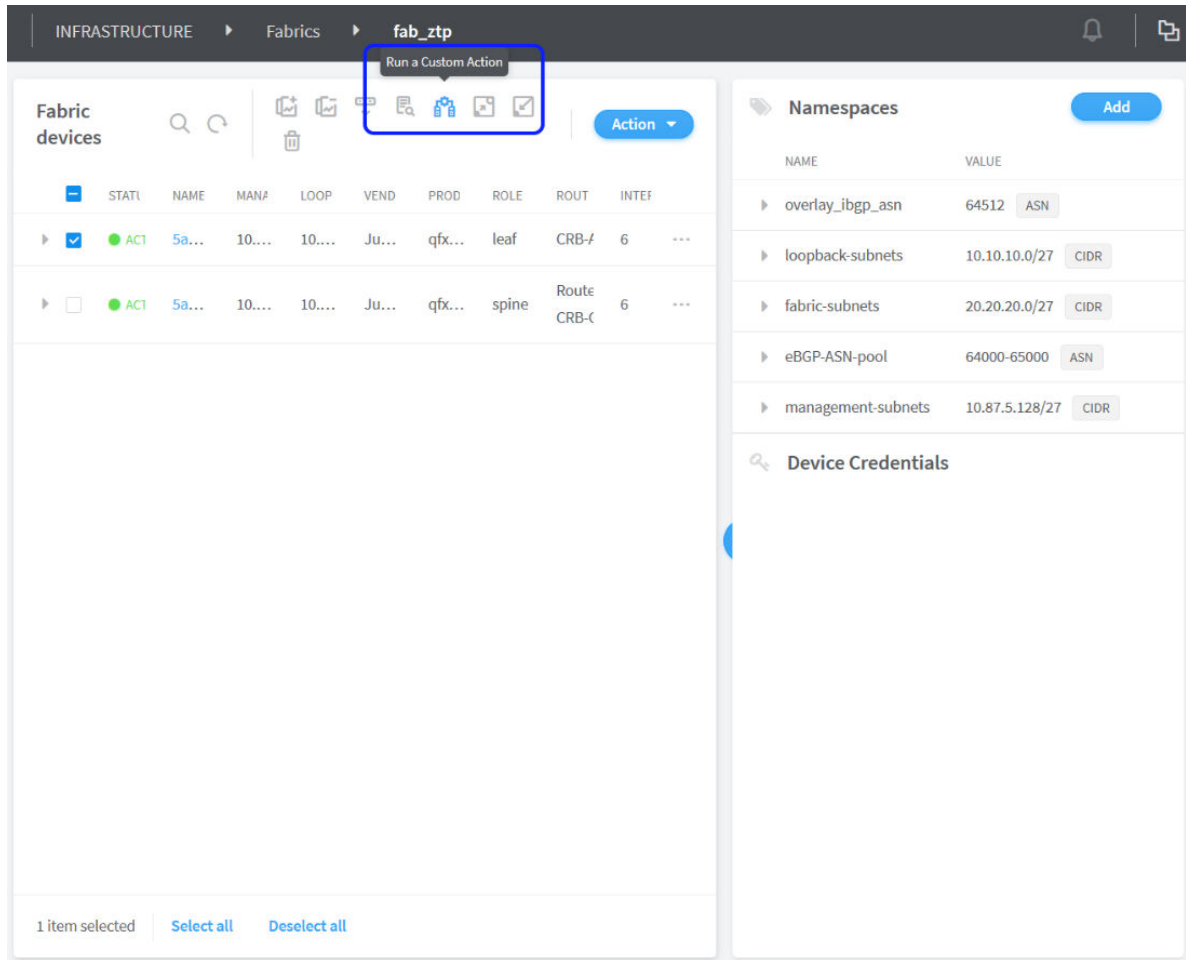
You can create a custom generic device operations command by adding a `job_template` object type in the `opt/contrail/fabric_ansible_playbooks/conf/predef_payloads.json` file. Follow these best practices when you define a new generic device operations command.

- Make sure that `template_type` is set to `device_operation`, which identifies this template as a generic device operation job template.
- Create a new `job_template` for every generic device command you need to execute. Specify the command name in the `job_template_name` field so that it is easy to identify the command.
- Make sure that the generic device operation `job_templates` references to the playbook `/opt/contrail/fabric_ansible_playbooks/operational_command.yml`.
- Any change to the `predef_payloads.json` requires a restart of the `config_api_1_xxxx` docker.

To run a generic device operations command:

1. Navigate to **Infrastructure > Fabrics > *fabric name***.
2. Select the fabric devices and click the **Run a Custom Action** button as shown in [Figure 55 on page 170](#).

Figure 55: Select fabric Devices



3. Click the operation that you want to perform and click **Next**.

Figure 56: Choose an Operation

The screenshot displays the 'Device Operation' interface with a breadcrumb trail: INFRASTRUCTURE > Fabrics > fab_ztp > Device Operation. The interface is divided into three steps: STEP 1 Operation (active), STEP 2 Parameters, and STEP 3 Operation Status. A progress bar at the top indicates the current step.

The 'Choose an operation' panel on the left lists several options, with 'Show Interface Details' highlighted in blue. The other options are: Show chassis information, Show mac mobility, Test overlay connectivity, Show current or rollback configuration, Check incoming multicast traffic, Show operations information, and Search using IP or MAC.

The 'Devices chosen for Show Interface Details' panel on the right displays a table with the following data:

DEVICE	MANAGEMENT IP	VENDOR NAME	PRODUCT NAME
5a12-qf9	10.XX.XX.XX	Juniper	qfx5100-48s-6q

At the bottom right of the interface, there are 'Cancel' and 'Next' buttons.

4. Select a **Sub Operation Type** and enter the filters.

Details of the devices selected are displayed as shown in [Figure 57 on page 172](#).

Figure 57: Devices Selected for the Operation

INFRASTRUCTURE ▸ Fabrics ▸ fab_ztp ▸ Device Operation

STEP 1 Operation STEP 2 Parameters STEP 3 Operation Status

Sub Operation Type
Choose a sub-operation*
 Show runtime interfaces ^
 Show configured interfaces
 Show interfaces by names
 Show runtime Interfaces

Filter*
 Filter Expression:
 Filter Type:
 + Add

Devices chosen for Show Interface Details

DEVICE	MANAGEMENT IP	VENDOR NAME	PRODUCT NAME
5a12-qfx9	10.XX.XX.XX	Juniper	qfx5100-48s-6q

Interface Details

Previous Cancel Execute

5. Click **Execute**.

Information about the selected device is displayed as shown in [Figure 58 on page 173](#).

Figure 58: Generic Device Operation Command Results

STEP 1 Operation STEP 2 Parameters STEP 3 Operation Status

Operation is being run on the following devices

DEVICE	STATUS
bng-contrail-qfx-1k2	●
bng-contrail-qfx-1k2-4	●
contrail-qfx5110-1	●

Results for bng-contrail-qfx-1k2

ADDRESS FAMILY	INTERFACE	OPERATIONAL STATUS	ADMIN STATUS	INPUTRATE (BPS)	OUTPUTRATE (BPS)
	xe-0/0/32:0				
	xe-0/0/32:1				
	xe-0/0/32:2				
	xe-0/0/32:3				
	xe-0/0/34:0				
	xe-0/0/35:3				
	ae1				
	lo0				
	lo0				
	irb				

Previous Cancel Finish

6. Click **Finish** to complete the operation.

Adding DHCP Server Information for Virtual Networks and Logical Routers

IN THIS SECTION

- [Topology | 174](#)
- [Steps to Add DHCP Server Information | 176](#)
- [Steps to Remove CSN Information | 178](#)

In a Contrail-automated multi-tenant data center EVPN or VXLAN fabric, the tenant administrator needs to ensure that all departments use corporate Dynamic Host Configuration Protocol (DHCP) servers for

endpoint IP and workload IP address assignment. Starting in Contrail Networking Release 1908, tenant administrators can define a set of DHCP server IP addresses while configuring virtual networks and logical routers on a multi-tenant data center fabric. After DHCP relay in each virtual network and logical router, Contrail Networking configures these defined IP addresses on the IP fabric.

In earlier releases, a Contrail services node (CSN) is used to provide DHCP and Domain Name System (DNS) services to bare metal servers. With Contrail Networking Release 1908, you can directly add DHCP server information by adding the server IP address in the **Overlay > Logical Router** page of the Contrail Command user interface (UI). The DHCP server that you use must be located in the same virtual network as that of the bare metal server or reachable through the Internet (inet.0).

Contrail Networking does not support the use of a DHCP server and a CSN at the same time. When you use a DHCP server, you must not provision a CSN and must remove existing CSNs. However, when you provision a CSN again, ensure that you remove DHCP server information and reprovision the bare metal server to enable the CSN to manage IP addresses.

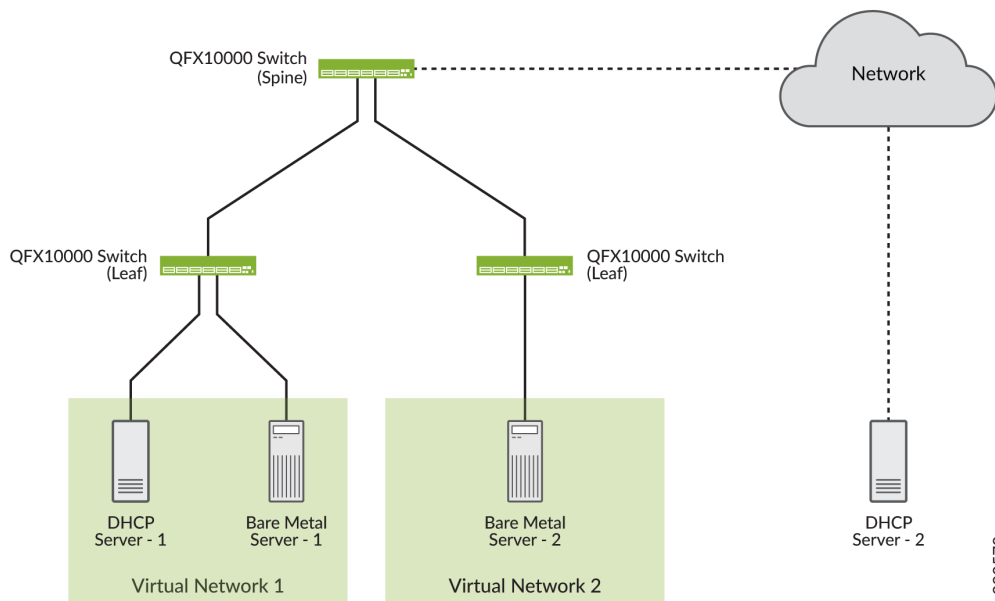
NOTE: This feature is supported only on QFX5000 and QFX10000 series devices running Junos OS Release 18.4R2 or later.

Topology

Consider the following scenarios as shown in [Figure 59 on page 175](#):

- The DHCP server (DHCP Server - 1) and bare metal server (Bare Metal Server - 1) are located in the same network (Virtual Network - 1)
- The DHCP server (DHCP Server - 2) and bare metal server (Bare Metal Server - 2) are not located in the same network

Figure 59: Sample Topology



Depending on whether the DHCP server is located in the same virtual network or connected remotely to a routed network through the underlay, the following configuration is pushed to the leaf switch or spine switch.

- If the DHCP server is located in the same virtual network, the following configuration is pushed:

```

set routing-instances <vrfname> forwarding-options dhcp-relay forward-only
set routing-instances <vrfname> forwarding-options dhcp-relay forward-only-replies
set routing-instances <vrfname> forwarding-options dhcp-relay server-group
DHCP_SERVER_GROUP <dhcp-server-ip>
set routing-instances <vrfname> forwarding-options dhcp-relay active-server-group
DHCP_SERVER_GROUP
set routing-instances <vrfname> forwarding-options dhcp-relay group RELAY_DHCP_SERVER_GROUP
interface <irb>
set routing-instances <vrfname> forwarding-options dhcp-relay overrides relay-source lo0

```

- If the DHCP server is not located in the same virtual network, the following (additional) route configuration is pushed:

```

set routing-instances <vrfname> routing-options static route <dhcp-server-ip> next-table
inet.0

```

```

set routing-instances <vrfname> routing-options interface-routes rib-group inet <ribgroup>
set routing-options rib-groups <ribgroup> import-rib <vrfname>.inet.0
set routing-options rib-groups <ribgroup> import-rib inet.0
set forwarding-options dhcp-relay forward-only-replies

```

Steps to Add DHCP Server Information

IN THIS SECTION

- [Adding DHCP Server Information to an Existing Logical Router | 176](#)
- [Adding DHCP Server Information while Creating a Logical Router | 177](#)

DHCP server information can be added to an existing logical router if the DHCP server is located in the same network as that of the virtual network and the logical router. You can edit the logical router and add the server information in the **Overlay > Logical Router** page of the Contrail Command UI.

You can also create a new logical router by using the Contrail Command UI. You can add DHCP information and associate virtual networks from the **Create Logical Router** page.

These topics provide information to add DHCP server information by using the Contrail Command UI.

NOTE: Ensure that you remove existing CSNs before you provision the DHCP server.
For more information, see "[Steps to Remove CSN Information](#)" on page 178.

Adding DHCP Server Information to an Existing Logical Router

Follow these steps to add DHCP server information to an existing logical router by using the Contrail Command UI:

1. Click **Overlay > Logical Routers**.
The Logical Routers page is displayed.
2. Select the router you want to edit by selecting the check box next to the name of the logical router, and click the **Edit** icon.
The Edit Logical Router page is displayed.
3. From the DHCP Relay Server(s) section, click **+Add**.

The **DHCP IP Address** field is displayed as seen in [Figure 60 on page 177](#).

Figure 60: DHCP IP Address Field

4. Enter the IP address of the DHCP server in the **DHCP IP Address** field.

NOTE: Ensure that you remove existing CSNs before you provision the DHCP server. For more information, see ["Steps to Remove CSN Information" on page 178](#).

5. Click **Save**.

The DHCP server IP address is now added to the logical router.

Adding DHCP Server Information while Creating a Logical Router

Follow these steps to add DHCP server information while creating a logical router by using the Contrail Command UI:

1. Click **Overlay > Logical Routers**.

The Logical Routers page is displayed.

2. Click **Create**.

The Create Logical Routers page is displayed.

3. Enter the following information as given in [Table 29 on page 177](#).

Table 29: Create Logical Router

Field	Action
Name	Enter a name for the logical router.
Admin State	Select Up as the admin state.
Extend to Physical Router	Select the physical router you want to extend the logical router to by selecting a router from the Extend to Physical Router list.

Table 29: Create Logical Router (Continued)

Field	Action
Logical Router Type	Select a logical router type from the Logical Router Type list.
Connected networks	Select the network(s) you want to connect the logical router to by selecting the network(s) from the Connected networks list.
Public Logical Router	Click Public Logical Router check box to enable the logical router to function as a public logical router.

- From the DHCP Relay Server(s) section, click **+Add**.

The **DHCP IP Address** field is displayed as seen in [Figure 61 on page 178](#).

Figure 61: DHCP IP Address Field

- Enter the IP address of the DHCP server in the **DHCP IP Address** field.

NOTE: Ensure that you remove existing CSNs before you provision the DHCP server. For more information, see ["Steps to Remove CSN Information" on page 178](#).

- (Optional) Click **+Add** to add more DHCP IP addresses.
- Click **Create**.

The logical router is now created and is listed in the Logical Routers page.

Steps to Remove CSN Information

Follow these steps to remove CSN information from the Contrail Command UI.

- Click **Infrastructure > Cluster**.

The Overview tab of the Cluster page is displayed.

2. Click the **Cluster Nodes** tab.

The Cluster AIO Nodes page is displayed.

3. Click the **Service Nodes** tab.

The list of CSNs are displayed.

4. To delete a CSN, hover over the name of the CSN and click the **Remove** icon.

The **Delete Service Nodes?** confirmation message is displayed.

5. Click **Delete** to confirm.

The CSN information is removed.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1908	Starting in Contrail Networking Release 1908, tenant administrators can define a set of DHCP server IP addresses while configuring virtual networks and logical routers on a multi-tenant data center fabric.

Return Material Authorization

IN THIS SECTION

- [Move a Device to RMA State | 180](#)
- [Replace a Device in RMA State with a New Device | 181](#)
- [Getting Started with a New Device | 182](#)

Contrail Networking Release 1907 supports Return Material Authorization (RMA). RMA is the process followed for repairing or replacing a defective device. You can create an RMA for a device after a Juniper Technical Assistance Center (JTAC) engineer has confirmed that the device is defective and has to be replaced or repaired. The device can then be replaced or repaired as per the standard service-level agreement (SLA) drawn at the time of purchase.

Once you find out that a device is defective, contact JTAC to determine whether the device has to be replaced or repaired. After JTAC confirms that the device has to be replaced or repaired, you can move

the device to RMA state. A device that is in RMA state cannot be configured and can be removed from the network.

With Contrail Networking Release 1908, Contrail supports upgrading a device that is replaced in a data center fabric to the image version specified (in the **OS Version** field) during the initial zero-touch-provisioning onboarding process.

Move a Device to RMA State

This topic provides instructions to move a device to RMA state by using the Contrail Command user interface.

1. Click **Infrastructure > Fabrics**.

The Fabrics page is displayed.

2. Click the name of the fabric you want to edit.

The Fabric devices page is displayed listing all the devices configured on the fabric.


3. Select the check box next to the name of the device you want to move to the RMA state.

You can select multiple devices at a time.

NOTE: Starting with Contrail Networking Release 1907, you can view the RMA state of a device from the Fabric Devices page. The statuses are **ACTIVE** and **RMA**.

4. Click the **Return Material Authorization** icon as shown in [Figure 62 on page 180](#) to move the selected device to the RMA state.

Figure 62: Return Material Authorization Icon



The following confirmation message is displayed:

Put the selected devices into RMA state?


5. Click **Confirm** to move the device to RMA state.

The device is now in RMA state.

Once you move a device from the Active state to RMA state, any configuration that you make on the fabric does not apply to the device that is in RMA state.

6. (Optional) To move a device from RMA state back to Active state,
 - a. Select the check box next to the name of the device you want to move to the Active state.
 - b. Click **Move Out of RMA** icon as shown in [Figure 63 on page 181](#).

Figure 63: Move Out of RMA Icon



The following confirmation message is displayed:

Reactivate selected device?

- c. Click **Confirm** to move the device to Active state.
The device is now in Active state.

Replace a Device in RMA State with a New Device

This topic provides instructions to replace a device in RMA state with a new device by using the Contrail Command user interface.

1. Click **Infrastructure > Fabrics**.
The Fabrics page is displayed.
2. Click the name of the fabric you want to edit.
The Fabric devices page is displayed listing all the devices on the fabric.
3. To select the device you want to replace, select the check box next to the name of the device.

NOTE: You can only replace devices that are in RMA state. To move a device to RMA state, see "[Move a Device to RMA State](#)" on page 180.

4. Click **Action** list and select **RMA Replacement**.
The RMA Replacement page is displayed.

The name and the serial number of the device in RMA state are displayed in the fields that are greyed out.

5. Enter the serial number of the new device in the **Serial Number** field.
6. Click **Replace** to confirm.

The Fabric Devices page is displayed listing the device you just added.

Getting Started with a New Device

After you have replaced the device in RMA state with a new device,

1. The Dynamic Host Configuration Protocol (DHCP) server allots an IP address to the new device.
2. The Contrail Networking Controller then discovers the temporary IP address of the new device from the DHCP leases table by using the serial number of the new device.
3. The Contrail Networking Controller pushes the initial configuration, including the old device's management IP address, to the new device.

The Contrail Networking Controller communicates with the device using the old IP address.

4. The new serial number and MAC address are saved in the `physical_router` object in the database.
5. The new device image is upgraded.

NOTE: With Contrail Networking Release 1908, the new device is upgraded to the device image version specified (in the **OS Version** field) during the initial zero-touch-provisioning onboarding process. For more information, see the **Provisioning Option - New Fabric** section of the "[Create a Fabric](#)" on page 7 topic.

6. Finally, the Contrail Networking Controller pushes the saved underlay and overlay configuration to the new device. Any configuration changes made while the device was in RMA state will also be pushed to the new device.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1908	With Contrail Networking Release 1908, the new device is upgraded to the device image version specified (in the OS Version field) during the initial zero-touch-provisioning onboarding process.

RELATED DOCUMENTATION

[Fabric Overview | 4](#)

[Create a Fabric | 7](#)

Approaches to Enable External Connectivity for Overlay Networks

Contrail Enterprise Multicloud (CEM) supports both QFX Series and MX Series devices. You can connect an overlay network to an external network by using either a QFX Series device or an MX Series device.

[Table 30 on page 183](#) lists the differences in configuration when you use a QFX Series device (with EVPN configured) and an MX Series device (with L3VPN configured).

Table 30: Enabling External Connectivity for Overlay Networks

Action	Use Case	QFX (EVPN)	MX (L3VPN)
Extending a Virtual Network	Enabling external connectivity to a layer 3 network. Uses L3VPN.	<ol style="list-style-type: none"> 1. Integrated Routed and Bridging (IRB) interface is created in inet.0. 2. No virtual routing and forwarding (VRF) instances are created. 	<ol style="list-style-type: none"> 1. Virtual switch with bridge domains (BD) are created. 2. IRB is created in VRF. 3. Configure static route to 0/0 inside the VRF. Apply appropriate filter to redirect traffic to the VRF.
		Verdict —Does not help to route traffic between the Internet and the virtual network.	Verdict —The right approach is when Layer 3 VPN (L3VPN) routing instance is used and no Source Network Address Translation (SNAT) is used.

Table 30: Enabling External Connectivity for Overlay Networks (Continued)

Action	Use Case	QFX (EVPN)	MX (L3VPN)
Extending an SNAT-LR	Enable external connectivity to a layer 3 network along with SNAT configuration. Uses L3VPN.	<ol style="list-style-type: none"> 1. IRB is created in inet.0. 2. No VRFs are created. 3. No Service Physical Interface Card (PIC). 	<ol style="list-style-type: none"> 1. Requires a Service PIC for SNAT. 2. Virtual switch, BDs, and VRFs are created. 3. IRB is created in VRF. 4. Configure static route to 0/0 inside the VRF. Apply appropriate filter to redirect traffic to the VRF.
		Verdict —Does not help to route traffic between the Internet and the virtual network	Verdict —The right approach is when Service PIC is present, L3VPN is used, and SNAT is used.
Extending a VXLAN-LR	Enable external connectivity from multiple layer 3 networks connected to a logical router. Uses EVPN.	<ol style="list-style-type: none"> 1. Virtual local area network (VLAN) created. 2. IRB created in VRF. 3. Type 5 route advertised in VRF. 4. Configure static route to 0/0 inside the VRF. Apply appropriate filter to redirect traffic to the VRF. 	<ol style="list-style-type: none"> 1. Virtual switch with BDs created. 2. Two VRFs created . Same IRB created in 2 VRFs, causing CommitError. 3. Static route to 0/0 inside the VRFs and appropriate filter to redirect traffic into the VRFs. 4. Virtual switch instances created with IRB.
		Verdict —The right approach is when VXLAN-LR is used.	Verdict —Does not help to route traffic between the Internet and the virtual network

Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles

IN THIS SECTION

- [Hardware Platforms and Associated Roles | 185](#)
- [Hardware Platforms and Associated Node Profiles and Roles | 191](#)

The following tables list the hardware platforms supported by Contrail Networking along with the node profiles and roles that can be configured on them.

Hardware Platforms and Associated Roles

The following tables list the supported hardware platforms and the roles that can be configured on them. The Contrail Networking releases and the corresponding Junos OS releases indicate the minimum release versions that must be installed on the hardware platforms to configure the required roles on them.

- For a list of QFX devices, see [Table 31 on page 186](#).
- For a list for MX devices, see [Table 32 on page 189](#).
- For a list of SRX devices, see [Table 33 on page 190](#).

Table 31: Supported QFX Series Switches

QFX Device	Supported from Contrail Networking Release Supported from Junos OS Releases												
	Physical Roles		Overlay Roles					Gateway Roles			Special Role		
	Leaf	Spine	CRB - Access	CRB -GW	CRB - MCast-GW	ERB - UCast-GW	lean	DC Gateway	DCI Gateway	PNF Service Chaining	Route Reflector	AR_Replicator	AR_Client
QFX 510-XX models	5.0.2 17.3R3	5.0.2 17.3R3	5.0.2 17.3R3				5.0.2 17.3R3				5.0.2 17.3R3		R1907 18.4R2
QFX 511-0-48 S-4C QFX 511-0-32 Q	5.0.2 17.3R3	5.0.2 17.3R3	5.0.2 17.3R3	5.0.2 18.1R3		5.1 18.1R3	5.0.2 17.3R3				5.0.2 17.3R3		R1907 18.4R2
QFX 512-0-48 Y-8C	5.1 18.4R2	5.1 18.4R2	5.1 18.4R2	5.1 18.4R2		5.1 18.4R2	5.1 18.4R2				5.1 18.4R2		R1907 18.4R2
QFX 512-0-48 T-6C	2008 20.2R1-S1	2008 20.2R1-S1	2008 20.2R1-S1	2008 20.2R1-S1		2008 20.2R1-S1	2008 20.2R1-S1				2008 20.2R1-S1		2008 20.2R1-S1

Table 31: Supported QFX Series Switches (Continued)

QFX Device	Supported from Contrail Networking Release													Supported from Junos OS Releases	
	Physical Roles		Overlay Roles					Gateway Roles			Special Role				
	Leaf	Spine	CRB - Access	CRB -GW	CRB - MCast-GW	ERB - UCast-GW	lean	DC Gateway	DCI Gateway	PNF Service Chaining	Route Reflector	AR_Replicator	AR_Client		
QFX 5120-32C	1909 19.1R2	1909 19.1R2	1912 19.1R2	2003 19.1R3		1912 19.1R2	1909 19.1R2				1909 19.1R2		R1909 19.1R2		
QFX 5200-32C-32Q	5.0.2 17.3R3	5.0.2 17.3R3	5.0.2 17.3R3				5.0.2 17.3R3				5.0.2 17.3R3		R1907 18.4R2		
QFX 5210-64C	5.0.2 17.3R3	5.0.2 17.3R3	5.0.2 17.3R3				5.0.2 17.3R3				5.0.2 17.3R3		R1907 18.4R2		
QFX 10002-36Q	5.0.2 17.3R3	5.0.2 17.3R3	5.0.2 17.3R3	5.0.2 17.3R3	5.1 17.3R3	5.1 18.1R3	5.0.2 17.3R3	5.1 18.1R3	2005 18.4R2-S3	5.1 18.1R3	5.0.2 17.3R3	R1907 18.4R2	R1907 18.4R2		

Table 31: Supported QFX Series Switches (Continued)

QFX Device	Supported from Contrail Networking Release													Supported from Junos OS Releases	
	Physical Roles		Overlay Roles					Gateway Roles			Special Role				
	Leaf	Spine	CRB - Access	CRB -GW	CRB - MCast-GW	ERB - UCast-GW	lean	DC Gateway	DCI Gateway	PNF Service Chaining	Route Reflector	AR_Replicator	AR_Client		
QFX 100 02-7 2Q	5.0.2 17.3R3	5.0.2 17.3R3	5.0.2 17.3R3	5.0.2 17.3R3	5.1 17.3R3	5.1 18.1R3	5.0.2 17.3R3	5.1 18.1R3		5.1 18.1R3	5.0.2 17.3R3	R19.07 18.4R2	R19.07 18.4R2		
QFX 100 08															
QFX 100 16															
QFX 100 02-6 0C	2003 19.1R3	2003 19.1R3		2003 19.1R3		2003 19.1R3	5.1 18.4R2	2003 19.1R3	2003 19.1R3	2003 19.1R3	2003 19.1R3		2003 19.1R3		

Table 32: Supported MX Series Routers

MX Device	Supported from Contrail Networking ReleaseSupported from Junos OS Releases								
	Physical Roles		Overlay Roles				Gateway Roles		Special Role
	Leaf	Spine	ERB-UCAST-Gateway	CRB-Gateway	CRB-MCAST-Gateway	null	DC-Gateway	DCI-Gateway	Route-Reflector
MX80	5.0.2 17.3R3	5.0.2 17.3R3				5.1 18.1R3	5.0.2 17.3R3		
MX240, MX480, MX960	5.0.2 17.3R3	5.0.2 17.3R3	2003 18.4R2-S3	2003 18.4R2-S3	2003 18.4R2-S3	5.1 18.1R3	2003 (without SNAT) and 2005 (with SNAT) 18.4R2-S3	2005 18.4R2-S3	5.0.2 17.3R3
MX2008, MX2010, MX2020	5.0.2 17.3R3	5.0.2 17.3R3	2003 18.4R2-S3	2003 18.4R2-S3	2003 18.4R2-S3	5.1 18.1R3	2003 18.4R2-S3 (without SNAT)		5.0.2 17.3R3
MX10003		5.1 18.1R3		2003 18.4R2-S3	2003 18.4R2-S3				5.0.2 17.3R3

Table 32: Supported MX Series Routers (Continued)

MX Device	Supported from Contrail Networking ReleaseSupported from Junos OS Releases								
	Physical Roles		Overlay Roles				Gateway Roles		Special Role
	Leaf	Spine	ERB-UCAST-Gateway	CRB-Gateway	CRB-MCAST-Gateway	null	DC-Gateway	DCI-Gateway	Route-Reflector
MX204, MX10008, MX10016	5.1 18.1R3	5.1 18.1R3	2003 18.4R2-S3	2003 18.4R2-S3	2003 18.4R2-S3	5.1 18.1R3	2003 18.4R2-S3 (without SNAT)		5.0.2 17.3R3
JNP10008, JNP10016	5.1 18.1R3	5.1 18.1R3	2003 18.4R2-S3	2003 18.4R2-S3	2003 18.4R2-S3	5.1 18.1R3	2003 18.4R2-S3 (without SNAT)		5.0.2 17.3R3

Table 33: Supported SRX Series Services Gateways

SRX Device	Supported from Contrail Networking Release
	Physical Role
	PNF
SRX4600, SRX4200, SRX4100, SRX5800, SRX5600, SRX5400, vSRX	5.1

Hardware Platforms and Associated Node Profiles and Roles

The following tables list the supported hardware platforms and the associated node profiles. The table also lists the roles that can be configured on these devices.

- For a list of QFX devices, see [Table 34 on page 191](#).
- For a list of MX devices, see [Table 35 on page 202](#).
- For a list of SRX devices, see [Table 36 on page 202](#).

Table 34: Supported QFX Series Switches

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
QFX10002-36Q	juniper-qfx10k	Leaf	CRB-Access, CRB-Gateway, DC-Gateway, Route-Reflector, ERB-UCAST-Gateway, DCI-Gateway, CRB-MCAST-Gateway, PNF-Servicechain, AR-Client, AR-Replicator

Table 34: Supported QFX Series Switches (Continued)

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
		Spine	lean, CRB-Access, CRB-Gateway, DC-Gateway, Route-Reflector, DCI-Gateway, CRB-MCAST-Gateway, PNF-Servicechain, AR-Client, AR-Replicator

Table 34: Supported QFX Series Switches (Continued)

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
QFX10002-72Q, QFX10016, QFX10008	juniper-qfx10k	Leaf	CRB-Access, CRB-Gateway, DC-Gateway, Route-Reflector, ERB-UCAST-Gateway, CRB-MCAST-Gateway, PNF-Servicechain, AR-Client, AR-Replicator

Table 34: Supported QFX Series Switches (Continued)

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
		Spine	leaf, CRB-Access, CRB-Gateway, DC-Gateway, Route-Reflector, CRB-MCAST-Gateway, PNF-Servicechain, AR-Client, AR-Replicator

Table 34: Supported QFX Series Switches (Continued)

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
QFX10002-60C	juniper-qfx10k	Leaf	CRB Access, CRB-Gateway, DC-Gateway, Route-Reflector, ERB-UCAST-Gateway, DCI-Gateway, AR Client NOTE: AR-Replicator role is not supported on Junos OS Release 19.2R2.

Table 34: Supported QFX Series Switches (Continued)

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
		Spine	lean, CRB-Gateway, DC-Gateway, Route-Reflector, ERB-UCAST-Gateway, DCI-Gateway, PNF-Servicechain NOTE: Contrail Networking Release 1909 supports QFX1000 2-60C device running Junos OS Release 19.1R2 and later. QFX1000 2-60C device works only if enterprise style of

Table 34: Supported QFX Series Switches (*Continued*)

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
			<p>configuration is enabled. To enable enterprise style of configuration, select the VLAN-ID Fabric Wide Significance checkbox when onboarding the QFX1000 2-60C device. For more information, see Create a Fabric.</p>
<p>QFX5100-XX models QFX5200-32C-32Q, QFX5210-64C</p>	<p>juni per- qfx5 k- lean</p>	<p>Leaf</p>	<p>CRB- Access, Route- Reflector, AR-Client</p>
		<p>Spin e</p>	<p>lean, Route- Reflector, AR-Client</p>

Table 34: Supported QFX Series Switches (Continued)

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
QFX5110-48S-4C, QFX5110-32Q	juniper-qfx5k	Leaf	CRB-Access, Route-Reflector, ERB-UCAST-Gateway, AR-Client
		Spine	lean, CRB-Access, CRB-Gateway, Route-Reflector, AR-Client
QFX5120-48Y-8C	juniper-qfx5k	Leaf	CRB-Access, CRB-Gateway, Route-Reflector, ERB-UCAST-Gateway, AR-Client
		Spine	lean-spine, Route-Reflector, AR-Client

Table 34: Supported QFX Series Switches (Continued)

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
QFX5120-48T-6C	juniper-qfx5120	Leaf	CRB-Access, CRB-Gateway, Route-Reflector, AR-Client, ERB-UCAST-Gateway
		Spine	lean-spine, Route-Reflector, AR-Client

Table 34: Supported QFX Series Switches (Continued)

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
QFX5120-32C	juniper-qfx5120	Leaf	CRB-Access, CRB-Gateway, Route-Reflector, AR-Client, ERB-UCAST-Gateway NOTE: Contrail Networking Release 1909 supports QFX5120-32C device running Junos OS Release 19.1R2 and later.

Table 34: Supported QFX Series Switches (*Continued*)

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
		Spine	lean-spine, Route-Reflector, AR-Client NOTE: Starting in Contrail Networking Release 1910, a QFX5120-32C device can be used in lean-spine deployments.
QFX5220-32CD QFX5220-128C NOTE: Starting with Contrail Networking Release 2011, you can configure QFX5220-XX devices running Junos OS Release 20.2R2 and later as a spine and superspine.	juniper-qfx5220	Spine Superspine	lean, Route-Reflector lean, Route-Reflector

Table 35: Supported MX Series Routers

MX Device	Node Profile	Physical Roles	Routing Bridging Roles
MX80	juniper-mx	Leaf	DC-Gateway, Route-Reflector
		Spine	DC-Gateway, null, Route-Reflector
MX240, MX480	juniper-mx	Leaf	DC-Gateway, Route-Reflector, DCI-Gateway, ERB-UCAST-Gateway
		Spine	DC-Gateway, null, Route-Reflector, DCI-Gateway, CRB-Gateway, CRB-MCAST-Gateway
MX204, MX960, MX2008, MX2010, MX2020, MX10008, MX10016	juniper-mx	Leaf	DC-Gateway, Route-Reflector, ERB-UCAST-Gateway
		Spine	DC-Gateway, null, Route-Reflector, CRB-Gateway, CRB-MCAST-Gateway
MX10003	juniper-mx	Spine	Route-Reflector, CRB-Gateway, CRB-MCAST-Gateway
JNP10008, JNP10016	juniper-mx	Leaf	Route-Reflector, DC-Gateway, ERB-UCAST-Gateway
		Spine	Route-Reflector, DC-Gateway, CRB-Gateway, CRB-MCAST-Gateway, null

Starting with Contrail Networking Release 2003, you can configure CRB-Gateway, ERB-UCAST-Gateway, and CRB-MCAST-Gateway routing-bridging roles on MX Series routers.

Table 36: Supported SRX Series Services Gateways

SRX Devices	Node Profile	Physical Roles	Routing Bridging Roles
SRX5400, SRX5600, SRX4600, SRX4100, SRX1500, SRX240H-POE, SRX5800, SRX4200	juniper-srx	PNF	PNF-Servicechain

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2011	Starting with Contrail Networking Release 2011, you can configure QFX5220-XX devices running Junos OS Release 20.2R2 and later as a spine and superspine.
2003	Starting with Contrail Networking Release 2003, you can configure CRB-Gateway, ERB-UCAST-Gateway, and CRB-MCAST-Gateway routing-bridging roles on MX Series routers.
1910	Starting in Contrail Networking Release 1910, a QFX5120-32C device can be used in lean-spine deployments.
1909	Contrail Networking Release 1909 supports QFX10002-60C device running Junos OS Release 19.1R2 and later.
1909	Contrail Networking Release 1909 supports QFX5120-32C device running Junos OS Release 19.1R2 and later.

4

CHAPTER

Managing Data Center Devices

Data Center Interconnect | 206

Logical Router Interconnect | 213

Configuring Data Center Gateway | 220

Virtual Port Groups | 236

Configuring Virtual Port Groups | 238

How to Detach and Attach a VMI to VPG | 246

Using Static, eBGP, PIM, and OSPF Protocols to Connect to Third-Party Network Devices | 258

Configuring Storm Control on Interfaces | 278

Creating Port Profiles, Storm Control Profiles, sFlow Profiles, or Telemetry Profiles by Cloning | 285

Configuring EVPN VXLAN Fabric with Multitenant Networking Services | 289

Edge-Routed Bridging for QFX Series Switches | 291

Activating Maintenance Mode on Data Center Devices | 293

Viewing the Network Topology | 295

Viewing Hardware Inventory of Data Center Devices | 303

Viewing Configuration of Devices Deployed in Contrail Fabric | 305

Detecting and Managing Manual CLI Configuration Changes | 308

Certificate Lifecycle Management Using Red Hat Identity Management | 314

Collapsed Spine Architecture | 318

Data Center Interconnect

IN THIS SECTION

- [Understanding Data Center Interconnect | 206](#)
- [Data Center Interconnect Deployment Topologies | 207](#)
- [Creating Data Center Interconnect | 208](#)

Contrail Networking supports the automation of data center interconnect (DCI) of two different data centers.

These topics provide information on data center interconnect deployment topologies and how you can create a data center interconnect.

Understanding Data Center Interconnect

You can automate data center interconnect (DCI) of two different data centers. Multiple tenants connected to a logical router in a data center can exchange routes with tenants connected to a logical router in another data center. All BGP routers in a data center should peer with local route reflectors and not with BGP routers on another fabric. Contrail Networking Release 5.1 supports layer three interconnect of data centers that exist in different fabrics. Starting in Contrail Networking Release 2011, layer 2 DCI functionality is also supported. Contrail Networking defines elements (spine switch and leaf switch) that belong to a data center.

A single Contrail Networking cluster can manage multiple data center pods that are composed of two-tier IP fabric. These data center pods are used to provision overlay layer 2 and layer 3 networking services as virtual networks and logical routers.

Contrail Networking automates the interconnection of logical routers (Layer 3 VRF) in each pod. A DCI object represents the extension of a logical router from one data center pod to another by using EVPN VXLAN Type 5 routes. These logical routers that are extended to the devices in each fabric are assigned DCI-Gateway role. The routing policies are configured on both pods to ensure EVPN type 5 routes are exchanged across the data center.

NOTE: The gateway devices must support DCI-Gateway routing bridging role.

Starting in Contrail Networking Release 2005, you can configure DCI-Gateway routing-bridging role on MX240, MX480, MX960, and MX10003 devices.

Data Center Interconnect Deployment Topologies

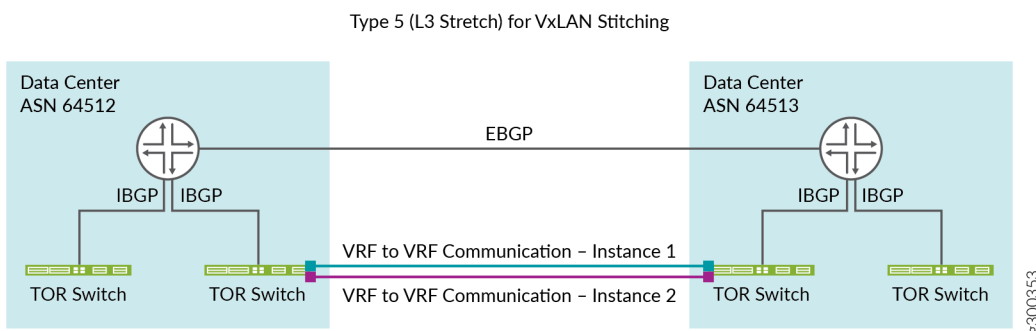
IN THIS SECTION

- DCI using EBGP | 207
- DCI using IBGP | 208

Contrail Networking supports the following data center interconnect (DCI) deployment topologies.

DCI using EBGP


Figure 64: DCI using EBGP Connection



DCI using EBGP connection establishes an EBGP connection between two data centers. The data centers are configured with two different autonomous system (AS) numbers as depicted in [Figure 64 on page 207](#).

DCI using IBGP

Figure 65: DCI using IBGP Connection



DCI using IBGP connection establishes an IBGP connection between two data centers. The data centers are configured with the same autonomous system (AS) numbers as depicted in [Figure 65 on page 208](#).

Creating Data Center Interconnect

IN THIS SECTION

- [Onboard Brownfield Devices | 209](#)
- [Create Virtual Network | 209](#)
- [Create Logical Routers | 210](#)
- [Create DCI | 210](#)

These topics provide step-by-step instructions to create data center interconnect.

Prerequisites

Before you start creating data center interconnect, ensure that:

- Junos OS 18.1 or later is installed
- Data center pods that Contrail Networking automates must have IP reachability
- Logical routers and client virtual networks are connected

- Logical routers extended to the devices in each fabric are assigned DCI-Gateway role
- BGP sessions between loopback addresses are reachable
- Underlay connectivity is enabled
- There is a route reflector on each data center that Contrail Networking is peering to

Follow these steps to create a data center interconnect.

Onboard Brownfield Devices

Follow the steps provided in the ["Onboard Brownfield Devices" on page 57](#) topic to onboard devices and assign roles to devices.

See Table 1 for an example configuration of how you can assign roles to a device. When you configure a QFX series device as a data center gateway, ensure that you assign DC-Gateway role to the spine and leaf device.

Table 37: Assign Roles to Devices

Device	Physical Role	Routing-Bridging Role
Spine devices	spine	CRB-Gateway, Route-Reflector, CRB-MCAST-Gateway, DCI-Gateway
Leaf devices	leaf	CRB-Access, DCI-Gateway

Create Virtual Network

Follow the steps provided in the ["Create Virtual Network" on page 82](#) topic to create virtual networks.

Before you begin, ensure that you

- Do not add network policy while creating the virtual network.

You can create the network policy and add it to the virtual network after you create the virtual network. For more information on creating a network policy, see ["Create Network Policy" on page 92](#).

- Have created a Network IPAM. For more information on creating a network IPAM, see ["Create Network IPAM" on page 94](#).

After you have created the virtual network and the network policy, follow these steps to attach the network policy to the virtual network.

1. Navigate to **Overlay>Virtual Networks**.

The All networks page is displayed.

2. To select the virtual network you want to add the policy to, select the check box next to the name of the virtual network. Then click the **Edit** icon at the end of the row.

The Edit Virtual Network page is displayed.

3. Select the network policy from the Network Policies list and click **Save**.

The policy is now added and the All networks page is displayed.

Create Logical Routers

Follow the steps provided in the "[Create Logical Routers](#)" on page 90 topic to configure logical routers.

While creating logical router, ensure that you

- Select **VXLAN Routing** as the Logical Router Type.
- Select the virtual network(s) from the Connected Networks list.
- Select the physical router (Spine device) to which you want to extend the logical router.

Create DCI

Follow these steps to create a DCI of two different data centers by using the Contrail Command user interface (UI).

1. Navigate to **Overlay > Interconnects**.

The Data Center Interconnect page is displayed.

2. Click **Create**.

The Create Data Center Interconnect page is displayed.

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

4. Select DCI mode.


You can select L2 or L3 DCI mode.

If you have selected **L2** as the DCI mode, the Fabric field, Available Virtual Networks table, and Selected Virtual Networks table are displayed. See [Figure 66 on page 211](#).

Contrail Networking Release 2011 supports layer 2 DCI functionality.

Figure 66: L2 DCI Mode

DCI name*

DCI Mode  L2 L3

Fabric

Available Virtual Networks

Add all <

NAME	VLAN	SUBNETS
No Virtual Networks found		

Selected Virtual Networks

Remove all

NAME	VLAN	SUBNETS
No Virtual Networks added		

Enter the following information.

- a. Select the fabrics that the data centers are a part of, from the Fabric list.
The available virtual networks that are part of Contrail are listed in the Available Virtual Networks table.
- b. From the Available Virtual Networks table, select the virtual networks you want included in the DCI by clicking the arrow next to each listed virtual network.
The virtual networks that you selected are displayed in the Selected Virtual Networks table.
- c. Click **Create** to create the L2 DCI mode.

If you have selected **L3** as the DCI mode, the Connections section is displayed. See [Figure 67 on page 212](#).

Figure 67: L3 DCI Mode

DCI name*

DCI Mode ⓘ
 L2 L3

Connections ⓘ

Select logical router*	Fabric	Extend to Physical Router (RB role = DCI-Gateway)* ⓘ
<input type="text" value="DC1-master-LR"/>	<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>

[+ Add](#)

[Create](#) [Cancel](#)

Enter the following information.

- a. Select a logical router from the **Select logical router** list.
- b. Select fabric from the **Select fabric** list.
- c. Select the physical router to which you want to extend the logical router to, from the **Extend to Physical Router (RB Role = DCI-Gateway)** list.
- d. Repeat steps 4.a through 4.c to create the next connection.
- e. Click **Create** to create the L3 DCI mode.

The DCI is now created and is listed in the Data Center Interconnect page.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2011	Starting in Contrail Networking Release 2011, layer 2 DCI functionality is also supported.

2005

Starting in Contrail Networking Release 2005, you can configure DCI-Gateway routing-bridging role on MX240, MX480, MX960, and MX10003 devices.

RELATED DOCUMENTATION

[VXLAN Data Center Interconnect Using EVPN Overview](#)

Logical Router Interconnect

IN THIS SECTION

- [Understanding Logical Router Interconnect | 213](#)
- [Creating Logical Router Interconnect | 215](#)

Contrail Networking Release 2005 supports the logical router interconnect feature, which enables interconnection of logical routers deployed in the same fabric.

These topics enable you to understand the way logical router interconnect works and create a logical router interconnect.

Understanding Logical Router Interconnect

In Contrail Networking Release 2005, the logical router interconnect feature enables a logical router in a fabric to interconnect with other logical routers deployed in the same fabric. The logical router interconnect feature is supported by logical routers that are deployed on QFX Series fabric devices. You can create logical router interconnection by leaking routes from a source logical router to multiple destination logical routers. In releases prior to release 2005, you could not leak routes from one logical router to other logical routers that are deployed in the same fabric.

In order to leak routes, you must select one logical router as a source logical router and one or more logical routers as destination logical routers in the Contrail Command user interface (UI). A source logical router leaks routes to multiple destination routers according to the type of **Export Source** you choose in the Contrail Command UI.

In the Contrail Command UI, you can choose **Routing Policy** or **Virtual Network** as an **Export Source**.

- If you select **Routing Policy** as the **Export Source**, you must ensure that the routing policy is configured for the QFX Series device, where the source logical router is deployed. Contrail Networking uses the routing policy terms to leak routes from a source logical router to all destination logical routers. If an incoming route matches the routing policy terms for a QFX Series device, that route is leaked from the source logical router to all destination logical routers deployed on the physical device.
- If you select **Virtual Network** as the **Export Source**, you must select the virtual networks that are connected to the source logical router. Contrail Networking automatically creates a routing policy where the `route-filter` value is assigned according to the subnets of the selected virtual networks. If an incoming route matches the routing policy terms of this routing policy, the source logical router can leak routes to all the destination logical routers.

If the source logical router and all destination logical routers are deployed on the same QFX Series device in a fabric, Contrail Networking uses `rib-groups` to leak routes from source logical routers to all destination logical routers. `rib-groups` import all routing policies as `import-policy` and all routing instances as `import-rib`.

If source logical router and destination logical routers are deployed on different QFX Series devices in a fabric, Contrail Networking uses `vrf-export` routing policies on source logical router and `vrf-import` routing policies on destination logical routers.

NOTE:

- The logical router interconnect feature is not supported by MX Series devices.
- The logical router interconnect feature does not support IPv6 routes.

SEE ALSO

rib-groups

vrf-import

Creating Logical Router Interconnect

IN THIS SECTION

- [Create a Fabric and Deploy Logical Routers on the Fabric Devices | 215](#)
- [Create a Routing Policy for QFX Series Devices | 215](#)
- [Creating Logical Router Interconnect | 217](#)

Contrail Networking Release 2005 supports interconnection between logical routers deployed in the same fabric. Logical router interconnect is supported by logical routers deployed on QFX Series devices that exist in the same fabric. Before configuring logical router interconnect, you must create a fabric, deploy logical routers on the fabric devices, and create a routing policy for the fabric device, where the logical router is deployed.

Create a Fabric and Deploy Logical Routers on the Fabric Devices

- Follow the procedure described in the ["Create a Fabric" on page 7](#) topic to create a fabric.
- Follow the procedure described in the ["Create Logical Routers" on page 90](#) topic to create logical routers on fabric devices.

The logical router interconnect feature is supported by logical routers that are deployed on QFX Series devices.

Create a Routing Policy for QFX Series Devices

Create a routing policy that is used by the source logical router to leak routes to all the destination logical routers. You must ensure that the routing policy terms are supported by the QFX Series device of the logical router.

1. Navigate to **Overlay>Routing>Routing Policies**.
2. Click **Add** to create a new routing policy.

Alternatively, you can also edit an existing routing policy for a QFX Series device. To edit an existing policy, select a policy from the displayed list and click the **Edit (pencil)** icon.

The **Create Routing Policy** page is displayed.

3. In the **Create Routing Policy** page, enter routing policy information according to the guidelines provided in [Table 38 on page 216](#).
4. Click **Create** to save the routing policy terms and create a routing policy for the logical router.

The **Routing Policies** tab is displayed listing the newly created policy.

Table 38: Create Routing Policy for QFX Series Device

Field	Guidelines
Name	Enter a name for the routing policy in the Name field.
Type	Select Physical Device or vRouter . You can create a routing policy for the type of device you select. Select Physical Device to create a routing policy for a QFX Series device.
Term(s)	
<i>From</i>	Select the matching conditions to be satisfied by the incoming routes.
Click Add Route filter . The Route Filter and Type fields are displayed.	
Route Filter	Enter an IP prefix address as a route filter in the Route Filter field.
Type	Select one or more types of prefix. If an incoming route satisfies the prefix match condition, the route is processed.

Table 38: Create Routing Policy for QFX Series Device (Continued)

Then

Select the actions to be performed on the matching routes. The supported actions and the values are:

Action	Value
action	<p>Reject-Reject the route that matches this term. No more terms are evaluated after hitting this term.</p> <p>Accept-Accept the route that matches this term. No more terms are evaluated after hitting this term.</p> <p>Next-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.</p>

NOTE:

- You must assign **Route Filter** as a match condition in a routing policy, if you want to use the routing policy as an export source in logical router interconnect.
- You don't need to assign a value to the **Community** field. Contrail automatically scans the route target value of a source logical router and assigns the route target community member to the routing policy.

Creating Logical Router Interconnect

Follow these steps in the Contrail Command UI to enable logical router interconnect between logical routers deployed in the same fabric.

1. Navigate to **Overlay > Interconnects > LR Interconnects**.

Click **Create** to configure a new logical router interconnect.

2. Enter values in the **Create LR Interconnect** page according to the guidelines provided in [Table 39 on page 218](#).

3. Click **Create** to create the logical router interconnect between a group of logical routers.

The **LR Interconnect** tab is displayed listing the newly created logical router interconnect.

Table 39: Create Logical Router Interconnect

Field	Action
Name	Enter a name for the logical router interconnect.
Description	Enter a description for the logical router Interconnect.
Select fabric	Select a fabric to create logical router Interconnect among the logical routers deployed on the QFX Series devices in the fabric.
<i>Source</i>	
Select logical router	Select a logical router from the list to assign the logical router as a source logical router.
Export Source	Select Routing Policy to leak routes using a routing policy. Select Virtual Network to leak routes using the list of tenant virtual networks connected to the source logical router.
Routing Policy	If you select Routing Policy as Export Source , select the routing policy for the source logical router from the list.
Available Virtual Networks	If you select Virtual Network as Export Source , the Available Virtual Networks list displays the list of virtual networks connected to the source logical router.
Selected Virtual Networks	The Selected Virtual Networks list displays the virtual networks you can use to leak routes from source logical router to all the destination logical routers.
<i>Destination</i>	
Select logical router	Select a logical router from the list to assign the it as a destination logical router.

Table 39: Create Logical Router Interconnect (*Continued*)

Field	Action
Extend to Physical Router	<p>Select the QFX Series device from the Extend to Physical Router list, where the destination logical routers are deployed.</p> <p>The destination logical routers extends the leaked routes to this QFX Series device.</p>
Add	Click Add to assign more destination logical routers (Optional).

NOTE:

- You cannot leak routes from destination logical routers to source logical router. Contrail Networking Release 2005 supports only unidirectional interconnection between source logical router and destination logical routers.
- You cannot assign a primary logical router as a source logical router and assign a public logical router as a destination logical router.
- You cannot assign a public logical router as a source logical router and assign a primary logical router as a destination logical router.

RELATED DOCUMENTATION

| [Creating Routing Policies for QFX Series Devices in Contrail Networking](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2005	Contrail Networking Release 2005 supports the logical router interconnect feature, which enables interconnection of logical routers deployed in the same fabric.

Configuring Data Center Gateway

IN THIS SECTION

- [Configuring QFX Series Devices as Data Center Gateway | 220](#)
- [Configuring MX Series Routers as Data Center Gateway | 234](#)

You can configure a QFX series device and an MX series router as a Data Center Gateway (DC-GW). DC-GW is an overlay role that is assigned to a QFX series switch or an MX series router to:

- Extend private network
- Extend public routable network

You can extend private network and extend public routable network with EVPN Type 5.

For more information on supported QFX series and MX series devices, see "[Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles](#)" on page 185.

Configuring QFX Series Devices as Data Center Gateway

IN THIS SECTION

- [Onboard Brownfield Devices | 221](#)
- [Add Bare Metal Server | 221](#)
- [Create Tenant Virtual Network | 223](#)
- [Add CSN Nodes | 230](#)
- [Create Logical Routers | 231](#)
- [Verification | 233](#)

You can configure a QFX series device as a DC-GW. For more information on supported QFX series devices, see ["Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185](#).

As an example, follow these steps to configure a QFX10000 device as a DC-GW.

Onboard Brownfield Devices

Follow the steps provided in the ["Onboard Brownfield Devices" on page 57](#) topic to onboard fabric devices and assign roles to devices.

See [Table 40 on page 221](#) for an example configuration of how you can assign roles to a device.

Table 40: Assign Roles to Devices

Device	Physical Role	Routing-Bridging Role
Spine devices QFX10000	spine	CRB-Gateway, Route-Reflector, CRB-MCAST-Gateway, DC-Gateway
Leaf devices	leaf	CRB-Access

Ensure that you assign the DC-Gateway role to the QFX10000 device as shown in [Table 40 on page 221](#).

Add Bare Metal Server

Follow these steps to add an existing bare metal server (BMS) by using the Contrail Command UI:

1. Click **Workloads>Instances**.
The Instances page is displayed.
2. Click **Create** to create a new instance.
The Create Instance page is displayed.
3. Select **Existing Baremetal Server** as the Server Type.
4. Enter the following information in the Create Existing Baremetal Server pane:

Table 41: Add Existing Bare Metal Server Information

Field	Action
Instance Name	Displays the name of the BMS instance.

Table 41: Add Existing Bare Metal Server Information (Continued)

Field	Action
Baremetal Node	Select a bare metal node.
Interface	Select an interface from the list.
IP Address	Enter IP address of the instance.
Virtual Network	Select a virtual network from the list.
VLAN ID	Enter VLAN ID.
Select Security Groups	Select default security group from the list.
Port Profile	Select a port profile from the list.
Native/Untagged	Select this check box to receive untagged packets without native VLAN ID.

Figure 68: Existing Bare Metal Server

Server Type ⓘ

Virtual Machine
 New Baremetal Server
 Existing Baremetal Server

Create Existing Baremetal Server

Instance Name*

Baremetal Node*

Associate interfaces

Interface	IP Address	Virtual Network*	VLAN ID*
<input type="text"/>	<input type="text" value="Enter valid IPv4"/>	<input type="text"/>	<input type="text" value="1"/>
Select Security Groups ⓘ	Port Profile	<input type="checkbox"/> Native/Untagged	
<input type="text"/>	<input type="text"/>		

+ Add

+ Add

5. Click **Create** to confirm.

Create Tenant Virtual Network

A virtual network is a collection of endpoints, such as virtual machine instances, that can communicate with each other. You can also connect virtual networks to your on-premises network. A virtual network in a EVPN VXLAN data center corresponds to a bridge domain for one tenant in a multi-tenant data center fabric.

Follow these steps to create a tenant virtual network from the Contrail Command user interface (UI).

1. Navigate to **Overlay>Virtual Networks**.

The All Networks page is displayed.

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

The Create Virtual Network page is displayed.

3. Enter a name for the network in the **Name** field.

4. Select VN Fabric Type.

Select **Routed** to enable routed virtual network functionality. A routed virtual network represents a layer 3 subnet between the fabric (border gateway) and the third-party physical network device.

For more information, see ["Using Static, eBGP, PIM, and OSPF Protocols to Connect to Third-Party Network Devices" on page 258](#).

Select **Switched** (default option) for tenant virtual network on leaf, bare metal server, or vRouter.

5. Select network policies from the **Network Policies** list. You can select more than one network policy.

Network policies provide connectivity between virtual networks by allowing or denying specified traffic. They define the access control lists to virtual networks. To create a new network policy, navigate to **Overlay>Network Policies**.

For more information on creating network policies, see ["Create Network Policy" on page 92](#).

NOTE: You can attach a network policy to the virtual network after you have created the virtual network.

6. Select any one of the following preferred allocation mode.
 - Flat subnet only
 - Flat subnet preferred
 - (Default) User defined subnet only
 - User defined subnet preferred

An allocation mode indicates how you choose a subnet. You select **Flat subnet only** or **Flat subnet preferred** allocation mode when the subnet is shared by multiple virtual networks. However, you select **(Default) User defined subnet only** or **User defined subnet preferred** allocation mode when you want to define a subnet range.

7. Enter subnet information as given in [Table 42 on page 224](#).

Table 42: Subnet Information

Field	Action
Network IPAM	Select the IP address management method that controls IP address allocation, DNS, and DHCP for the subnet.
CIDR	Enter the overlay subnet CIDR.
Allocation Pools	Enter a list of ranges of IP addresses for vRouter-specific allocation.
Gateway	Enter the gateway IP address of the overlay subnet. This field is disabled by default. To configure this field, uncheck Auto Gateway.
Service Address	Specify the user configured IP address for DNS Service instead of the default system allocated one.
Auto Gateway	This check box is enabled by default and gateway address is allocated by the system. When this box is unchecked, gateway address is user configurable.
DHCP	Select this check box if you want Contrail to provide DHCP service.
DNS	Select this check box if you want the vRouter agent to provide DNS service.

8. Enter host route information.

Host routes are a list of prefixes and next hops that are passed to the virtual machine through DHCP.

- a. **Route Prefix**—Enter a full CIDR value with an IP address and a subnet mask. For example, 10.0.0.0/24.

- b. **Next Hop**—Enter next hop address.

9. Enter floating IP pool information.

A floating IP address is an IP address (typically public) that can be dynamically assigned to a running virtual instance. You can configure floating IP address pools in project networks, then allocate floating IP addresses from the pool to virtual machine instances in other virtual networks.

- a. **Pool Name**—Enter pool name.
 - b. **Projects**—Select project from the list.
10. Enter fat flows information. See [Table 43 on page 225](#).

You can apply fat flows to all VMIs under the configured VN. Fat flows help reduce the number of flows that are handled by Contrail.

Table 43: Configure Fat Flow

Field		Action
Protocol		Select the application protocol.
Port		Enter a value between 0 through 65,535. Enter 0 to ignore both source and destination port numbers. NOTE: If you select ICMP as the protocol, the Port field is not enabled.
Ignore Address		Configure fat flows to support aggregation of multiple flows into a single flow by ignoring source and destination ports or IP addresses. If you select Destination, only the Prefix Aggregation Source fields are enabled. If you select Source, only the Prefix Aggregation Destination fields are enabled. If you select the None (selected by default), both Prefix Aggregation Source and Prefix Aggregation Destination fields are enabled.
Prefix Aggregation Source	Source Subnet	Enter the source IP address. Ensure that the source subnet of the flows match. For example, enter 10.1.0.0/24 to create fat flows with 10.1.0.0/24 as the subnet. The valid subnet mask range is /8 through /32. NOTE: For packets from the local virtual machine, source refers to the source IP of the packet. For packets from the physical interface, source refers to the destination IP of the packet.

Table 43: Configure Fat Flow (Continued)

Field		Action
	Prefix	<p>Enter source subnet prefix length.</p> <p>The prefix length you enter is used to aggregate flows matching the source subnet. For example, when the source subnet is 10.1.0.0/16 and prefix length is 24, the flows matching the source subnet is aggregated to 10.1.x.0/24 flows. The valid the prefix length range is /(subnet mask of the source subnet) through /32.</p>
Prefix Aggregation Destination	Destination Subnet	<p>Enter the destination IP address.</p> <p>Ensure that the destination subnet of the flows match. Enter 10.1.0.0/24 to create fat flows with 10.1.0.0/24 as the subnet. The valid subnet mask range is /8 through /32.</p> <p>NOTE: For packets from the local virtual machine, destination refers to the destination IP of the packet. For packets from the physical interface, destination refers to the source IP of the packet.</p>
	Prefix	<p>Enter the destination subnet prefix length.</p> <p>The prefix length you enter is used to aggregate flows matching the destination subnet. For example, when the source subnet is 10.1.0.0/16 and prefix length is 24, the flows matching the source subnet is aggregated to 10.1.x.0/24 flows. The valid prefix length range is /(subnet mask of the destination subnet) through /32.</p>

11. Enter routing policy and bridge domain information as given below.

- a. Select routing policy from the **Routing Policies** list.

To create a routing policy, navigate to **Overlay>Routing>Routing Policy**.

- b. Define a list of route target prefixes.

Enter an IP address in the ASN field and Target in the range 0 through 65,535, or ASN in the range 1 through 65,535 and Target in the range 1 through 4,294,967,295 if 4-byte ASN is disabled. If 4-byte ASN is enabled, enter ASN in the range 1 through 4,294,967,295 and Target in the range 0 through 65,535.

- c. Define export route targets.

You can advertise the matched routes from the local virtual routing and forwarding (VRF) table to the MPLS routing table.

Enter an IP address in the ASN field and Target in the range 0 through 65,535, or ASN in the range 1 through 65,535 and Target in the range 1 through 4,294,967,295 if 4-byte ASN is disabled. If 4-byte ASN is enabled, enter ASN in the range 1 through 4,294,967,295 and Target in the range 0 through 65,535.

d. Define import route targets.

Import the matched routes from the MPLS routing table and to the local virtual routing and forwarding (VRF) table.

Enter an IP address in the ASN field and Target in the range 0 through 65,535, or ASN in the range 1 through 65,535 and Target in the range 1 through 4,294,967,295 if 4-byte ASN is disabled. If 4-byte ASN is enabled, enter ASN in the range 1 through 4,294,967,295 and Target in the range 0 through 65,535.

e. Enter bridge domain information. See [Table 44 on page 227](#).

A bridge domain is a set of logical interfaces that share the same flooding or broadcast characteristics.

Table 44: Bridge Domains

Field	Action
Name	Enter a name for the Layer 2 or Layer 3 bridge domain.
I-SID	Enter a Service Identifier in the range from 1 through 16777215.
MAC Learning	Enable or disable MAC learning. MAC learning is the process of obtaining the MAC addresses of all the nodes in a virtual network. It is enabled by default.
MAC Limit	Configure the maximum number of MAC addresses that can be learned.
MAC Move Limit	Configure the maximum number of times a MAC address move occurs in the MAC move time window. A MAC move is when a MAC address appears on a different physical interface or within a different unit of the same physical interface.

Table 44: Bridge Domains (Continued)

Field	Action
Time Window (secs)	Configure the period of time over which the MAC address move occurs. The default period is 10 seconds.
Aging Time (secs)	Configure the MAC table aging time, the maximum time that an entry can remain in the Ethernet Switching table before it is removed. The default time period is 300 seconds.

12. Enter advanced configuration information as given in [Table 45 on page 228](#).

Table 45: Advanced Configuration

Field	Action
Admin State	Select the administrative state of the virtual network.
Reverse Path Forwarding	Enable or disable Reverse Path Forwarding (RPF) check for the virtual network.
Shared	Select to share the virtual network with all tenants.
External	Select the check box to make the virtual networks reachable externally.
Allow Transit	Select to enable the transitive property for route imports.
Mirroring	Select to mark the virtual network as a mirror destination network.
Flood Unknown Unicast	Select to flood the network with packets with unknown unicast MAC address. By default, the packets are dropped.

Table 45: Advanced Configuration (*Continued*)

Field	Action
Multiple Service Chains	Select to allow multiple service chains within two networks in a cluster.
IP Fabric Forwarding	Select to enable fabric based forwarding.
Forwarding Mode	Select the packet forwarding mode for the virtual network.
Extend to Physical Router(s)	Select the physical router to which you want to extend the logical router. The physical router provides routing capability to the logical router.
Static Route(s)	Select the static routes to be added to this virtual network.
QoS	Select the QoS to be used for this forwarding class.
Security Logging Object(s)	Select the security logging object configuration for specifying session logging criteria.
ECMP Hashing Fields	Configure one or more ECMP hashing fields. When configured all traffic destined to that VN will be subject to the customized hash field selection during forwarding over ECMP paths by vRouters.
PBB Encapsulation	Select to enable Provider Backbone Bridging (PBB) EVPN tunneling on the network.
PBB ETree	Select to enable PBB ETree mode on the virtual network which allows L2 communication between two end points connected to the vRouters. When the check box is deselected, end point communication happens through an L3 gateway provisioned in the remote PE site.

Table 45: Advanced Configuration (*Continued*)

Field	Action
Layer2 Control Word	Select to enable adding control word to the Layer 2 encapsulation.
SNAT	Select to provide connectivity to the underlay network by port mapping.
MAC Learning	Enable or disable MAC learning. MAC learning is the process of obtaining the MAC addresses of all the nodes in a virtual network. It is enabled by default.
Provider Network	Select the provider network. The provider network specifies VLAN tag and the physical network name.
IGMP enable	Enable or disable IGMP.
Multicast Policies	Select the multicast policies. To create a policy, navigate to Overlay>Multicast Policies .
Max Flows	Enter the maximum number of flows permitted on each virtual machine interface of the virtual network.

13. Click **Create.**

The All Networks page is displayed. The virtual network that you created is displayed on this page.

Add CSN Nodes

Follow these steps to add CSN Nodes to the fabric by using the Contrail Command UI:

Navigate to the EVPN fabric you provisioned.

1. Click the fabric name, and then click the fabric device.

The Fabric Device page is displayed.

2. Enter the following information:

Table 46: Add CSN Node to Fabric Device Information

Field	Action
Management IP	Enter management IP address.
VTEP Address	Enter VTEP address.
Loopback IP	Enter loopback IP address.
BGP Router	Select BGP router from the list.
Virtual Router Type	Select virtual router type from the list.
Existing CSN	Select existing CSN from the list.

3. Click **Save** to confirm changes to the fabric.

Create Logical Routers

A logical router replicates the functions of a physical router. It connects multiple virtual networks. A logical router performs a set of tasks that can be handled by a physical router, and contains multiple routing instances and routing tables.

Follow these steps to create a logical router (LR).

1. Navigate to **Overlay>Logical Routers** and click **Create**.
The Create Logical Routers page is displayed.
2. Enter the following information as given in [Table 47 on page 231](#).

Table 47: Create a Logical Router

Field	Action
Name	Enter a name for the Logical Router.

Table 47: Create a Logical Router (Continued)

Field	Action
Admin State	<p>Select the administrative state that you want the device to be in when the router is activated.</p> <p>Up is selected by default.</p>
Logical Router Type	Select SNAT Routing or VXLAN Routing from the list.
Choose Fabric	Select the fabric that you are associating this logical router to.
Connected Networks	Select the networks that you want to connect this logical router to.
Extend to Physical Router	<p>Select the physical router(s) to which you want to extend virtual networks or routed virtual networks to, from the Extend to Physical Router list.</p> <p>A physical router provides routing capability to the logical router.</p>
Reconfigure Physical Routers	<p>This link is enabled when you select a routed virtual network from the Connected networks list. Click Reconfigure Physical Router to reconfigure a physical router that you want to extend a virtual network to.</p> <p>For more information, refer to the Create Logical Routers section of the "Using Static, eBGP, PIM, and OSPF Protocols to Connect to Third-Party Network Devices" on page 258 topic.</p>
Public Logical Router	(Optional) Select this check box if you want the logical router to function as a public logical router.
NAT	<p>Select this check box to enable Network Address Translation (NAT).</p> <p>This check box is disabled by default.</p>
VxLAN Network Identifier	<p>Enter VXLAN network identifier in the range from 1 through 16,777,215.</p> <p>This field is disabled by default.</p>

Table 47: Create a Logical Router (Continued)

Field	Action
DHCP IP Address	<p>Enter DHCP relay server IP address.</p> <p>You can add more than one IP address. To add another address, click +Add.</p>
Route Target(s)	<p>Click +Add to add route targets.</p> <p>Enter Autonomous System (AS) number in the ASN field.</p> <ul style="list-style-type: none"> • Enter ASN in the range of 1-4,294,967,295, when 4 Byte ASN is enabled in Global Config. • Enter ASN in the range of 1-65,535, when 4 Byte ASN is disabled. • You can also add suffix <i>L</i> or <i>l</i> (<i>lower-case L</i>) at the end of a value in the ASN field to assign an AS number in 4-byte range. Even if the value provided in the ASN field is in the range of 1-65,535, adding <i>L</i> or <i>l</i> (<i>lower-case L</i>) at the end of the value assigns the AS number in 4-byte range. If you assign the ASN field a value in the 4-byte range, you must enter a value in the range of 0-65,535 in the Target field. <p>Enter route target in the Target field.</p> <ul style="list-style-type: none"> • Enter route target in the range of 0-65,535, when 4 Byte ASN is enabled and ASN field is assigned a 4-byte value. • Enter route target in the range of 0-4,294,967,295, when the ASN field is assigned a 2-byte value.

3. Click **Create** to create the logical router.

The Logical Routers page is displayed.

NOTE: The router_interface object (Virtual Port) is created as part of the LR creation and VN extension to Spines workflow. While planning the IP address for spines, you must be aware that an extra one IP address is required for the router_interface object which gets created automatically.

Verification

EVPN type 5 configuration is pushed to QFX10000 switch as a DC-GW.

Figure 69: EVPN Type 5 Configuration

```

set groups _contrail_overlay_evpn_ Interfaces irb unit 5 proxy-macip-advertisement
set groups _contrail_overlay_evpn_ Interfaces irb unit 5 family inet address 10.x7.x8.xx/28 virtual-gateway-address 10.x7.7x.1xx
set groups _contrail_overlay_evpn_ protocols evpn vni-options vni 5 vrf-target target:64512:8000004
set groups _contrail_overlay_evpn_ protocols evpn encapsulation vxlan
set groups _contrail_overlay_evpn_ protocols evpn extended-vni-list all
set groups _contrail_overlay_evpn_ policy-options policy-statement _contrail_vn-public-12-5-import term t1 from community target_64512_8000004
set groups _contrail_overlay_evpn_ policy-options policy-statement _contrail_vn-public-12-5-import term t1 then accept
set groups _contrail_overlay_evpn_ policy-options policy-statement _contrail_vn-public-12-5-export term t1 then community add target_64512_8000004
set groups _contrail_overlay_evpn_ policy-options policy-statement _contrail_vn-public-12-5-export term t1 then accept
set groups _contrail_overlay_evpn_ policy-options policy-statement _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6-import term t1 from community target_64512_8000005
set groups _contrail_overlay_evpn_ policy-options policy-statement _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6-import term t1 then accept
set groups _contrail_overlay_evpn_ policy-options policy-statement _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6-export term t1 then community add target_64512_8000005
set groups _contrail_overlay_evpn_ policy-options policy-statement _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6-export term t1 then accept
set groups _contrail_overlay_evpn_ policy-options community target_64512_8000004 members target:64512:8000004
set groups _contrail_overlay_evpn_ policy-options community target_64512_8000005 members target:64512:8000005
set groups _contrail_overlay_evpn_ switch-options vtep-source-interface lo0.0
set groups _contrail_overlay_evpn_ switch-options route-distinguisher x.5.x.5:1
set groups _contrail_overlay_evpn_ switch-options vrf-import _contrail_vn-public-12-5-import
set groups _contrail_overlay_evpn_ switch-options vrf-import _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6-import
set groups _contrail_overlay_evpn_ switch-options vrf-export _contrail_vn-public-12-5-export
set groups _contrail_overlay_evpn_ switch-options vrf-export _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6-export
set groups _contrail_overlay_evpn_ switch-options vrf-target Target:64512:1
set groups _contrail_overlay_evpn_ vlans bd-5 vlan-id 5
set groups _contrail_overlay_evpn_ vlans bd-5 l3-interface irb.5
set groups _contrail_overlay_evpn_ vlans bd-5 vxlan vni 5
set groups _contrail_overlay_evpn_type5_ Interfaces lo0 unit 1006 family inet address 12x.x.0.1/32
set groups _contrail_overlay_evpn_type5_ forwarding-options family inet filter input redirect_to_public_vrf_filter
set groups _contrail_overlay_evpn_type5_ protocols evpn default-gateway no-gateway-community
set groups _contrail_overlay_evpn_type5_ firewall family inet filter redirect_to_public_vrf_filter term term-100 then routing-instance _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6
set groups _contrail_overlay_evpn_type5_ firewall family inet filter redirect_to_public_vrf_filter term default-term then accept
set groups _contrail_overlay_evpn_type5_ routing-instances _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6 instance-type vrf
set groups _contrail_overlay_evpn_type5_ routing-instances _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6 interface lo0.1006
set groups _contrail_overlay_evpn_type5_ routing-instances _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6 interface irb.5
set groups _contrail_overlay_evpn_type5_ routing-instances _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6 vrf-import _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6-import
set groups _contrail_overlay_evpn_type5_ routing-instances _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6 vrf-export _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6-export
set groups _contrail_overlay_evpn_type5_ routing-instances _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6 routing-options static route 0.0.0.0/0 next-hop-inet 0
set groups _contrail_overlay_evpn_type5_ routing-instances _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6 protocols evpn ip-prefix-routes advertise direct-next-hop
set groups _contrail_overlay_evpn_type5_ routing-instances _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6 protocols evpn ip-prefix-routes encapsulation vxlan
set groups _contrail_overlay_evpn_type5_ routing-instances _contrail__contrail_lr_internal_vn_23640071-2302-4728-8424-5528df330ae8__-13-6 protocols evpn ip-prefix-routes vni 100

```

Configuring MX Series Routers as Data Center Gateway

IN THIS SECTION

- Onboard Brownfield Devices | 234
- Create Virtual Network | 235

You can configure an MX series router as a DC-GW. You must ensure that you assign the DC-Gateway routing-bridging role to the MX series router during device onboarding. For more information on supported MX series routers, see ["Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles"](#) on page 185.

Follow these steps to configure an MX series router as a DC-GW.

Onboard Brownfield Devices

Follow the steps provided in the ["Onboard Brownfield Devices"](#) on page 57 topic to onboard fabric devices and assign roles to devices.

Ensure that you also assign DC-Gateway routing-bridging role to the MX series router (spine device) while assigning device roles.

Create Virtual Network

After you have onboarded fabric devices and assigned roles to devices, you create a virtual network and extend it to the MX series router.

Follow these steps to create a virtual network and extend it to MX series router.

1. Navigate to **Overlay>Virtual Networks** and click **Create**.

The Create Virtual Network page is displayed.

2. Enter a name for the network in the **Name** field.
3. Select VN Fabric Type.

Select **Routed** to enable routed virtual network functionality. A routed virtual network represents a layer 3 subnet between the fabric (border gateway) and the third-party physical network device. For more information, see ["Using Static, eBGP, PIM, and OSPF Protocols to Connect to Third-Party Network Devices" on page 258](#).

Select **Switched** (default option) for tenant virtual network on leaf, bare metal server, or vRouter.

4. Enter subnet information as given in [Table 48 on page 235](#).

Table 48: Subnet Information

Field	Action
Network IPAM	Select the IP address management method that controls IP address allocation, DNS, and DHCP for the subnet.
CIDR	Enter the overlay subnet CIDR.

5. Click **Advanced** to view the advance configuration section.
6. Select the **External** check box to make the virtual network reachable externally.
7. Select the MX series router from the Extend to Physical Router(s) list.
8. Click **Create** to save configuration.

The MX series router is now configured as a DC-GW.

After you configure an MX series router as a DC-GW, you can enable DNAT. For more information on enabling DNAT in a DC-GW, see ["Destination Network Address Translation for Bare Metal Servers" on page 406](#).

RELATED DOCUMENTATION

[Fabric Overview | 4](#)

[Edge-Routed Bridging for QFX Series Switches | 291](#)

[Destination Network Address Translation for Bare Metal Servers | 406](#)

Virtual Port Groups

In Contrail Networking, a virtual port group (VPG) is a group of one or more physical interfaces attached to one or more virtual network interfaces. Each virtual network interface object corresponds to a VLAN ID and is attached to a Virtual Network (VN).


A virtual network interface is a virtual representation of a physical network interface, which may correspond directly to a network interface controller. In network virtualization, this enables a system to store the information and operate on the virtual interfaces independently, without involving the physical interfaces.

A VPG enables you to select multiple interfaces on the same device or on different devices. A VPG is similar to the link aggregation group (LAG) but supports both LAG and multihoming depending on whether you select the interfaces on the same devices or on different devices. A LAG is automatically created if you select more than one interface on the same device.

In LAG configuration, one or more physical interfaces on a switch (a QFX Series device) become members of a link aggregation group (LAG). The LAG is connected to the virtual network interface of a VN, where the bare metal server is deployed.

[Figure 70 on page 237](#) shows how the interfaces belonging to two devices are grouped using a LAG configuration. **VPG1** is a traditional LAG configuration and **VPG2** is a single-homed LAG configuration. **VPG3** is a virtual port group, which groups the physical interfaces on two QFX Series devices using multi-homed LAG configuration.

Figure 70: Virtual Port Group



Depending on whether **VLAN-ID Fabric-Wide Significance** field is selected or not, the behavior of virtual port groups is different in enterprise style (**VLAN-ID Fabric-Wide Significance** option enabled) and service provider style (**VLAN-ID Fabric-Wide Significance** option disabled) configurations.

In enterprise style configurations, the field **VLAN-ID Fabric-Wide Significance** is enabled and you can associate one VLAN ID only to one virtual network. Once you assign a VLAN ID to a virtual network, the VLAN ID field is greyed out because the VLAN ID has a one-to-one correspondence with the virtual network and cannot be assigned to another virtual network. This is to ensure that the same VLAN ID is not associated with more than one virtual network within the same enterprise style fabric. Also, you can use only one untagged VLAN within the same VPG. Once you select a virtual network, you cannot select the same virtual network again unless it is an untagged virtual network. However, the VLAN ID must be the same.

In service provider style configuration, the field **VLAN-ID Fabric-Wide Significance** is disabled and there is no restriction on the VLAN IDs to be assigned to virtual networks.

Unlike in enterprise style configuration, you can select the same virtual network twice. However, you must assign different VLAN ID to each to make it clearer since you can select the same virtual network twice in different VPGs.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1909	Depending on whether VLAN-ID Fabric-Wide Significance field is selected or not, the behavior of virtual port groups is different in enterprise style (VLAN-ID Fabric-Wide Significance option enabled) and service provider style (VLAN-ID Fabric-Wide Significance option disabled) configurations.

RELATED DOCUMENTATION

| [Configuring Virtual Port Groups | 238](#)

Configuring Virtual Port Groups

This topic describes how to create virtual port groups (VPGs) from Contrail Command UI. Contrail Networking Release 2008 introduces a redesigned VPG-creation workflow. To create a VPG, perform the steps described in "[No Link Title](#)" on page 238 if you are using release 2008 later and those described in "[No Link Title](#)" on page 243 if you are using releases 2003 and 2005.

- **For release 2008:**

In Contrail Networking Release 2008, you can create a VPG without attaching VLANs. You have the ability to add VLANs after the VPG is created. In scaled setups, there can be a large number of VLANs, making it very hard to manage inside the create or edit Virtual Port Group pages. Release 2008 simplifies the assignment of VLANs by introducing a dedicated page for management. The VPG creation workflow comprises two steps with the first step being configuration of the VPG. Only when the configuration step is completed successfully can you assign the VLANs which is the second step.

To create virtual port groups in Contrail Command in release 2008:

1. Navigate to **Overlay > Virtual Port Group > Create Virtual Port Group**.

The **New Virtual Port Group** wizard is displayed.

2. Enter a name for the virtual port group in the **Virtual Port Group Name** field.
3. Select the fabric from the **Fabric Name** list.

The available physical interfaces on the devices in the selected fabric are listed.

- From the **Available Physical Interface** box, select the physical interfaces to be included in the virtual port group by clicking the arrow next to each physical interface. The available physical interfaces are the interfaces available on TORs that are already onboarded.

The selected interfaces are displayed in the **Assigned Physical Interface** box.

If you select more than one interface on the same TOR as shown in [Figure 75 on page 244](#), a link aggregation group (LAG) is automatically created on the device.

- Select a security group from the **Security Groups** list.

For enterprise style fabric configuration, attach a security group to the virtual port group. The policies defined in the security group is assigned to all the ports in the virtual port group. For service provider style fabric configuration, you can attach a security group to every VLAN.

- Assign a port profile to the virtual port group by selecting a port profile from the **Port Profile** list.

A port profile functions like a container that can support multiple port-related configurations, and allows you to apply those configurations by attaching them to the port profile.

- Click **Next** to create the VPG. If VPG creation fails, an error message is displayed. If VPG creation is successful, you will be directed to the second step in the process, in which you can add the VLANs.
- (Optional) You can assign VLANs in this step of the wizard. You can also add VLANs in the **Overlay > Virtual Port Group** page (see ["10" on page 241](#)). To add VLANs here, enter the information as shown in [Table 49 on page 239](#).

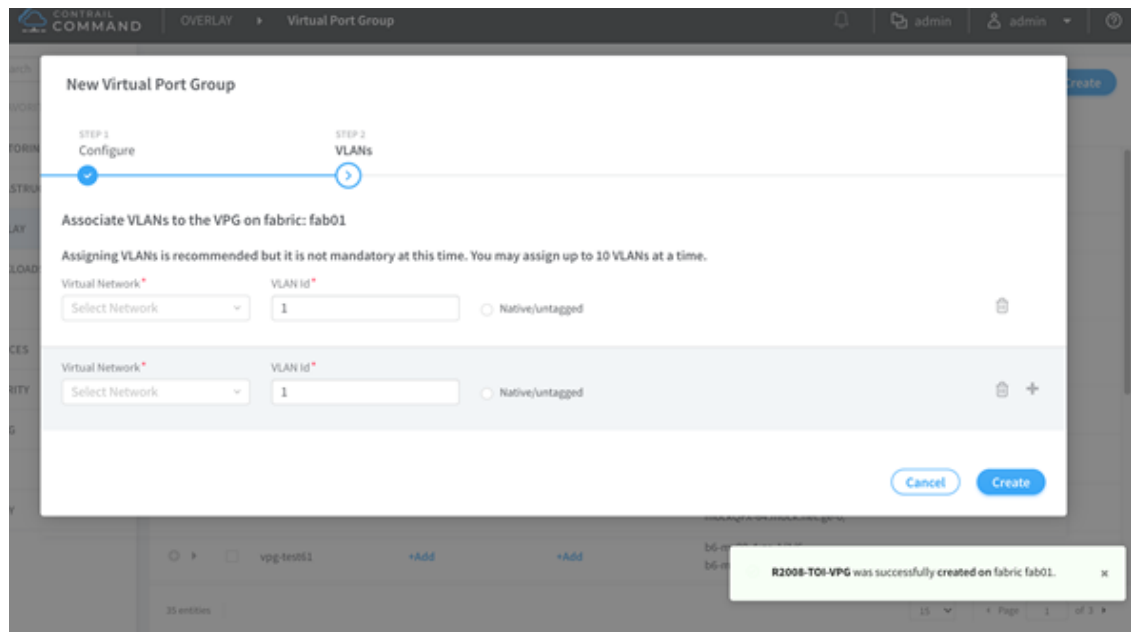
Table 49: Enter VLAN Information

Field	Action
Virtual Network	Select the virtual network to which the virtual port group belongs.
VLAN ID	Enter the VLAN ID and network to which the VLAN is associated. If you enable the VLAN-ID Fabric-Wide Significance option when creating a fabric, you can associate one VLAN ID to only one virtual network. This ensures that the same VLAN ID is not associated with more than one virtual network within the same enterprise style fabric.

Table 49: Enter VLAN Information (Continued)

Field	Action
Native/untagged	Select this check box to allow a native/untagged virtual network (optional). You can assign only one native/untagged VLAN in a virtual port group.
Security Group	This field is available only in service provider style fabric configuration. Select a security group from the Security Groups list. You can attach a security group to each VLAN.

Figure 71: Assign VLANs



9. Click **Create**.

The newly created virtual port group is displayed in the Virtual Port Group page with details of the interfaces as shown in [Figure 72 on page 241](#).

Figure 72: Virtual Port Groups

NAME	VLANs	PHYSICAL INTERFACES	PORT PROFILE
PavelTest01	0 23	b6-mx80-4-re-0/0/0 b6-mx80-4-ge-1/2/8 + 1 more	
vpg-internal-0	1	b6-mx80-4-ae0	
vpg-test70	150 0 + 1 more +Add	mockQFX-64.mock.net-ge-0	
tst-tst	+Add	b6-mx80-4-ge-1/3/5 b6-mx80-4-ge-1/3/2 + 8 more	
tst-13454	+Add	b6-mx80-4-ge-1/3/3 b6-mx80-4-ge-1/2/9	
vpg-test64	+Add	mockQFX-64.mock.net-ge-0 mockQFX-64.mock.net-ge-0	
vpg-test62	+Add	mockQFX-64.mock.net-ge-0 mockQFX-64.mock.net-ge-0	
vpg-test61	+Add	b6-mx80-4-ge-1/1/6 b6-mx80-4-ge-1/0/10	

10. (Optional) To assign VLANs if not previously configured or to edit configured VLANs, perform one of the following steps.

- To edit or add only VLANs, click a VLAN or click **Add** next to the VPG name. The VLANs assignment page is displayed.
- To edit VPG information and/or edit VLANs, select a VPG and click the edit (pencil) icon. The **Edit VPG** page is displayed.

Edit the VPG information as required. Click **Save** to save the changes and remain on this page. Alternatively, click **Save and assign new VLANs** to save the changes and assign VLANs. The VLANs assignment page is displayed.

Figure 73: Edit VPG

Virtual Port Group name: R2008-TOI-VPG
 Fabric name: fab01
 Security Groups: default
 Port Profile: port-profile-1

Available Physical Interface

DISPLAY NAME	PHYSICAL ROUTER
ge-1/2/1	b6-mx80-4
ge-1/1/11	b6-mx80-4
ge-1/1/1	b6-mx80-4
ge-1/2/3	b6-mx80-4

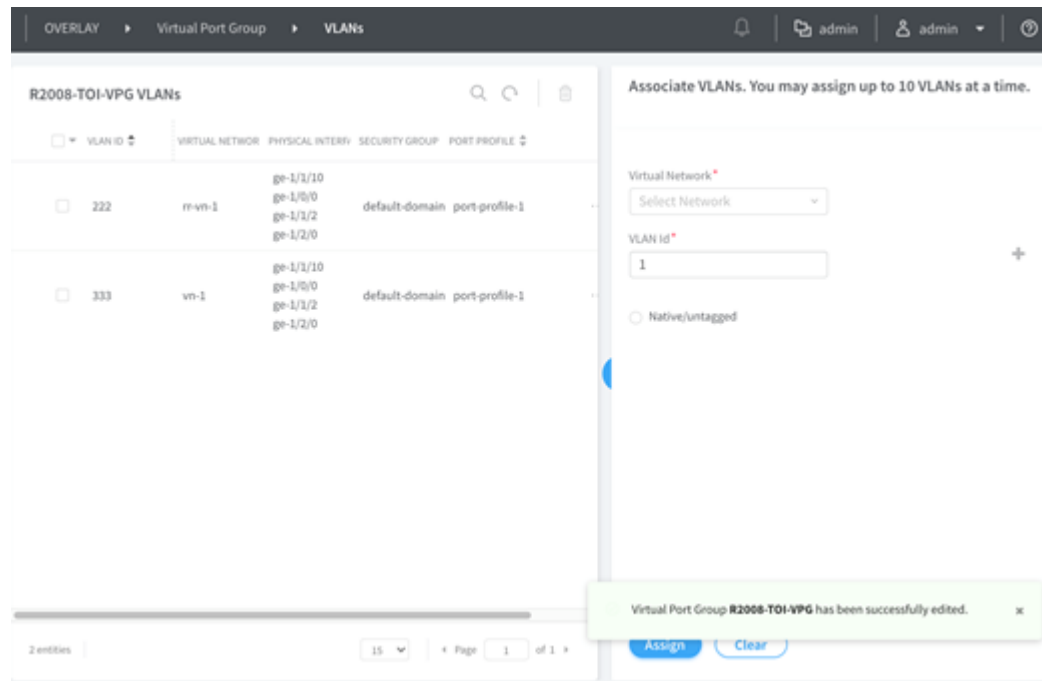
Assigned Physical Interface

DISPLAY NAME	PHYSICAL ROUTER
ge-1/1/10	b6-mx80-4
ge-1/2/0	b6-mx80-4
ge-1/0/0	b6-mx80-4
ge-1/1/2	b6-mx80-4

Buttons: Cancel, Save and assign new VLANs, Save

- The VLANs assignment page has two panels. The left panel lists all currently configured VLANs, if any. The right panel enables you to assign additional VLANs. Enter VLAN information and click **Assign** to attach the VLANs. The VLANs appear in the left panel. You can attach up to 10 VLANs at a time. You can also edit existing VLANs from this page. Successful and failed attempts at assigning and editing are indicated through success or error message pop-ups.

Figure 74: Edit VLANs



For better visibility, you can hide the right panel by clicking the blue expansion icon. You can also use this page to delete individual VLANs and bulk delete multiple VLANs.

- **For releases 2003 and 2005:**

To create virtual port groups in Contrail Command using releases 2003 and 2005:

1. Navigate to **Overlay > Virtual Port Group > Create Virtual Port Group**.

The Create Virtual Port Group page is displayed.

2. Enter a name for the virtual port group in the **Virtual Port Group Name** field.

3. Select virtual port group type.

With Contrail Networking Release 2003, you can create a routed virtual port group from the Contrail Command UI. Select the **Routed** option button to create a routed virtual port group. Select **Layer 2** option button to create a virtual port group.

4. Select the fabric from the **Fabric Name** list.

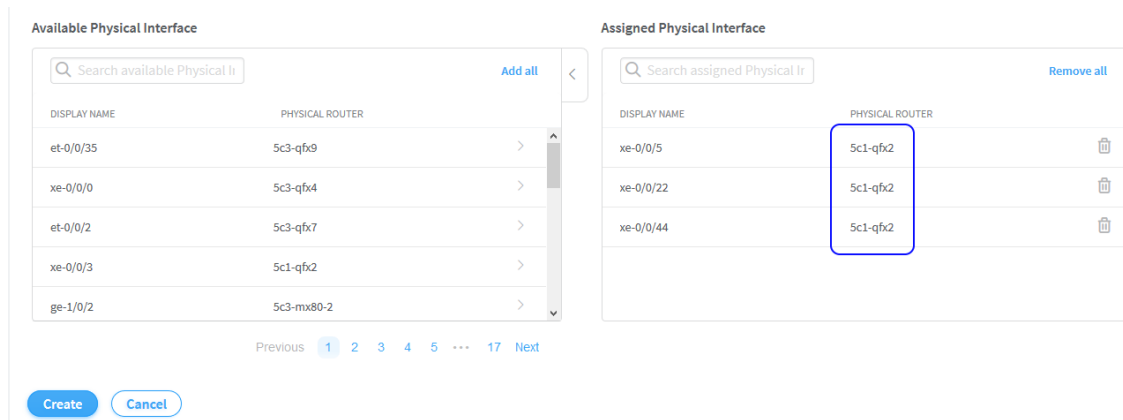
The available physical interfaces on the devices in the selected fabric are listed.

5. From the **Available Physical Interface** box, select the physical interfaces to be included in the virtual port group by clicking the arrow next to each physical interface. The available physical interfaces are the interfaces available on TORs that are already onboarded.

The selected interfaces are displayed in the **Assigned Physical Interface** box.

If you select more than one interface on the same TOR as shown in [Figure 75 on page 244](#), a link aggregation group (LAG) is automatically created on the device.

Figure 75: Select Interfaces on the Same TOR



- Assign a security group to the virtual port group by selecting a security group from the **Security Groups** list.

The policies defined in the security group is assigned to all the ports in the virtual port group.

- Select and assign a port profile from the **Port Profile** list.

A port profile functions like a container that can support multiple port-related configurations, and allows you to apply those configurations by attaching them to the port profile.

- Enter the following information as given in [Table 50 on page 244](#).

Table 50: Enter VLAN Information

Field	Action
Network	Select the virtual network to which the virtual port group belongs.
VLAN ID	Enter the VLAN ID and network to which the VLAN is associated. If you enable the VLAN-ID Fabric-Wide Significance option when creating a fabric, you can associate one VLAN ID to only one virtual network. This ensures that the same VLAN ID is not associated with more than one virtual network within the same enterprise style fabric.

Table 50: Enter VLAN Information (Continued)

Field	Action
Display Name	Enter the VLAN name. If the Auto Display Name field is selected, this field is autogenerated from the virtual port group name.
Auto Display Name	Select Auto Display Name if you want the VLAN name to be autogenerated from the virtual port group name.
Native/untagged	Select this check box to allow a native/untagged virtual network (optional). You can assign only one native/untagged VLAN in a virtual port group.

9. Click **Create**.

The newly created virtual port group is displayed on the Virtual Port Group page with details of the interfaces and the TORs as shown in [Figure 76 on page 245](#).

Figure 76: Virtual Port Groups

NAME	VLAN IDS	TOR PORT VLAN IDS	PHYSICAL INTERFACES	VIRTUAL NETWORK
vpg-internal-0		4094	ge-0/0/9:contrail-qfx5110-6	right_vn_1
vpg-test		4094	fxp0:contrail-srx5600-2 xe-0/0/32:2:bng-contrail-qfx-1... 1 more	left_vn_13

No items selected

You can delete a virtual port group by clicking the delete icon against the virtual port group. To delete a virtual port group, you must first remove the referenced VMI and the associated BMS instance from the virtual port group.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2008	In Contrail Networking Release 2008, you can create a VPG without attaching VLANs. You have the ability to add VLANs after the VPG is created.
2003	With Contrail Networking Release 2003, you can create a routed virtual port group from the Contrail Command UI. Select the Routed option button to create a routed virtual port group.

RELATED DOCUMENTATION

| [Virtual Port Groups](#) | [236](#)

How to Detach and Attach a VMI to VPG

IN THIS SECTION


- [Detaching a Tagged VMI from a VPG](#) | [247](#)
- [Attaching a Tagged VMI to a VPG](#) | [251](#)
- [Detaching an Untagged VMI from a VPG](#) | [253](#)
- [Attaching an Untagged VMI to a VPG](#) | [256](#)

<Optional>This topic describes how to detach and attach a tagged as well as untagged Virtual Machine Interface (VMI) to a Virtual Port Group (VPG).

Before performing this procedure, it is mandatory to:


- Create a VPG in Contrail Command user interface. For more information about configuring VPG, see ["Virtual Port Groups" on page 236](#).

Figure 77: Creating a Virtual Port Group




- Ensure a tagged VMI (VLAN) already exists in VPG. If a tagged VMI does not exist, then create one following step 10 and 11 from ["Configuring Virtual Port Groups"](#) on page 238.

Figure 78: Creating a tagged VMI in VPG



- Similarly, ensure an untagged (native) VMI (VLAN) exists in VPG. If it does not exist, then create an untagged VMI.

Figure 79: Creating an Untagged VMI in VPG



Detaching a Tagged VMI from a VPG

Before detaching a tagged VMI from a VPG, you need to know how VMI is stored in Cassandra Config database. Use Config Editor from Contrail UI, or Config API port, '8082' to view the VMI. The following code snippet shows how a VMI is stored.

```
{
  "virtual-machine-interface": {
```

```

"virtual_machine_interface_bindings": {
  "key_value_pair": [
    {
      "value": "vpg_1",
      "key": "vpg"
    },
    {
      "value": "{\\"local_link_information\\":
[{\\"port_id\\":\\"xe-0/0/47\\",\\"switch_id\\":\\"xe-0/0/47\\",
\\"switch_info\\":\\"5b11-qfx10\\",\\"fabric\\":\\"test-fabric\\"}]}",
      "key": "profile"
    },
    {
      "value": "baremetal",
      "key": "vnic_type"
    },
    {
      "value": "vrouter",
      "key": "vif_type"
    }
  ]
},
"virtual_port_group_back_refs": [
  {
    "to": [
      "default-global-system-config",
      "test-fabric",
      "vpg_1"
    ],
    "href": "http://10.87.69.8:8082/virtual-port-group/350196b7-1334-4c34-9f4f-b8796449a2a1",
    "attr": null,
    "uuid": "350196b7-1334-4c34-9f4f-b8796449a2a1"
  }
],
"name": "vpg_1-11",
"display_name": "vpg_1-11",
"uuid": "c6298b0d-1907-43f0-81f3-e0d9f43774bf",
"virtual_machine_interface_properties": {
  "sub_interface_vlan_tag": 11
},
"virtual_port_group_refs": [
  {
    "to": [

```

```

    "default-global-system-config",
    "test-fabric",
    "vpg_1"
  ],
  "href": "http://10.87.69.8:8082/virtual-port-group/350196b7-1334-4c34-9f4f-b8796449a2a1",
  "attr": {
    "vlan_tag": 11
  },
  "uuid": "350196b7-1334-4c34-9f4f-b8796449a2a1"
}
]
}
}

```

To view how VPG is stored in the Cassandra Config database, use Config Editor from Contrail UI, or Config API port, '8082'. The following code snippet shows how VPG is stored in the Cassandra Config database.

```

{
  "virtual-port-group": {
    "virtual_machine_interface_back_refs": [
      {
        "to": [
          "default-domain",
          "admin",
          "vpg_1-11"
        ],
        "href": "http://10.87.69.8:8082/virtual-machine-interface/c6298b0d-1907-43f0-81f3-e0d9f43774bf",
        "attr": {
          "vlan_tag": 11
        },
        "uuid": "c6298b0d-1907-43f0-81f3-e0d9f43774bf"
      }
    ],
    "display_name": "vpg_1",
    "uuid": "350196b7-1334-4c34-9f4f-b8796449a2a1",
    "virtual_machine_interface_refs": [
      {
        "to": [
          "default-domain",
          "admin",

```

```

        "vpg_1-11"
      ],
      "href": "http://10.87.69.8:8082/virtual-machine-interface/c6298b0d-1907-43f0-81f3-
e0d9f43774bf",
      "attr": null,
      "uuid": "c6298b0d-1907-43f0-81f3-e0d9f43774bf"
    }
  ],
}
}

```

To detach a tagged VMI from a VPG:

1. Create a **detach.json** file and then run the curl command with 'detach.json' as parameter:

```

# cat detach.json
{
  "virtual-machine-interface": {
    "virtual_machine_interface_bindings": {
      "key_value_pair": [
        {
          "value": "true",
          "key": "detach_vpg"
        }
      ]
    }
  }
}

# cat detach_command.sh
#!/bin/bash
curl -X PUT -H "X-Auth-Token: $TOKEN"
http://10.87.69.9:8082/virtual-machine-interface/c6298b0d-1907-43f0-81f3-e0d9f43774bf
-H "Content-type: application/json" -d @detach.json

```


In this command,

- **https** or **http** is used depending on whether config-api has SSL enabled (https) or not (http).
- **-H "X-Auth-Token: \$TOKEN"** is used when RBAC is enabled.
- **\$TOKEN** is used as a placeholder for an Openstack token.

After detaching a VMI from the VPG, the following 3 changes are noticeable:

- In the Config Editor, VMI does not contain any “virtual_machine_interface_bindings”, any “virtual_port_group_refs” or any “virtual_port_group_back_refs”.
- In the Config Editor, VPG does not contain any “virtual_machine_interface_refs” or any “virtual_machine_interface_back_refs”.
- In Contrail Command UI, VMI has disappeared from VPG.

Figure 80: VMI disappeared from VPG



Attaching a Tagged VMI to a VPG

To attach a tagged VMI to a VPG:

1. Create an **attach.json** file and then run the curl command with 'attach.json' as parameter:

```
# cat attach.json
{
  "virtual-machine-interface": {
    "virtual_machine_interface_bindings": {
      "key_value_pair": [
        {
          "key": "attach_vpg",
          "value": "true"
        },
        {
          "key": "vpg",
          "value": "vpg_1"
        },
        {
          "key": "profile",
          "value": "{\"local_link_information\":
[{\\"port_id\\":\\"xe-0/0/47\\",\\"switch_id\\":\\"xe-0/0/47\\",
\\"switch_info\\":\\"5b11-qfx10\\",\\"fabric\\":\\"test-fabric\\"}]}"
```

```

    },
    {
      "key": "vnic_type",
      "value": "baremetal"
    },
    {
      "key": "vif_type",
      "value": "vrouter"
    }
  ]
}
}
}

# cat attach_command.sh
#!/bin/bash
curl -X PUT -H "X-Auth-Token: $TOKEN"
http://10.87.69.9:8082/virtual-machine-interface/c6298b0d-1907-43f0-81f3-e0d9f43774bf
-H "Content-type: application/json" -d @attach.json


```

NOTE: In order to add the value associated to the "profile" key in attach.json, you can retrieve it from Config Editor, VMI entry, "virtual_machine_interface_bindings" dictionary (when the VMI is created and it is attached to the VPG).

After attaching a VMI to the VPG, the following 3 changes are noticeable:

- In the Config Editor, VMI contains "virtual_machine_interface_bindings", "virtual_port_group_refs" and "virtual_port_group_back_refs".
- In the Config Editor, VPG contains "virtual_machine_interface_refs" and "virtual_machine_interface_back_refs".
- In the Contrail Command UI, VMI appears in VPG.

Figure 81: VMI appears in VPG



Detaching an Untagged VMI from a VPG

Before detaching an untagged VMI from a VPG, you need to know how VMI is stored in Cassandra Config database. Use Config Editor from Contrail UI, or Config API port, '8082' to view the VMI. The following code snippet shows how a VMI is stored.

```
{
  "virtual-machine-interface": {
    "virtual_machine_interface_bindings": {
      "key_value_pair": [
        {
          "value": "vpg_1",
          "key": "vpg"
        },
        {
          "value": "{\"local_link_information\": [{\"port_id\": \"xe-0/0/47\", \"switch_id\": \"xe-0/0/47\", \"switch_info\": \"5b11-qfx10\", \"fabric\": \"test-fabric\"}]}",
          "key": "profile"
        },
        {
          "value": "baremetal",
          "key": "vnic_type"
        },
        {
          "value": "vrouter",
          "key": "vif_type"
        },
        {
          "value": "15",
          "key": "tor_port_vlan_id"
        }
      ]
    }
  }
}
```



```

    ]
  },
  "virtual_port_group_back_refs": [
    {
      "to": [
        "default-global-system-config",
        "test-fabric",
        "vpg_1"
      ],
      "href": "http://10.87.69.8:8082/virtual-port-group/350196b7-1334-4c34-9f4f-b8796449a2a1",
      "attr": null,
      "uuid": "350196b7-1334-4c34-9f4f-b8796449a2a1"
    }
  ],
  "name": "vpg_1-15-untagged",
  "display_name": "vpg_1-15-untagged",
  "uuid": "48359942-31e3-4743-9a1a-f31600246aa3"
  "virtual_machine_interface_properties": {
    "sub_interface_vlan_tag": 0
  },
  "virtual_port_group_refs": [
    {
      "to": [
        "default-global-system-config",
        "test-fabric",
        "vpg_1"
      ],
      "href": "http://10.87.69.8:8082/virtual-port-group/350196b7-1334-4c34-9f4f-b8796449a2a1",
      "attr": {
        "native_vlan_tag": 15
      },
      "uuid": "350196b7-1334-4c34-9f4f-b8796449a2a1"
    }
  ]
}
}
}

```

In the above code snippet, an untagged (native) VLAN (VMI) has "tor_port_vlan_id" in addition to other fields present in a tagged VLAN (VMI).

To view how VPG is stored in the Cassandra Config database, use Config Editor from Contrail UI, or Config API port, '8082'. The following code snippet shows how VPG is stored.

```
{
  "virtual-port-group": {
    "virtual_machine_interface_back_refs": [
      {
        "to": [
          "default-domain",
          "admin",
          "vpg_1-15-untagged"
        ],
        "href": "http://10.87.69.8:8082/virtual-machine-interface/48359942-31e3-4743-9a1a-f31600246aa3",
        "attr": {
          "native_vlan_tag": 15
        },
        "uuid": "48359942-31e3-4743-9a1a-f31600246aa3"
      }
    ],
    "display_name": "vpg_1",
    "uuid": "350196b7-1334-4c34-9f4f-b8796449a2a1",
    "virtual_machine_interface_refs": [
      {
        "to": [
          "default-domain",
          "admin",
          "vpg_1-15-untagged"
        ],
        "href": "http://10.87.69.8:8082/virtual-machine-interface/48359942-31e3-4743-9a1a-f31600246aa3",
        "attr": null,
        "uuid": "48359942-31e3-4743-9a1a-f31600246aa3"
      }
    ],
  }
}
```

To detach an untagged VMI from a VPG:

1. Create a **detach.json** file and then run the curl command with 'detach.json' as parameter. The detach.json is identical to detach.json from previous section. The curl command is similar, except that the URL in the curl command now contains the UUID of the native (untagged) VMI.

```
# cat detach_command.sh
#!/bin/bash
curl -X PUT -H "X-Auth-Token: $TOKEN"
http://10.87.69.9:8082/virtual-machine-interface/48359942-31e3-4743-9a1a-f31600246aa3
-H "Content-type: application/json" -d @detach.json
```


In this command,

- **https** or **http** is used depending on whether config-api has SSL enabled (https) or not (http).
- **-H "X-Auth-Token: \$TOKEN"** is used when RBAC is enabled.
- **\$TOKEN** is used as a placeholder for an Openstack token.

After detaching a VMI from the VPG, following 3 changes are noticeable:

- In the Config Editor, VMI does not contain any “virtual_machine_interface_bindings”, any “virtual_port_group_refs” or any “virtual_port_group_back_refs”.
- In the Config Editor, VPG does not contain any “virtual_machine_interface_refs” or any “virtual_machine_interface_back_refs”.
- In the Contrail Command UI, VMI has disappeared from VPG.

Figure 82: VMI disappeared from VPG



Attaching an Untagged VMI to a VPG

To attach an untagged VMI to a VPG:

1. Create a **attach.json** file and then run the curl command with 'attach.json' as parameter.

```
# cat attach.json
{
  "virtual-machine-interface": {
    "virtual_machine_interface_bindings": {
      "key_value_pair": [
        {
          "key": "attach_vpg",
          "value": "true"
        },
        {
          "key": "vpg",
          "value": "vpg_1"
        },
        {
          "key": "profile",
          "value": "{\\"local_link_information\\":
[[\\"port_id\\":\\"xe-0/0/47\\",\\"switch_id\\":\\"xe-0/0/47\\",\\"switch_info\\":\\"5b11-qfx10\\",
\\"fabric\\":\\"test-fabric\\"}}]"
        },
        {
          "key": "vnic_type",
          "value": "baremetal"
        },
        {
          "key": "vif_type",
          "value": "vrouter"
        },
        {
          "key": "tor_port_vlan_id",
          "value": "15"
        }
      ]
    }
  }
}

# cat attach_command.sh
#!/bin/bash
curl -X PUT -H "X-Auth-Token: $TOKEN"
```


```
http://10.87.69.9:8082/virtual-machine-interface/48359942-31e3-4743-9a1a-f31600246aa3
-H "Content-type: application/json" -d @attach.json
```

NOTE: In order to add the value associated to the "profile" key in attach.json, you can retrieve it from Config Editor, VMI entry, "virtual_machine_interface_bindings" dictionary (when the VMI is created and it is attached to the VPG).

After attaching a VMI to a VPG, the following 3 changes are noticeable:

- In the Config Editor, VMI contains "virtual_machine_interface_bindings", "virtual_port_group_refs" and "virtual_port_group_back_refs".
- In the Config Editor, VPG contains "virtual_machine_interface_refs" and "virtual_machine_interface_back_refs".
- In Contrail Command UI, VMI appears in VPG.

Figure 83: VMI appears in VPG



Using Static, eBGP, PIM, and OSPF Protocols to Connect to Third-Party Network Devices

IN THIS SECTION

- [Overview | 259](#)
- [Steps to Connect to a Third-Party Device | 260](#)

Starting in Contrail Networking Release 2003, you can use the Contrail Command user interface (UI) to connect the border gateway devices to third-party devices that are not managed by Contrail Networking. In earlier releases, Contrail Networking did not support connecting to an unmanaged third-party device. However, with this release, you can use the Contrail Command UI to configure border leaf devices or spine devices to connect to third-party devices.

Overview

You can use the Contrail Command user interface (UI) to connect to third-party devices that are not managed by Contrail Networking. If you want to connect two border leaf devices to a third-party device, you have to create routed virtual networks. You then create a virtual port group to connect the physical interfaces of the border leaf devices to the unmanaged third-party device. You then create logical routers and associate the routed virtual networks with the logical routers. You can configure eBGP and static routing protocols between the logical router and the unmanaged third-party device. Starting in Contrail Networking Release 2005, you can also configure PIM and OSPF routing protocols between the logical router and the unmanaged third-party device. However, the routing protocol that you can configure depends on the border gateway device that you use. Configuration is then pushed to the border gateway devices and a connection is formed between the fabric and the external network through the third-party unmanaged device.

For example, service chaining through a physical network function (PNF) such as a firewall or a loadbalancer are different use cases that use third-party devices. The PNF is outside the fabric, and you use Contrail Networking to configure border gateway devices to steer traffic through the PNF. Another example is to provide connectivity between tenant virtual networks that are in the fabric, and external layer 3 networks that are connected by a third-party router. This topic is an example of PNF service chaining between tenant virtual networks through an external firewall cluster.

The general workflow is as follows:

1. Create routed virtual networks. Define the routed virtual network for each security zone between border gateway devices and third-party device.

A routed virtual network represents a layer 3 subnet between the fabric (border gateway) and the external physical network function.

2. Configure virtual port groups. Connect the border gateway devices on the fabric to the third-party device by using EVPN Ethernet Segment Identifier-Link Aggregation Group (ESI-LAG) interface.
3. Configure logical routers and configure Integrated Routed and Bridging (IRB) interface by manually configuring the IP address of the IRB interface of the logical router.
4. Configure the routing protocols (eBGP, static, PIM, or OSPF) between the logical router and the external physical network function (for example, a firewall).

NOTE: The routing protocol that you can configure depends on the border gateway device that you use.

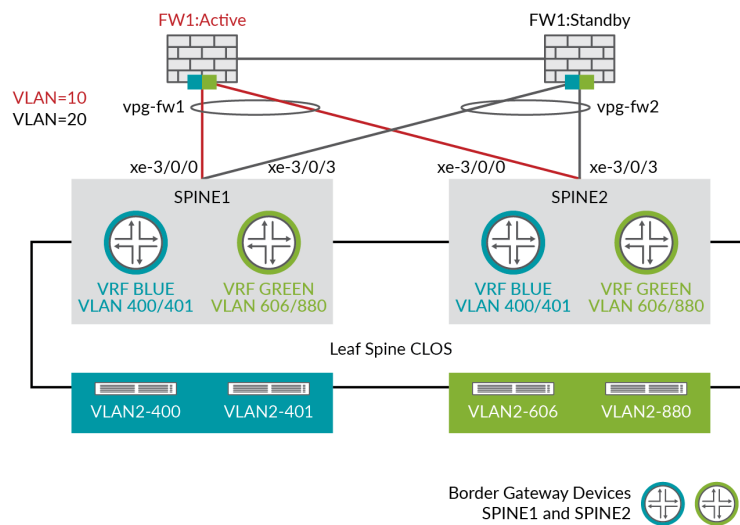
Steps to Connect to a Third-Party Device

IN THIS SECTION

- Topology | 260
- Before You Begin | 261
- Create Routed Virtual Networks | 262
- Create Routed Virtual Port Groups | 264
- Create Logical Routers | 267

Topology

Figure 84: Connecting to an Unmanaged Third-Party Device



8300955

Figure 84 on page 260 is an example topology of how you can connect to an unmanaged third-party (Firewall) device.

In this topology, there are two firewall zones, blue and green. There are two virtual networks associated with each zone. VLAN2-400 and VLAN2-401 are the virtual networks in the blue zone. VLAN-606 and VLAN-880 are the virtual networks in the green zone. Each zone has a logical router that is created on spine devices. These spine devices are configured with CRB-Gateway roles. Traffic between the virtual network subnets within the same zone is routed within the logical router VRF routing table on the spine.

This topology shows that the traffic between the green zone and the blue zone passes through the firewall cluster which is connected to the two spine devices. The spine devices also act as the border gateway devices. The two PNF firewalls are clustered in the A-S (active-standby) mode and is represented as a single device. The routing protocol on the logical router on the QFX device to the firewall cluster, will use the virtual gateway IP (VIP) of the firewall cluster as the next-hop (BGP peer) address.

The logical routers are created on the border gateway devices. You extend the tenant virtual networks and the routed virtual networks to the logical router.

The two blue tenant virtual networks and the two green tenant virtual networks are connected to the border gateway devices. Traffic from the blue tenant virtual networks has to pass through the third-party firewall (FW1:Active) to reach the green tenant virtual networks.

Follow these steps to connect to a third-party device not managed by Contrail Networking.

Before You Begin

Before you begin, ensure that you have

1. Discovered and onboarded fabric devices.
2. Assigned physical and routing-bridging roles to devices.
3. The tenant virtual networks are already created and are associated with TOR switches.

The VPGs are already created on the TOR switches.

4. • You can use a QFX Series device with CRB-Gateway, CRB-MCAST-Gateway, DCI-Gateway, or DC-Gateway roles assigned, as border gateway device.

However, in this topology, the ERB-UCAST-Gateway role cannot be configured on the border gateway device. You will not be able to reconfigure physical routers and configure routed virtual network properties if you assign ERB-UCAST-Gateway role to border gateway device.

NOTE: We recommend that you use QFX10000 series switches in border leaf and spine roles for the border function in an EVPN fabric.

If you use a QFX Series switch as the border gateway device, you can configure eBGP, static, PIM, or OSPF routing protocols on the border gateway device.

- Starting in Contrail Networking Release 2008, you can use MX240, MX480, MX960, or MX10003 device as a border leaf and spine device, to connect to third-party network devices. You can assign CRB-Gateway, CRB-MCAST-Gateway, DCI-Gateway, or DC-Gateway roles to these border gateway devices.

If you use an MX Series router as the border gateway device, you can configure eBGP and static routing protocols on the border gateway device.

Create Routed Virtual Networks

IN THIS SECTION

- [Procedure | 262](#)

Procedure

Step-by-Step Procedure

A routed virtual network is used to extend a fabric that is managed by Contrail Networking to a third-party device. By creating a routed virtual network, you create layer 3 links between border gateway devices and third-party devices. You can run eBGP, static, PIM, or OSPF routing protocols between these border gateway devices and third-party devices.

NOTE: If you use an MX240, MX480, MX960, or MX10003 device as border gateway devices, you can only configure eBGP and static routing protocols on the border gateway devices.

Contrail Networking Release 2003 supports creating a routed virtual network from the **Overlay>Virtual Networks** page of the Contrail Command UI.

You create two routed virtual networks, VN-BLUE and VN-GREEN, for this workflow. Follow these steps to create a routed virtual network:

1. Navigate to **Overlay>Virtual Networks**.

The All Networks page is displayed.

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

The Create Virtual Network page is displayed.

3. Enter VN-BLUE in the Name field.
4. Select **Routed** VN Fabric Type to create a routed virtual network.

With Contrail Networking Release 2003, you select **Routed** from the VN Fabric Type field to create a routed virtual network. See [Figure 85 on page 263](#).

Figure 85: Create Routed Virtual Network

OVERLAY > Virtual Networks > Create Virtual Network

Network Tags Permissions

Name* ?
Enter Name

VN Fabric Type ?
Routed Switched

Network Policies ?
Select Network Policies

Allocation Mode ?

5. In the Subnets section, click **+Add** to add a new subnet.

Enter the CIDR for this routed virtual network in the CIDR field.

You do not configure Allocation Pools, Gateway, DNS or DHCP for the subnet CIDR of a routed virtual network. These fields are greyed out by default.

NOTE: Do not configure options such as Host Routes; Floating IP Pools; Fat Flows; Routing, Bridging, and Policies; and Advanced options for a routed virtual network. Fields with default values can be left as is.

6. Click **Create** to create the routed virtual network.

The routed virtual network that you created is displayed on the Virtual Networks page.

Repeat steps "2" on page 263 through "6" on page 264 to create the Green-VN routed virtual network.

Create Routed Virtual Port Groups

IN THIS SECTION

- Procedure | 264

Procedure

Step-by-Step Procedure

A virtual port group (VPG) is a group of one or more physical interfaces attached to one or more virtual machine interfaces (VMI). Each VMI object corresponds to a VLAN ID and is attached to a Virtual Network. For more information, see "[Virtual Port Groups](#)" on page 236.

The VPG workflow provides a framework for you to configure the connection between the external device (PNF) and the border gateway devices of the fabric by using the EVPN ESI-LAG. The VPG is configured as a routed VPG and is associated with the routed virtual networks.

Contrail Networking Release 2003 supports creating a routed virtual port group from the **Overlay>Virtual Port Group** page of the Contrail Command UI. A routed VPG can only contain routed virtual networks.

Follow these steps to create a routed virtual network:

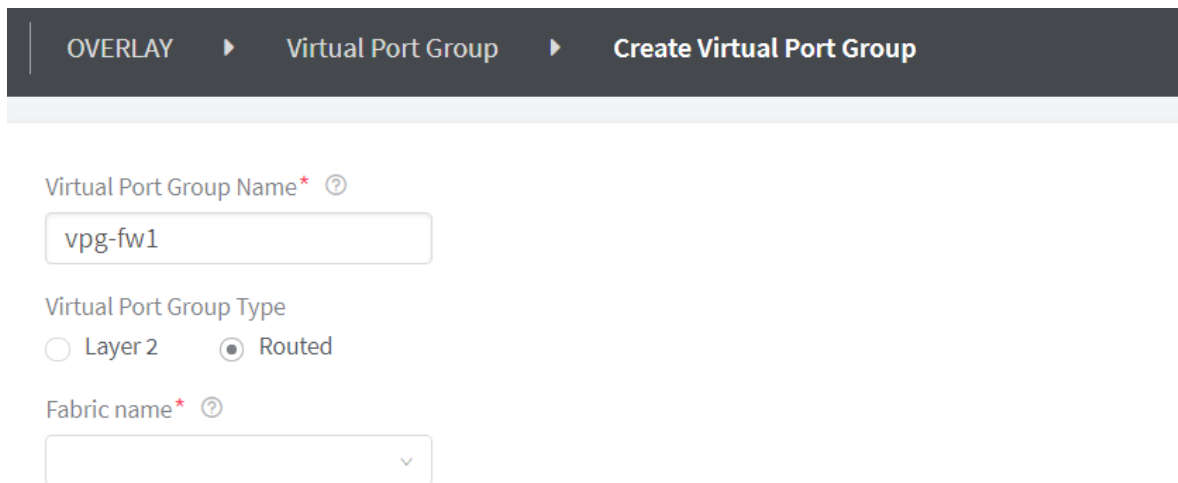
1. Navigate to **Overlay>Virtual Port Group>Create Virtual Port Group**.

The Create Virtual Port Group page is displayed.

2. Enter a name for the virtual port group in the **Virtual Port Group Name** field.
3. Select **Routed** virtual port group type.

With Contrail Networking Release 2003, you can create a routed virtual port group from the Contrail Command UI. Select the **Routed** option button to create a routed virtual port group. Select **Layer 2** option button to create a virtual port group. See [Figure 86 on page 265](#).

Figure 86: Create Routed Virtual Port Group



OVERLAY > Virtual Port Group > Create Virtual Port Group

Virtual Port Group Name* ⓘ
vpg-fw1

Virtual Port Group Type
 Layer 2 Routed

Fabric name* ⓘ
▼

4. Select the fabric from the **Fabric Name** list.

The available interfaces and physical routers in the selected fabric are listed in the Available Physical Interface box.

5. From the **Available Physical Interface** box, select the physical interfaces to be included in the virtual port group by clicking the arrow next each physical interface. The available physical interfaces are the interfaces available on TORs that are already onboarded. It also includes the associated physical router.

The selected interfaces are displayed in the **Assigned Physical Interface** box.

Figure 87: Assign Physical Interface

OVERLAY > Virtual Port Group > Create Virtual Port Group

Virtual Port Group Name*

Virtual Port Group Type
 Layer 2 Routed

Fabric name*

Available Physical Interface

DISPLAY NAME	PHYSICAL ROUTER
ge-3/0/15	cswj-0dc4-0015
xe-1/0/6:1	cswj-0dc4-0016
et-1/0/3	cswj-0dc4-0015
et-0/0/49	cswj-0dc4-0013
et-0/0/1	cswj-0dc4-0001

Assigned Physical Interface

No Physical Interface matching current criteria

Create Cancel

6. Enter the following information in the VLAN section.

- a. Select the routed virtual network from the Network list.

For the blue network, select VN-BLUE. See [Figure 88 on page 267](#). For the green network, select VN-GREEN.

- b. Enter VLAN ID and the network to which the VLAN is associated.

For VN-BLUE virtual network, enter 10. For VN-Green virtual network, enter 20.

- c. Enter the VPG name. If the Auto Display Name field is selected, this field is autogenerated from the virtual port group name.
- d. Select **Auto Display Name** if you want the VLAN name to be autogenerated from the virtual port group name.
- e. Select the **Native/untagged** check box to allow untagged virtual network. You cannot select the same virtual network more than once if the check box is not selected.

Click **+Add** to add VN-GREEN routed virtual network information. Repeat steps provided in step "6" on page 266.

Figure 88: Select Network

The screenshot shows a configuration interface for selecting networks. It features two rows of configuration fields:

- Row 1:**
 - Network:** VN-BLUE (dropdown menu)
 - VLAN ID:** 10 (text input)
 - Display Name:** vpg-fw1-10 (text input)
 - Auto Display Name:** (checkbox)
- Row 2:**
 - Network:** VN-GREEN (dropdown menu)
 - VLAN ID:** 20 (text input)
 - Display Name:** vpg-fw1-20 (text input)
 - Auto Display Name:** (checkbox)

Additional elements include a '+ Add' button at the bottom left and a 'Native/untagged' checkbox on the left side of the first row.

7. Click **Create**.

The newly created routed virtual port group is displayed on the Virtual Port Group page.

Create Logical Routers

IN THIS SECTION

- [Procedure | 267](#)

Procedure

Step-by-Step Procedure

A logical router performs a set of tasks and replicates the functions that can be handled by a physical router. A logical router connects multiple virtual networks.

You create two logical routers, LR-BLUE and LR-Green, for this workflow. Follow these steps to create LR-Blue.

1. Navigate to **Overlay > Logical Routers** page.

The Logical Routers page is displayed.

2. Click **Create**.

The Create Logical Router page is displayed. See [Figure 89 on page 269](#).

3. Enter LR-Blue in the Name field.
4. Select **Up** from the Admin State list.

This is the administrative state that you want the device to be in when the router is activated.

5. Select **VXLAN Routing** from the Logical Router Type list.
6. Select the fabric from the Choose Fabric list.
7. Select the networks you want to connect to this logical router from the Connected Networks list.

When you want overlay tenant traffic to pass through a third-party device, select both tenant virtual networks and routed virtual network from the from the Connected Networks list.

In this step, select VN-BLUE as the routed virtual network and select VLAN2-400 and VLAN2-401 as its two associated tenant virtual networks. See [Figure 89 on page 269](#).

When you create the LR-Green logical router, you select VN-GREEN as the routed virtual network and select VLAN2-606 and VLAN2-880 as its two associated tenant virtual networks in this step.

Figure 89: Create Logical Router

OVERLAY ▸ Logical Routers ▸ Create Logical Router

Logical Router Tags Permissions

Name* ⓘ

Admin State ⓘ
 Up Down

Logical Router Type ⓘ

Choose Fabric*

Connected networks ⓘ

Extend to Physical Router* ⓘ
 [Reconfigure Physical Routers](#)

Public Logical Router ⓘ
 NAT

[Create](#) [Cancel](#)

Adding a Loopback Routed Virtual Network—Starting in Contrail Networking Release 2005, you can connect a loopback routed virtual network to a logical router. You can select a loopback routed virtual network from the Connected Networks list only if you have provided loopback subnet information during the fabric onboarding process. For more information, see ["Create a Fabric" on page 7](#). Loopback subnet information can also be added after fabric onboarding from the **Infrastructure>Fabric>Fabric Name>Namespaces** page.

NOTE: If you have not provided loopback subnet information during fabric onboarding, you will not be able to select a loopback routed virtual network from the Connected Networks list.

When you select a routed virtual network from the Connected Networks list, the letter **R** is automatically added before the name to denote that the virtual network selected is a routed virtual network. See [Figure 89 on page 269](#).

When you select a loopback routed virtual network from the Connected Networks list, the letter **L** is automatically added before the name to denote that the virtual network selected is a lookback routed virtual network. See [Figure 90 on page 270](#).

When you select a routed virtual network, the Reconfigure Physical Routers link is enabled.

Figure 90: Adding a Loopback Routed Virtual Network

OVERLAY > Logical Routers > Create Logical Router

Logical Router Tags Permissions

Name* ⓘ
LR-Blue

Admin State ⓘ
 Up Down

Logical Router Type ⓘ
VXLAN Routing

Choose Fabric*
fab_setup_3

Connected networks ⓘ
L fab_setup_3-overlay-loopback-network ×

Extend to Physical Router* ⓘ
5a12-qfx5 × [Reconfigure Physical Routers](#)

Public Logical Router ⓘ

NAT

8. Select the physical router(s) to which you want to extend virtual networks or routed virtual networks to, from the Extend to Physical Router list.

A physical router provides routing capability to the logical router. You can extend the networks that you selected in step "7" on page 268 to multiple physical routers.

9. After you extend the networks to the physical router, you can reconfigure the physical router by clicking **Reconfigure Physical Router**. See [Figure 89 on page 269](#).

Step-by-Step Procedure


After you click **Reconfigure Physical Router**, the Extend to Physical Routers setup page is displayed. See [Figure 91 on page 271](#).

Enter the following information for each router.

Figure 91: Extend to Physical Router

Extend to Physical Routers setup

R cswj-0dc4-0015 cswj-0dc4-0016

Routed Virtual Network*	Routed Interface IP Address*	Routing Protocol*	
VN-BLUE		eBGP	
Local ASN ⓘ	Peer ASN*	Peer IP*	
Import Route Policies	Export Route Policies	md5 Authentication Key (optional) ⓘ	<input type="checkbox"/> BFD ⓘ

[+ Add](#)

[Cancel](#) [Submit](#)

NOTE: Loopback routed virtual network can only be configured with eBGP routing protocol.

- a. Select the VN-BLUE routed virtual network from the Routed Virtual Network list.

If you select a loopback routed virtual network from the Routed Virtual Network list, the Routing Protocol list is disabled. The default protocol selected is **eBGP**. For example, see [Figure 92 on page 272](#).

Starting in Contrail Networking Release 2005, you can configure BFD with eBGP, static Route, OSPF, and PIM protocols.

NOTE: If you use an MX240, MX480, MX960, or MX10003 device as border gateway devices, you can only configure eBGP and static routing protocols on the border gateway devices.

- b. Enter IP address in the Routed Interface IP Address field.

The routed IP address that you enter must be from the subnet that you provided while creating the routed virtual network.

If you have selected a loopback routed virtual network from the Routed Virtual Network list, the Loopback IP Address field is displayed instead of the Routed Interface IP Address field. In this case, the loopback IP address that you enter must be from the loopback subnet that you provided during the fabric onboarding process.

Figure 92: Selecting a Loopback Routed Virtual Network

The screenshot shows a configuration form with the following fields and options:

- Routed Virtual Network***: A dropdown menu with the value "fab_setup_3-overlay-l..." selected.
- Loopback IP Address**: A text input field with the placeholder "Enter valid IP".
- Routing Protocol***: A dropdown menu with "eBGP" selected.
- Local ASN** (with a help icon): A text input field with the placeholder "Enter Local ASN".
- Peer ASN***: A text input field with the placeholder "Enter Peer ASN".
- Peer IP***: A text input field with the placeholder "Enter valid IPs".
- Import Route Policies**: A dropdown menu.
- Export Route Policies**: A dropdown menu.
- md5 Authentication Key (optional)**: A text input field.
- BFD** (with a help icon): An unchecked checkbox.
- MultiHop Options**: An unchecked checkbox.
- + Add**: A blue button at the bottom left.

- c. Select **eBGP**, **Static Routes**, **OSPF**, or **PIM** routing protocols from the Routing Protocol list.

NOTE: If you use an MX240, MX480, MX960, or MX10003 device as border gateway devices, you can only configure eBGP and static routing protocols on the border gateway devices.

- **Step-by-Step Procedure**

Follow these steps if you have either selected **eBGP** as the routing protocol ([Figure 91 on page 271](#)) or if you have selected a loopback routed virtual network from the Routed Virtual Network list ([Figure 90 on page 270](#)):

- i. Enter Autonomous System (AS) number in the Local ASN field.

This is the logical router (VRF routing instance) AS number. If you do not provide an AS number, the overlay iBGP AS number provided during fabric onboarding will be used.

- ii. Enter peer AS number in the Peer ASN field.

This is the AS number of the third-party device.

- iii. Enter the IP address of the unmanaged third-party device in the Peer IP field.

NOTE: The IP address that you enter must be from the subnet that you provided while creating the routed virtual network.

- iv. Select import routing policies that you want to apply to this eBGP session from the Import Route Policies list.

NOTE: You can add routing policies from the **Overlay>Routing>Routing Policies** page of the Contrail Command UI.

- v. Select export routing policies that you want to apply to this eBGP session from the Export Route Policies list.

- vi. (Optional) Enter md5 authentication key in the md5 Authentication Key field.

You can enter the key in plain text or in encrypted format that starts with $\$$.

- vii. (Optional) Select **BFD** check box to enable bidirectional forward detection.

After you enable BFD, provide detection time and multiplier information in the **Interval (ms)** and **Multiplier** fields respectively.

- viii. If you have selected another physical router in step ["8" on page 270](#), click the next tab to enter information.

Repeat steps ["9.a" on page 271](#) through ["9.c" on page 272](#).

ix. Click **Submit** to submit configuration information.

- **Step-by-Step Procedure**

Enter the following information if you have selected **Static Routes** as the protocol:

NOTE: In order to select **Static Routes** as the protocol, you must create an interface route table.

Follow these steps to create an interface route table:

- Navigate to **Overlay>Routing>Interface Route Tables** and click **Add**.

The Create Interface Route Tables page is displayed.

- Enter a name for the interface route table in the Name field.
- Click **Add** to add Prefix and Community information.
- (Optional) Click **Add** to define another prefix list.
- Click **Create** to create the interface route table.

i. Select the interface route table from the Routing Table list. See [Figure 93 on page 274](#).

The interface route table lists the prefixes that can be used in the static route configuration.

ii. Enter next hop IP address in the Remote Next Hop Address field.

iii. Click **Submit** to submit configuration information.

Figure 93: Static Route Routing Protocol

The screenshot shows a configuration form for a static route. It contains the following fields and controls:

- Routed Virtual Network***: A dropdown menu.
- Routed Interface IP Address***: A text input field with the placeholder text "Enter valid IP".
- Routing Protocol***: A dropdown menu currently showing "Static Routes".
- Routing Table***: A dropdown menu.
- Remote Next Hop Address***: A text input field.
- A trash icon is located to the right of the Routing Protocol dropdown.
- A "+ Add" button is located at the bottom left of the form.

- **Step-by-Step Procedure**

Enter the following information if you have selected **OSPF** as the routing protocol:

NOTE: You cannot configure OSPF routing protocol on the border gateway devices if you use MX240, MX480, MX960, or MX10003 devices as border gateway devices.

- i. Enter the area ID in the Area ID field. See [Figure 94 on page 276](#).
The area ID is provided in *x.x.x.x* format.
- ii. Select any one of the following area type.
 - NSSA Area—An NSSA Area or a not-so-stubby area allows external routes to be flooded within the area.
 - Regular Area—Regular areas are areas that are not backbone areas (Area ID: 0.0.0.0), Stub areas, or NSSA areas.
 - Stub Area—Stub areas are areas through which autonomous system (AS) external advertisements are not flooded.

For more information, see [Configuring OSPF Areas](#).

NOTE: If the area ID that you enter is 0.0.0.0, the Area Type field is disabled.

- iii. Enter hello interval in the Hello Interval [Seconds] field. The default value is 10 seconds.
- iv. Enter dead interval in the Dead Interval [Seconds] field. The default value is 40 seconds.
- v. Select import routing policies that you want to apply to this OSPF session from the Import Route Policies list.

NOTE: You can add routing policies from the **Overlay>Routing>Routing Policies** page of the Contrail Command UI.

- vi. Select export routing policies that you want to apply to this OSPF session from the Export Route Policies list.
- vii. (Optional) Enter md5 authentication key in the md5 Authentication Key field.
You can enter the key in plain text or in encrypted format that starts with *\$?*.
- viii. (Optional) Select **BFD** check box to enable bidirectional forward detection.

After you enable BFD, provide detection time and multiplier information in the **Interval (ms)** and **Multiplier** fields respectively.

- ix. Select **Redistribute Loopback** to advertise loopback IP address in an area type.
Redistribute loopback IP address can be selected only for one OSPF area type.
- x. Select **Originate Summary LSA** check box to flood link-state advertisement (LSA) summary in to the selected area type.
This check box is disabled if
 - the area ID that you enter is 0.0.0.0.
 - you do not define an area type.
- xi. Click **Submit** to submit configuration information.

Figure 94: OSPF Routing Protocol

- **Step-by-Step Procedure**

Enter the following information if you have selected **PIM** as the routing protocol:

NOTE: You cannot configure PIM routing protocol on the border gateway devices if you use MX240, MX480, MX960, or MX10003 devices as border gateway devices.

- i. Enter rendezvous point (RP) IP address in the RP IP Address field. See [Figure 95 on page 277](#).

The RP IP address is the IP address of the router that receives multicast traffic. This IP address is autopopulated for every physical router. If the same physical router is allocated an RP IP address when configuring another router virtual network, this RP IP Address field is autopopulated with the same RP IP address.

- ii. PIM Mode field is disabled. The default PIM mode is **Sparse-dense**.
- iii. (Optional) Select **BFD** check box to enable bidirectional forward detection.

After you enable BFD, provide detection time and multiplier information in the **Interval (ms)** and **Multiplier** fields respectively.

- iv. Select **Enable PIM on all interfaces** check box to enable PIM on all IRB interfaces of the VRF.
- v. Click **Submit** to submit configuration information.

Figure 95: PIM Routing Protocol

The screenshot shows a configuration form for the PIM Routing Protocol. It contains the following fields and controls:

- Routed Virtual Network***: A dropdown menu.
- Routed Interface IP Address***: A text input field with the placeholder "Enter valid IP".
- Routing Protocol***: A dropdown menu with "PIM" selected.
- RP IP Address* ⓘ**: A text input field with the placeholder "Enter valid IPs".
- PIM Mode***: A dropdown menu with "Sparse-dense" selected.
- BFD ⓘ**: A checkbox to enable bidirectional forward detection.
- Enable PIM on all interfaces**: A checkbox to enable PIM on all IRB interfaces.
- + Add**: A button to add a new configuration entry.

10. Enter the routed VXLAN ID in the VxLAN Network Identifier field,

11. Click **Create** to create the logical router.

Repeat steps "2" on page 267 through "11" on page 277 to create LR-Green logical router.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2008	Starting in Contrail Networking Release 2008, you can use MX240, MX480, MX960, or MX10003 device as a border leaf and spine device, to connect to third-party network devices.
2005	Starting in Contrail Networking Release 2005, you can also configure PIM and OSPF routing protocols between the logical router and the unmanaged third-party device.

2005	Starting in Contrail Networking Release 2005, you can connect a loopback routed virtual network to a logical router.
2005	Starting in Contrail Networking Release 2005, you can configure BFD with eBGP, static Route, OSPF, and PIM protocols.
2003	Starting in Contrail Networking Release 2003, you can use the Contrail Command user interface (UI) to connect the border gateway devices to third-party devices that are not managed by Contrail Networking.
2003	With Contrail Networking Release 2003, you select Routed from the VN Fabric Type field to create a routed virtual network.
2003	With Contrail Networking Release 2003, you can create a routed virtual port group from the Contrail Command UI. Select the Routed option button to create a routed virtual port group.

Configuring Storm Control on Interfaces

IN THIS SECTION

- [When Is a Traffic Storm Generated? | 279](#)
- [How Do You Recognize a Traffic Storm? | 279](#)
- [How Can You Use Storm Control Profiles to Manage a Traffic Storm? | 279](#)
- [Configuring Storm Control Profiles | 280](#)

Starting with Contrail Networking Release 1908, you can configure storm control on the access interfaces of a datacenter fabric managed by Contrail Networking. Storm control feature is supported in both greenfield and brownfield deployments with enterprise style configuration.

When Is a Traffic Storm Generated?

A traffic storm is generated when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own copy of the messages on the network. This, in turn, prompts further replications, creating a snowball effect. The network is suddenly flooded with packets, creating unnecessary traffic that leads to poor network performance or even a complete loss of network service. Storm control enables the switch to monitor traffic levels and to drop broadcast, multicast, and unknown unicast packets when a specified traffic level—called the *storm control level*—is exceeded, thus preventing packets from proliferating and degrading the LAN. As an alternative to having the switch drop packets, you can configure it to shut down interfaces or temporarily disable interfaces when the storm control level is exceeded.

How Do You Recognize a Traffic Storm?

To recognize a storm, you must be able to identify when traffic has reached an abnormal level. Suspect a storm when operations begin timing out and network response times slow down. Users might be unable to access expected services. Monitor the percentage of broadcast and unknown unicast traffic in the network when it is operating normally. This data can then be used as a benchmark to determine when traffic levels are too high. You can then configure storm control to set the level at which you want to drop broadcast and unknown unicast traffic.

How Can You Use Storm Control Profiles to Manage a Traffic Storm?

You can configure storm control on devices after Contrail Command is set up and all devices discovered. You attach storm control profile to a port profile and then apply the port profile to interfaces or virtual port groups. A port profile functions like a container that can support multiple port-related configurations, and allows you to apply those configuration by attaching them to the port profile. You can then apply the port profile on an interface or a virtual port group. In Contrail Networking Release 1908, you can attach only storm control profiles to port profiles.

You can define one storm control profile per port profile and one port profile per interface or virtual port group.

NOTE: Storm control profile feature is supported only on QFX5000 and QFX10000 series devices.

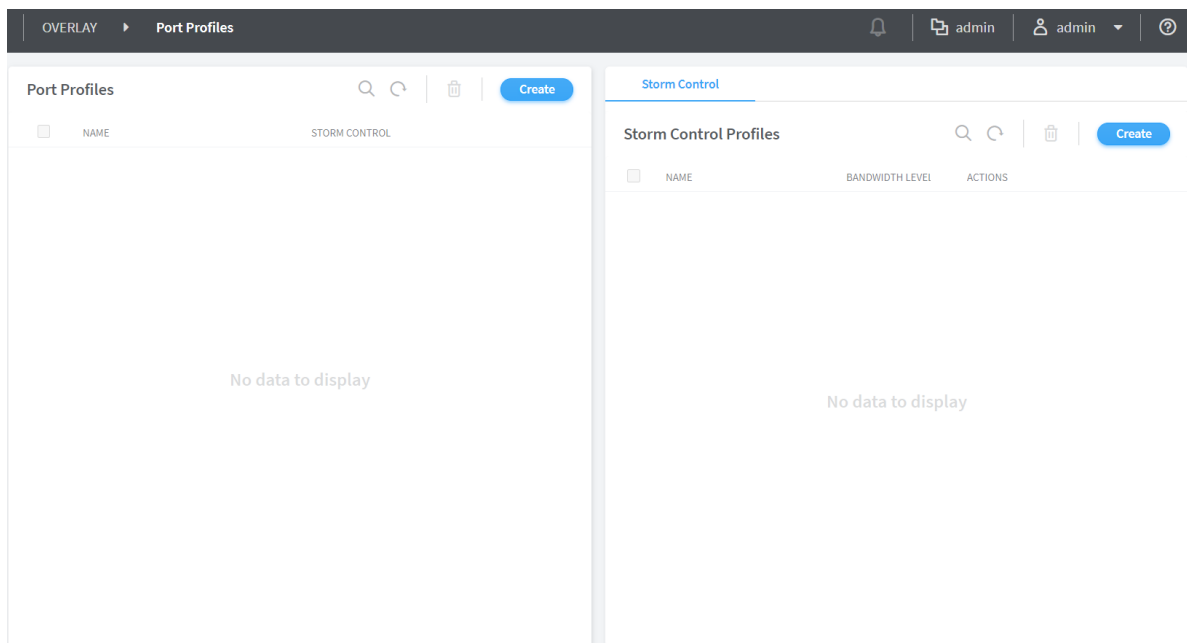
Configuring Storm Control Profiles

To enable storm control on an interface, you must first create a storm control profile, and then attach it to a port profile. You can then apply the port profile to an interface or a virtual port group (VPG). You can create port profiles and storm control profiles from the **Overlay > Port Profiles** page.

To create storm control profiles:

1. Click **Overlay > Port Profiles**.

Figure 96: Port Profiles



You must first create a storm control profile and then attach it to the port profile. You can attach the storm control profile to existing port profile or attach to a new port profile while creating it.

2. Click **Overlay > Port Profiles > Storm Control Profile > Create**.

Figure 97: Create Storm Control Profile

You must specify a storm control profile name and the threshold bandwidth percentage, after which the specified action is performed on the interface.

- **Bandwidth Level**— Enter the maximum value (in percentage) in the range 0–100. If the bandwidth utilized by broadcast, unknown unicast, or multicast (BUM) traffic exceeds this value, the action (default drop or configured **Interface shutdown**) specified in the storm control profile is applied on the interface. The default bandwidth level is 20%.
- **Actions**—Specify the action to be performed on the interface when the bandwidth utilization exceeds the specified bandwidth level. The default action is to drop the packets. For example, if you set a value 20% for the **Bandwidth Level** field, and specify an action **Interface Shutdown**, the interface shuts down when bandwidth utilization exceeds 20%.
- **Recovery timeout**—Specify a value in the range of 10–3600 for recovery timeout in seconds, after which the shut down interface needs to be brought up again. The default recovery timeout value is 600 seconds.
- **Traffic Types to Exclude**—Select the traffic types to be excluded from the storm control profile. By default, storm control is applied to all traffic types.

The multicast options **No multicast**, **No registered multicast**, and **No unregistered multicast** are mutually exclusive. That is, you can specify only one of these multicast options at a time.

3. Click **Create**.
4. Click **Overlay > Port Profiles > Create**.

Figure 98: Create Port Profile

The figure consists of two screenshots from the Contrail Command web interface. The top screenshot shows the 'Create Port Profile' page. The breadcrumb navigation is 'OVERLAY > Port Profiles > Create Port Profile'. The left sidebar shows a navigation menu with 'OVERLAY' selected. The main form contains the following fields and controls:

- Name***: A text input field with the placeholder 'Port Profile Name'.
- MTU**: A text input field with the placeholder 'bytes'.
- Storm Control Profile**: A dropdown menu with a '+ Add a profile' link next to it.
- Description**: A text area.
- Admin State**: A toggle switch set to 'On'.
- Enable LACP**: A checkbox with a 'LACP Type' dropdown and a 'LACP Interval' dropdown.
- Flow Control**: A checkbox.
- BFD Loop Protection**: A checkbox.
- QoS (CoS) Untrust Interface**: A checkbox.

At the bottom of the form are 'Create' and 'Cancel' buttons. The bottom screenshot shows the 'Create Profile Access' page. The breadcrumb navigation is 'OVERLAY > Port Profiles > Create Profile Access'. The left sidebar is not visible. The main form contains:

- Name***: A text input field.
- Storm Control Profile**: A dropdown menu with 'test-profile' selected.

At the bottom of the form are 'Create' and 'Cancel' buttons.

You must specify a port profile name and select a storm control profile from the profiles created in step "3" on page 282. You can attach only one storm control profile per port profile.

NOTE: If you want to delete a storm control profile, you must first remove it from the port profile. To delete a port profile, you must first detach the port profile from the VPG or the instance.

Starting with Contrail Networking Release 2008, additional port attributes are available for port profile objects including MTU, admin state, LACP, flow control, BPDU loop protection, and QoS (CoS) untrust interface. [Table 51 on page 283](#) provides detailed information on each attribute.

Table 51: Port Profile Attributes

Port Attributes	Description	Available Options
MTU	Sets maximum transmission unit in bytes.	Range: 256–9216
Description	Outlines interface description.	
Admin State	Changes the admin state of the interface.	
Enable LACP	Enables Link Aggregation Control Protocol (LACP).	<ul style="list-style-type: none"> • LACP Type—Active or Passive <ul style="list-style-type: none"> Active—Initiate transmission of packets. Passive—LACP packets are not exchanged with passive mode. • LACP Interval—Slow or Fast <ul style="list-style-type: none"> Slow—Receives packets every 30 seconds. Fast—Receives packets every second.
Flow Control	Enables flow control. Controlling the flow by pausing and restarting prevents buffers on the nodes from overflowing and dropping frames.	-

Table 51: Port Profile Attributes *(Continued)*

Port Attributes	Description	Available Options
BPDU Loop Protection	Increases the efficiency of STP, RSTP, and MSTP by preventing ports from moving into a forwarding state that would result in a loop opening up in the network.	-
QoS (CoS) Untrust Interface	Applies classifier based on 1P bits to all ethernet-switching ports.	-

Port profile objects enable users to customize configuration for devices and interfaces.

Click **Create**.

- After you create a port profile, you can assign it to interfaces or virtual port groups as shown in [Figure 99 on page 284](#).

Click **Overlay > Virtual Port Group > Create**

Figure 99: Attach Port Profile to VPG

The screenshot shows the 'Create Virtual Port Group' configuration interface. At the top, the breadcrumb navigation reads 'OVERLAY > Virtual Port Group > Create Virtual Port Group'. The main form includes the following elements:

- Virtual Port Group Name:** A text input field containing 'vpg-test'.
- VLAN Section:**
 - Tagged
 - VLAN id:** Input field with '1'.
 - TOR Port VLAN id:** Input field with '4094'.
 - Display Name:** Input field with 'vpg-test-4094'.
 - Auto Display Name
 - Network:** A dropdown menu.
- Security Groups:** A dropdown menu with 'default' selected.
- Port Profile:** A dropdown menu with 'test' selected, highlighted by a blue box.
- Fabric name:** A dropdown menu with 'fc7e14c5-91c6-491f-91...' selected.
- Available Physical Interface:** A section with a search bar and 'Add all' button. It displays 'No Physical Interface matching current criteria'.
- Assigned Physical Interface:** A section with a search bar and 'Remove all' button. It displays 'No Physical Interface matching current criteria'.
- Buttons:** 'Create' and 'Cancel' buttons at the bottom left.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2008	Starting with Contrail Networking Release 2008, additional port attributes are available for port profile objects including MTU, admin state, LACP, flow control, BPDU loop protection, and QoS (CoS) untrust interface.
1908	Starting with Contrail Networking Release 1908, you can configure storm control on the access interfaces of a datacenter fabric managed by Contrail Networking.

RELATED DOCUMENTATION

| [Configuring Virtual Port Groups](#) | 238

Creating Port Profiles, Storm Control Profiles, sFlow Profiles, or Telemetry Profiles by Cloning

Starting from Contrail Networking Release 2003, you can create a new port profile, storm control profile, sFlow profile, or a telemetry profile by cloning an existing one. You can modify the parameters in the cloned profile and save it. Cloning helps you quickly create a profile, without going through the process of defining each and every parameter. You can choose to modify those parameters that need to be modified.

NOTE: You cannot clone the default entities or profiles. However, you can create clones of an entity or a profile that you created earlier.

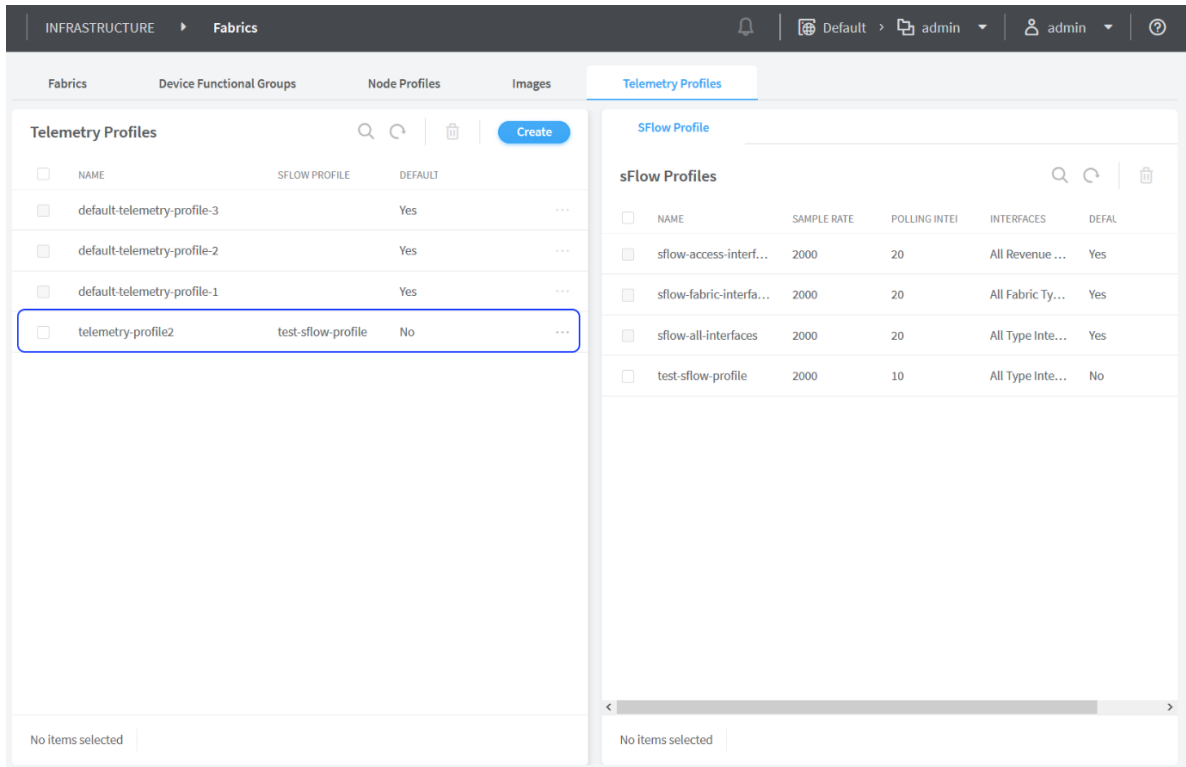
This topic shows how to clone a telemetry profile from Contrail command.

Follow these steps to clone a telemetry profile:

1. Navigate to **Infrastructure** > **Fabrics** and click the **Telemetry Profiles** tab. Follow the steps in [Configuring Contrail Insights Flows by Assigning Telemetry and sFlow Profiles to Devices](#) and create a telemetry profile.

The newly created telemetry profile is displayed as highlighted in [Figure 100 on page 286](#).

Figure 100: Telemetry Profile



2. Click the clone icon as highlighted in [Figure 101](#) on page 287.

Figure 101: Clone a Profile

The screenshot displays the 'Telemetry Profiles' management interface. The left pane shows a table of existing profiles, and the right pane shows the configuration for a selected 'SFlow Profile'.

Telemetry Profiles Table:

NAME	SFLOW PROFILE	DEFAULT	Actions
default-telemetry-profile-3		Yes	...
default-telemetry-profile-2		Yes	...
default-telemetry-profile-1		Yes	...
telemetry-profile2	test-sflow-profile	No	Clone, Edit, Delete

SFlow Profile Configuration Table:

NAME	SAMPLE RATE	POLLING INTEI	INTERFACES	DEFAL
sflow-access-interf...	2000	20	All Revenue ...	Yes
sflow-fabric-interfa...	2000	20	All Fabric Ty...	Yes
sflow-all-interfaces	2000	20	All Type Inte...	Yes
test-sflow-profile	2000	10	All Type Inte...	No

You can see that the name of the existing profile is prefixed with *Copy of* as highlighted in [Figure 102 on page 288](#). You can rename the profile if required.

Figure 102: Modify Parameters

INFRASTRUCTURE > Fabrics > Clone Telemetry Profile

Profile Name*

Copy of telemetry-profile2

sFlow Profile

test-sflow-profile [Create New](#)

test-sflow-profile Details	
Profile Name	test-sflow-profile
Sample Rate	2000
Sampling Direction	ingress
Polling Rate	10
Adaptive Sample Rate	300
Enabled Interfaces	all

[Create](#) [Cancel](#)

3. Modify the required parameters and click **Create**. The cloned profile is displayed as shown in [Figure 103 on page 289](#).

Figure 103: Cloned Profile

The screenshot shows the Contrail Insights web interface. The top navigation bar includes 'INFRASTRUCTURE' and 'Fabrics'. The main content area is divided into two panels. The left panel, titled 'Telemetry Profiles', contains a table with the following data:

<input type="checkbox"/>	NAME	SFLOW PROFILE	DEFAULT	...
<input type="checkbox"/>	default-telemetry-profile-3		Yes	...
<input type="checkbox"/>	default-telemetry-profile-2		Yes	...
<input type="checkbox"/>	default-telemetry-profile-1		Yes	...
<input type="checkbox"/>	Copy of telemetry-profile2	test-sflow-profile	No	...
<input type="checkbox"/>	telemetry-profile2	test-sflow-profile	No	...

The right panel, titled 'SFlow Profile', contains a table with the following data:

<input type="checkbox"/>	NAME	SAMPLE RATE	POLLING INTEL	INTERFACES	DEFAL
<input type="checkbox"/>	sflow-access-interf...	2000	20	All Revenue...	Yes
<input type="checkbox"/>	sflow-fabric-interfa...	2000	20	All Fabric Ty...	Yes
<input type="checkbox"/>	sflow-all-interfaces	2000	20	All Type Inte...	Yes
<input type="checkbox"/>	test-sflow-profile	2000	10	All Type Inte...	No

At the bottom of each panel, it says 'No items selected'.

RELATED DOCUMENTATION

[Configuring Contrail Insights Flows by Assigning Telemetry and sFlow Profiles to Devices](#)

Configuring EVPN VXLAN Fabric with Multitenant Networking Services

Junos OS supports different ways to configure an EVPN VXLAN fabric with multitenant networking services:

- Fabric-wide significance of a VLAN ID or enterprise style configuration

In this mode, Contrail Networking ensures that every Layer 2 Service or VLAN ID in a fabric is unique, and that there is a 1:1 mapping between the VLAN ID (4K VLANs per fabric) and the Virtual Extensible LAN Network Identifier (VNI). In most cases, 4K bridge domains are more than sufficient

for any enterprise deployment. Hence, fabric-wide significance of a VLAN ID implies that any VLAN being provisioned in an EVPN VXLAN fabric maps to a VNI in a 1:1 ratio.

- Local significance of a VLAN ID or service provider style configuration

In some Junos OS devices like MX Series, the VLAN ID used to connect an endpoint on a physical port is independent of the VNI associated to the VLAN. This means that the same virtual network or a Layer 2 bridge domain identifier can have more than one VLAN ID (as access interface) attached to the same VNI. This is especially relevant when a Juniper EVPN VXLAN fabric is used in a VMWare vCenter environment, where different ESXI hosts might use distinct distributed virtual switches and but locally significant VLAN IDs.

Contrail Networking Release 1908 enables you to select enterprise style of configuration for the CRB-Access role on QFX Series switch-interfaces. Once configured you can modify the enterprise style setting to service provider style of configuration. However, you cannot modify the service provider style to enterprise style of configuration without having to recreate the fabric.

NOTE: Contrail Networking Release 1909 supports QFX10002-60C device running Junos OS Release 19.1R2 and later. QFX10002-60C device works only if enterprise style of configuration is enabled. To enable enterprise style of configuration, select the **VLAN-ID Fabric Wide Significance** check box when onboarding the QFX10002-60C device. For more information, see ["Create a Fabric" on page 7](#).

CRB-Access, CRB-Gateway, AR-Client, DC-Gateway, Route-Reflector, ERB-UCAST-Gateway, DCI-Gateway, lean, and PNF-Servicechain routing bridging roles are supported on the QFX10002-60C device. However, AR-Replicator role is not supported on Junos OS Release 19.2R2. For more information, see ["Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185](#).

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1909	Contrail Networking Release 1909 supports QFX10002-60C device running Junos OS Release 19.1R2 and later.
1908	Contrail Networking Release 1908 enables you to select enterprise style of configuration for the CRB-Access role on QFX Series switch-interfaces.

RELATED DOCUMENTATION

[Create a Fabric | 7](#)

[Flexible Ethernet Services Encapsulation](#)

Edge-Routed Bridging for QFX Series Switches

IN THIS SECTION

- [Benefits of ERB | 292](#)
- [Optimization of IRB Interfaces Creation in ERB Switches | 292](#)

The edge-routed bridging (ERB) for QFX series switches feature configures the inter-VN unicast traffic routing to occur at the leaf (ToR) switches in an IP CLOS with underlay connectivity topology. The ERB feature introduces the **ERB-UCAST-Gateway** and **CRB-MCAST-Gateway** roles. ERB is supported on the following devices running Junos OS Release 18.1R3 and later:

- QFX5110-48S
- QFX5110-32Q
- QFX10002-36Q
- QFX10002-72Q
- QFX10008
- QFX10016

Contrail Networking supports assigning physical roles and routing bridging (overlay) roles to a networking device like a switch. The roles define the routing and bridging responsibilities of the device in the data center. A device can have one physical role and one or more routing bridging roles. In releases prior to Contrail Networking Release 5.1, Contrail Networking supports centrally-routed bridging (CRB) roles on data center devices. In CRB, when you configure the logical router to allow traffic to flow between Ethernet virtual network instances, the routing occurs at the spine device. Traffic is routed from the leaf to the spine and back. IRB interfaces are configured in the overlay at each spine device to route traffic between virtual networks. Contrail Networking Release 5.1, supports the **ERB-UCAST-Gateway** role in which the routing occurs at the leaf switch. The IRB interfaces are configured at the leaf switch to enable unicast traffic routing at the leaf switch.

Traffic is routed in lesser hops when routed at the leaf switches. For example, consider two bare metal servers belonging to two separate VNs. Unicast traffic between the VNs are routed at the leaf switch and doesn't need to flow to the spine and back. Traffic is routed through the shortest path.

When you configure the **ERB-UCAST-Gateway** role on the leaf switches, it is recommended that you also configure the **CRB-MCAST-Gateway** role for multicast traffic on the corresponding spine devices. The **CRB-MCAST-Gateway** role is also supported. While unicast traffic can be routed at the leaf switches, multicast traffic routing still occurs at the spine devices. The existing **CRB-Gateway** role is capable of routing both unicast and multicast traffic at the spine devices. However, in ERB, if leaf switches route the unicast traffic, configuring the **CRB-Gateway** role on the spine is unnecessary since unicast traffic will never reach the spine device. Instead, you must configure the spine devices with the **CRB-MCAST-Gateway** role to route multicast traffic when required.

Benefits of ERB

- Traffic is routed through the shortest path.
- When you extend a logical router to a physical router, you can extend the logical router to leaf switches as well. Previously, logical routers could only be extended to the spine devices.

Optimization of IRB Interfaces Creation in ERB Switches

In an ERB topology with **ERB-UCAST-Gateway** role configured on ERB switches, when creating a logical router, an integrated routing and bridging (IRB) interface is created for each associated VN on all ERB switches. IRB interfaces are created regardless of a local port or VPG in the associated VNs. This leads to a lot of unnecessary IRB and VRF configurations associated to the logical router being pushed to the ERB switches.

Starting from release 2011, Contrail Networking checks if a VPG is attached to the VN in an ERB switch with **ERB-UCAST-Gateway** role configured, before pushing IRB interface configuration associated with the logical router to the switch. Only if a VPG is attached to the VN, the IRB interface configuration is pushed to the ERB switch. The routing instance is not created if there are no VPGs.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2011	Starting from release 2011, Contrail Networking checks if a VPG is attached to the VN in an ERB switch with ERB-UCAST-Gateway role configured, before pushing IRB interface configuration associated with the logical router to the switch.

RELATED DOCUMENTATION

[Edge-Routed Bridging Overlay Design and Implementation](#)

[Fabric Overview | 4](#)

[Create a Fabric | 7](#)

[Configuring Data Center Gateway | 220](#)

Activating Maintenance Mode on Data Center Devices

Starting with Contrail Networking Release 1909, you can activate maintenance mode on spine and leaf devices in a data center fabric. In maintenance mode, traffic flowing through the device is drained out or rerouted to other devices so that you can perform maintenance activity on the device like replace line cards or fix any issue on the device.

Prior to Contrail Networking Release 1909, devices were placed in maintenance mode only when performing hitless software upgrade.

Activating Maintenance Mode

To activate maintenance mode on a data center device in a fabric.

1. Navigate to the **Infrastructure > Fabrics** page in Contrail Command. A list of fabrics is displayed in the **Fabrics** tab.
2. Click the **Fabrics** tab and select a data center fabric. The list of devices connected in a spine and leaf topology and corresponding details of each device in the selected fabric is displayed. The roles assigned to the devices are also displayed.
3. Click ... on the right side of a fabric device.

4. Click **Activate Maintenance Mode**. A page requesting confirmation to activate maintenance mode is displayed.
5. Click **Confirm** to confirm activation of maintenance mode on the device.
Alternatively, click **Cancel** to cancel activating the maintenance mode.
6. Select the health check parameters for the device in the **Parameters** tab.
The health check parameters confirm that the device and the network as a whole are stable to activate maintenance mode. By default, if health check fails for a particular device, then maintenance mode is not activated. You can deselect the **Abort on health check failure** check box to continue activation on the device even if the health check fails.
7. Click **Next**. The **Testing** page appears.
The **Testing** page validates and displays the result of the health check on the device for the parameters selected previously in the **Parameters** tab. If health check fails for the selected parameters, then you can go back to the previous page by clicking **Previous** and either change the value of the health check parameter or disable the parameter altogether. You can perform this step multiple times until health check passes for the device or you are able to determine that performing maintenance on the device is feasible.

Alternatively, you can click **Previous** and deselect the **Abort on health check failure** check box in the **Parameters** tab to continue maintenance mode activation on the device even if health check fails.
8. Click **Next**. The **Activating** page appears and the device is placed in maintenance mode.
9. Click **Finish** to exit the wizard. The **Fabrics** page appears and status of the device is listed as under **Maintenance Mode**.

Deactivating Maintenance Mode

Once maintenance activity on a data center device is completed, you can deactivate maintenance mode on the device and bring it back online. To deactivate maintenance mode on a data center device.

1. Navigate to the **Infrastructure > Fabrics** page in Contrail Command. A list of fabrics is displayed in the **Fabrics** tab.
2. Select a data center fabric. The list of devices in the selected fabric is displayed. The roles assigned to the devices are also displayed.
3. Click the ... on the right side of a fabric device which is under maintenance mode.
4. Click **Deactivate Maintenance Mode**. A page requesting confirmation to deactivate maintenance mode is displayed.
5. Click **Confirm** to confirm deactivation of maintenance mode on the device.
Alternatively, click **Cancel** to cancel deactivating the maintenance mode.
6. Select the health check parameters for the device in the **Parameters** tab.

By default, if health check fails for a particular device, then maintenance mode is not deactivated. You can deselect the **Abort on health check failure** check box to continue deactivation on the device even if the health check fails.

7. Click **Next**. The **Deactivating** page appears.

The device is taken out of maintenance mode and this page validates and displays the result of the health check on the devices for the parameters selected previously in the **Parameters** tab.

8. Click **Finish** to exit the wizard. The **Fabrics** page appears displaying the status of the device as **Active**.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1909	Starting with Contrail Networking Release 1909, you can activate maintenance mode on spine and leaf devices in a data center fabric.

RELATED DOCUMENTATION

[Performing Hitless Software Upgrade on Data Center Devices | 326](#)

[Terminating Ongoing Fabric Jobs | 113](#)

Viewing the Network Topology

Starting with Contrail Networking Release 1907, the Contrail Command UI provides visual representation of the network topology. All devices within a fabric are displayed in a single view.

NOTE: Topology view is supported in Contrail Insights (formerly AppFormix) version 2.19.11 and later.

The Topology view supports basic manipulations such as dragging nodes, zooming in and out, fitting to view, in addition to having different layout visualizations. User-edited network layout is saved in the database so any change in network devices layout is preserved across sessions.

Topology view displays the following:

- Network devices
- Hosts
- Compute instances in hosts
- Edges connecting network devices and Contrail Networking vRouter hosts but not bare metal server (BMS) alone.

Three views are supported:

- Horizontal
- Vertical
- Radial

The topology heatmap shows network and server resources used in real-time or historically. For example, data center operators can see the bytes per second or the link utilization inside the network, or the CPU resources being consumed by a specific server or virtual machine.

Select **Infrastructure > Fabrics > <fabric name> > Topology View**. From the Display drop-down list, select an option to filter data viewed. Mouse over the nodes and interfaces in the Topology View for more detail.

To change views, select an icon in the vertical tool bar in the Topology View panel. Following are example views. For descriptions of the Summary view options, see [Table 52 on page 302](#).

Figure 104: Horizontal View

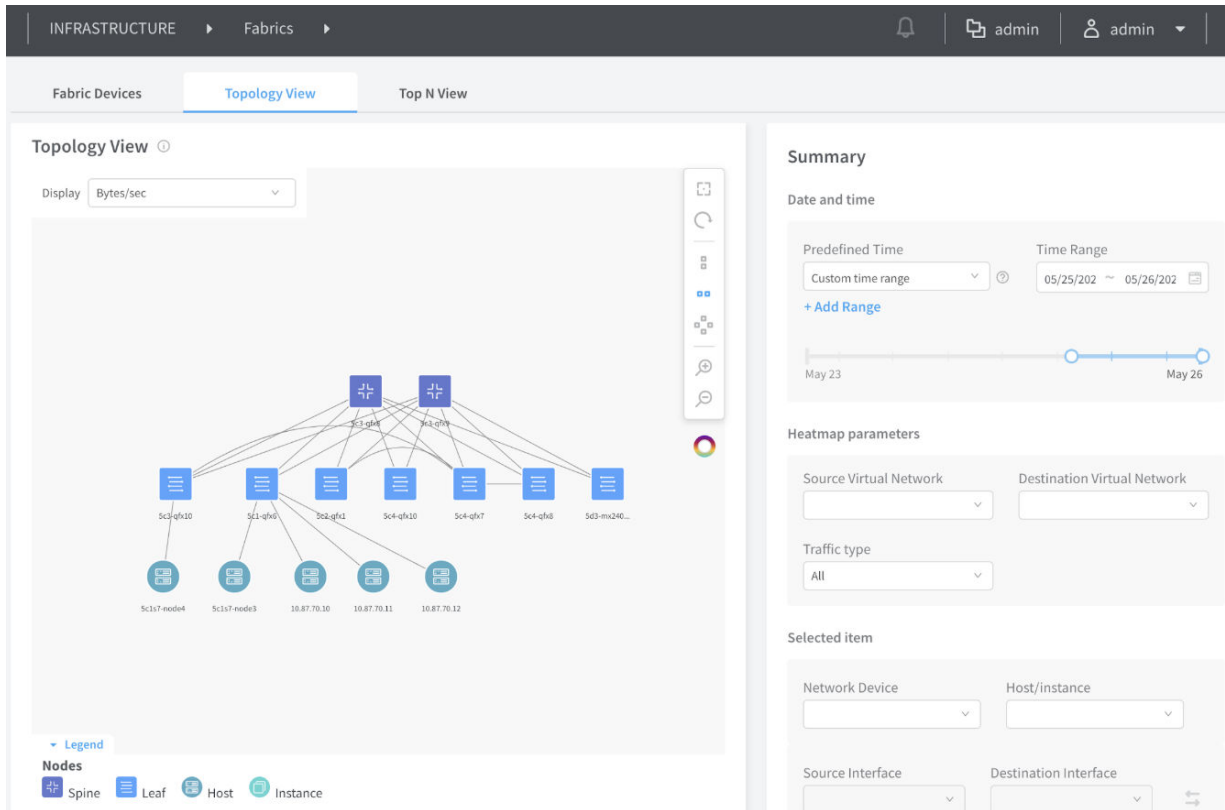


Figure 105: Horizontal View with Heatmap

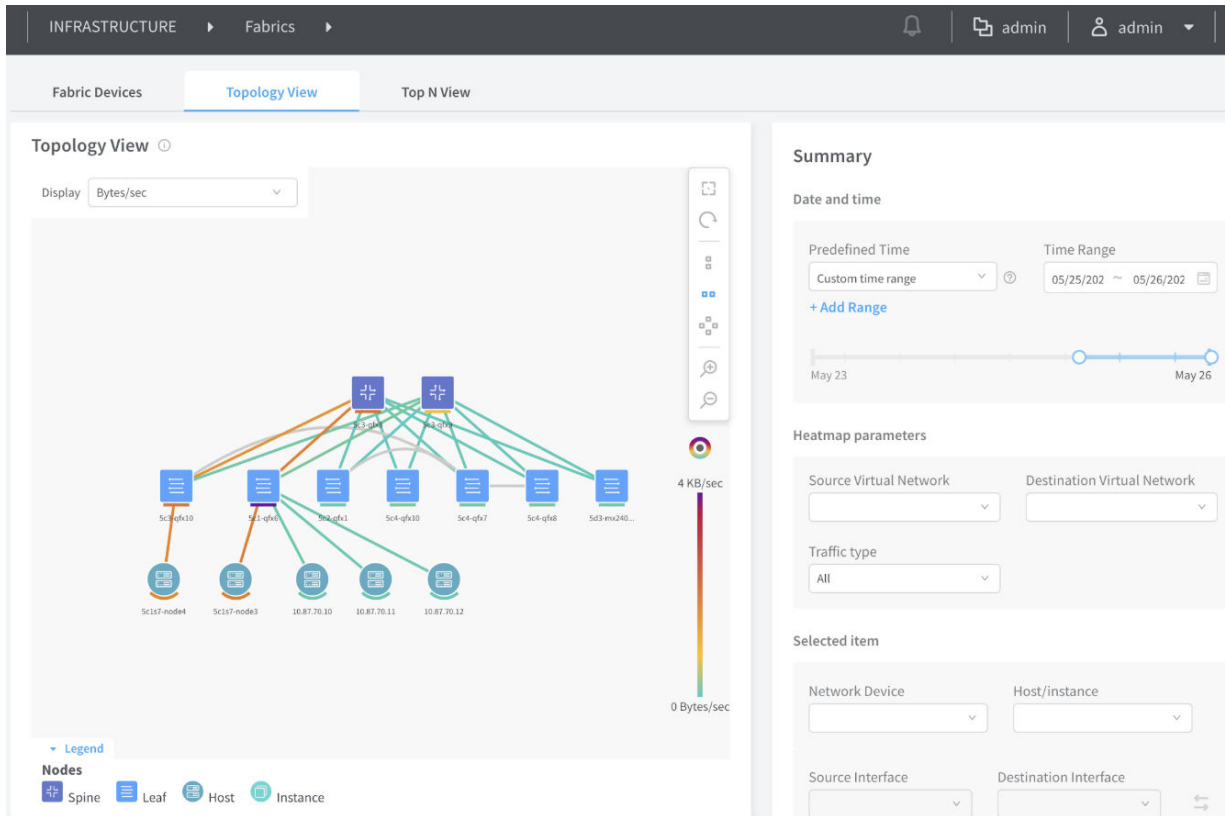


Figure 106: Horizontal View with VMs




Figure 107: Vertical View

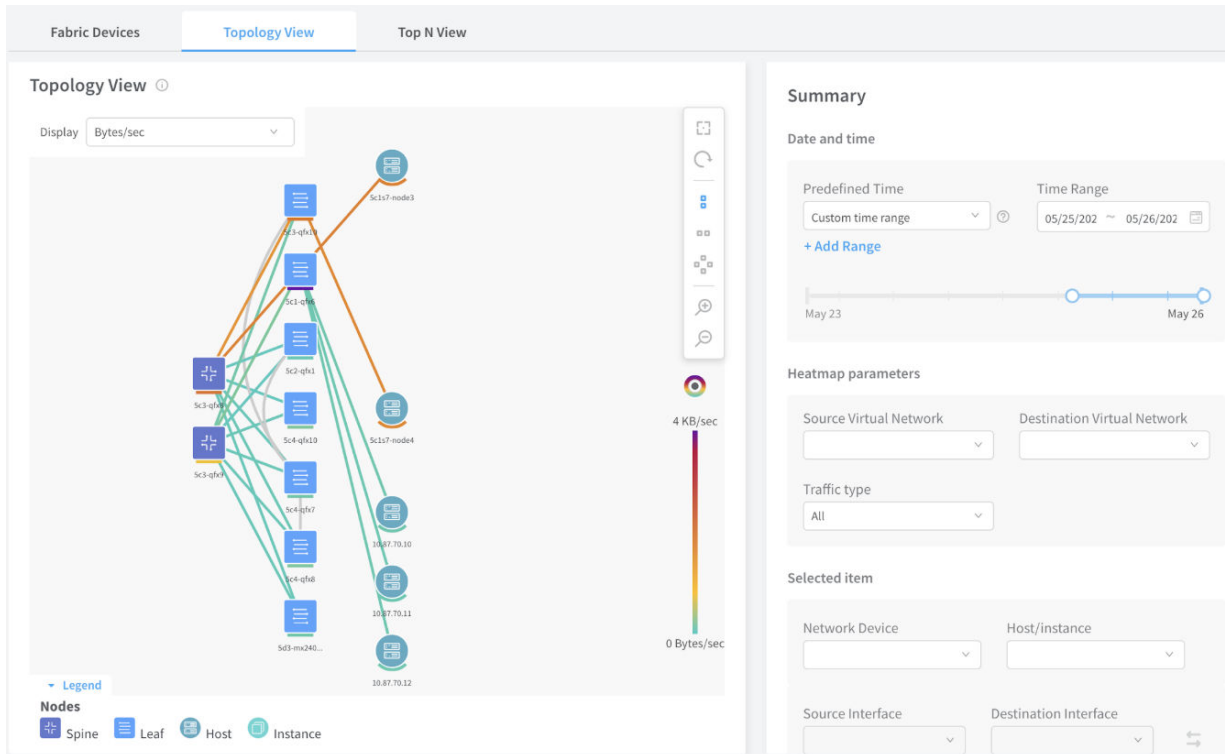


Figure 108: Vertical View with VMs

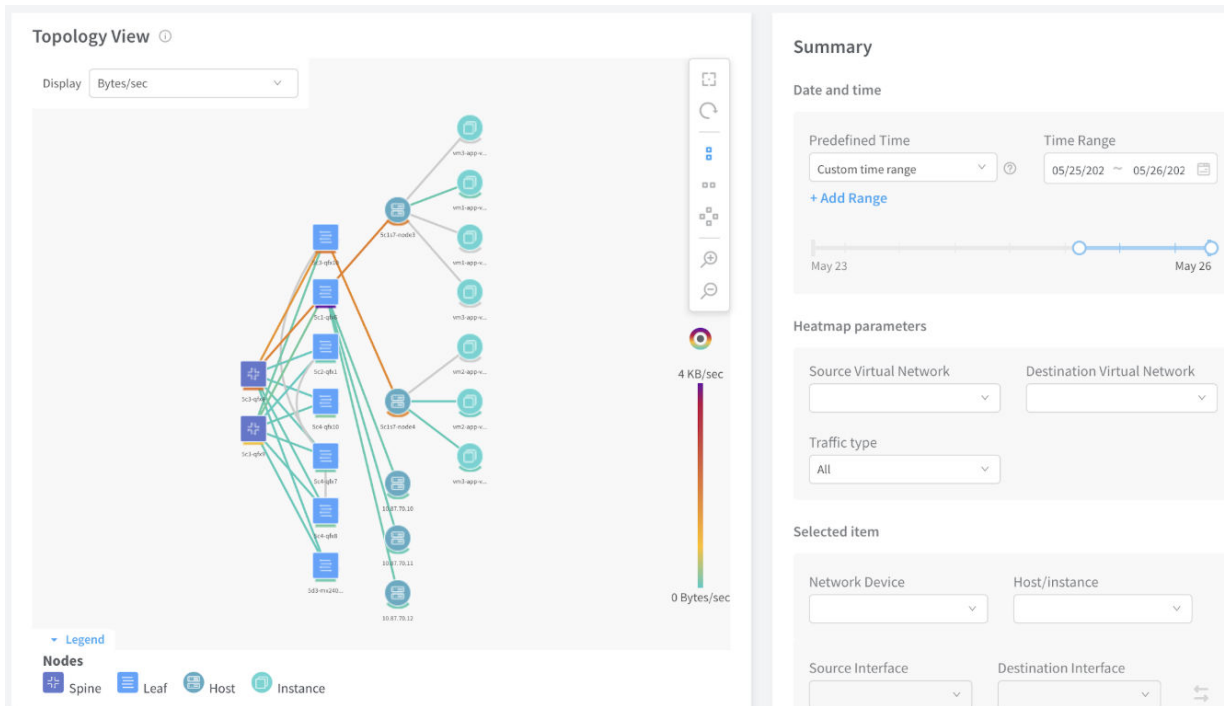


Figure 109: Radial View

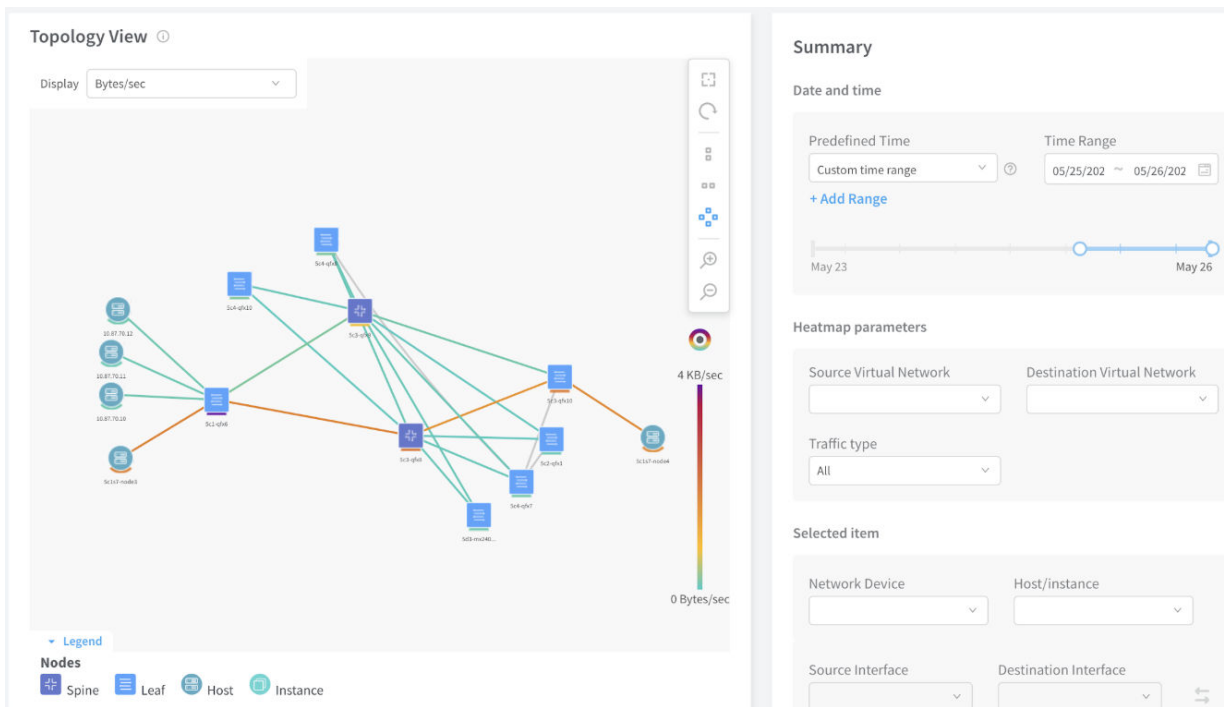


Figure 110: Radial View with VMs

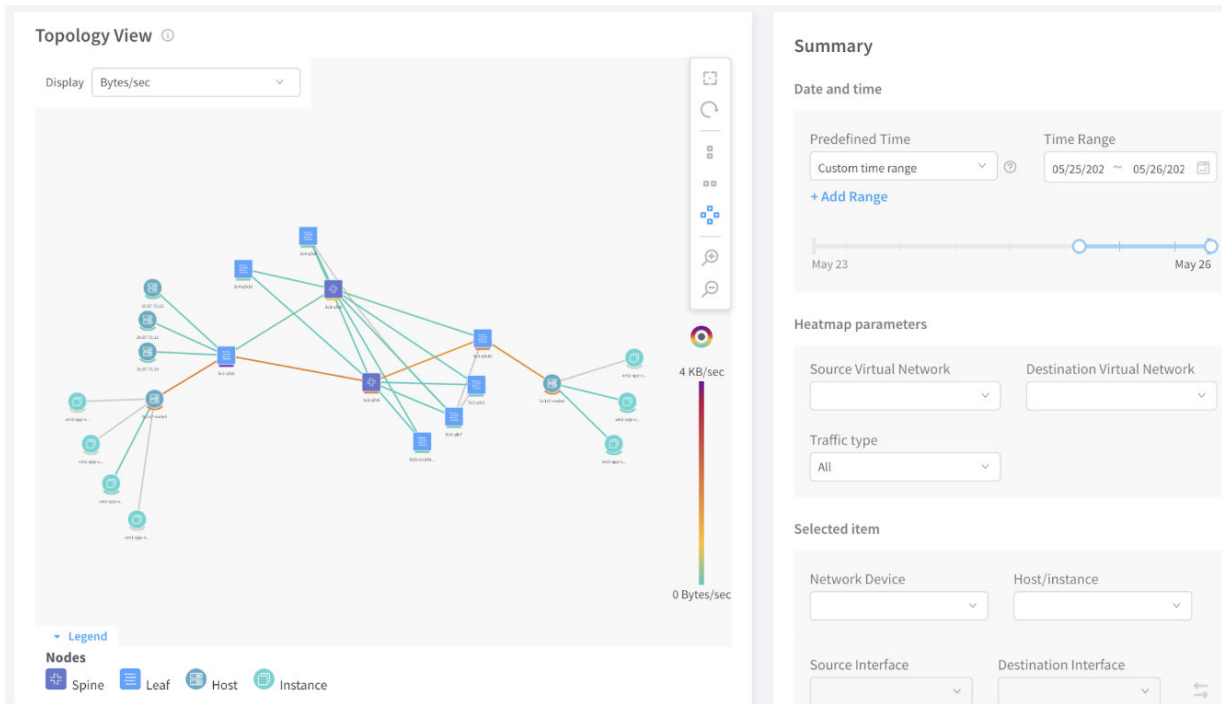


Table 52: Summary Descriptions

Summary Options	Description
Date and time	
Predefined Time	Select the period in the history for which data is to be displayed.
Time range	Use the calendar or type directly into the fields to select the desired start and end time. Additionally, you can select a time interval by dragging the mouse.
Heatmap parameters	
Source Virtual Network	Filter data with this source virtual network.
Destination Virtual Network	Filter data with this destination virtual network.

Table 52: Summary Descriptions (*Continued*)

Summary Options	Description
Traffic type	Filter data by traffic type, such as All, Multicast, or Other.
Selected item	
Network Device	Filter data passing through the network device.
Host/instance	Filter data with the host compute instance(s), such as Memory Usage or CPU Usage.
Source Interface	Filter the source interface on the selected network device.
Destination Interface	Filter the destination interface on the selected network device

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1907	Starting with Contrail Networking Release 1907, the Contrail Command UI provides visual representation of the network topology. All devices within a fabric are displayed in a single view.

RELATED DOCUMENTATION

| [Contrail Insights Flows in Contrail Command](#)

Viewing Hardware Inventory of Data Center Devices

In Contrail Networking Release 1909, you can view the hardware inventory of all data center devices deployed in a fabric using the Contrail Command user interface (UI). You can use the **Hardware Inventory** tab in Contrail Command to view the hardware inventory information. In releases prior to

Contrail Networking Release 1909, you had to use the `show chassis hardware` command on the CLI to view the hardware inventory of a data center device. The hardware inventory contains information about CPU, power supply, Flexible PIC Concentrators (FPCs), Physical Interface Cards (PICs), and so on installed in the router or switch chassis of the devices in the data center fabric. The hardware inventory information of a device is read and populated, when the fabric onboarding job is initiated after adding the device to a new or existing fabric; however, you can also view the hardware inventory information of the device in real-time.

To view the hardware inventory in the Contrail Command UI, perform the following steps:

1. Click a fabric in the **Infrastructure>Fabrics** page.

The **Fabric devices** page is displayed with a list of devices deployed in the fabric.

2. Click any device from the list to view the hardware inventory of the device.

The **Interfaces** page is displayed.

3. Click the **Hardware Inventory** tab.

The hardware inventory of the selected device is displayed. See [Figure 111 on page 304](#).

Figure 111: Hardware Inventory

MODULE	MODEL	MODEL NUMBER	PART NUMBER	VERSION	SERIAL NUMBER	DESCRIPTION
Chassis					AI41073718	QFX5120-32C
Pseudo CB 0						
Routing Engine 0		QFX5120-32C-CHAS	BUILTIN		BUILTIN	RE-QFX5120-32C
FPC 0		QFX5120-32C-CHAS	650-092914	REV 01	AI41073718	QFX5120-32C
CPU			BUILTIN		BUILTIN	FPC CPU
PIC 0		QFX5120-32C-CHAS	BUILTIN		BUILTIN	32X40G/32X100G-QSFP
Xcvr 0			740-032986	REV 01	QD513194	QSFP+-40G-SR4
Xcvr 1			740-032986	REV 01	QB500529	QSFP+-40G-SR4
Xcvr 2			740-032986	REV 01	QB160020	QSFP+-40G-SR4
Xcvr 3			740-032986	REV 01	QC260324	QSFP+-40G-SR4
Xcvr 4			740-061003	REV 01	1RC6232303N	QSFP28-END-100G-4x25G-DAC-B0...
Xcvr 5			740-061003	REV 01	1RC62324031	QSFP28-END-100G-4x25G-DAC-B0...
Xcvr 6			740-032986	REV 01	QG5001XG	QSFP+-40G-SR4
Power Supply 0		QFX520048Y-APSU-AO	640-084631	REV 01	1GG18240248	QFX520048Y-650W-AC-AFO

4. (Optional) Click the **Fetch** button if there is no inventory information available or you want to see an updated inventory information. See [Figure 111 on page 304](#).

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1909	In Contrail Networking Release 1909, you can view the hardware inventory of all data center devices deployed in a fabric using the Contrail Command user interface (UI).

Viewing Configuration of Devices Deployed in Contrail Fabric

Starting with Contrail Networking Release 2003, you can use the Contrail Command user interface (UI) to view the configuration information of each devices deployed in a fabric. The device configuration is information related to interface, encryption, physical role or routing-bridging roles assigned to the device, and so on. To view the configuration information of a fabric device, you have to perform the following steps in the Contrail Command UI:

1. Click a fabric in the **Infrastructure>Fabrics** page.
The **Fabric devices** page is displayed with a list of devices deployed in the fabric.
2. Select one or more devices from the list and click the **View Configuration** button. See [Figure 112 on page 305](#).

Figure 112: View Configuration Button in the Contrail Command UI

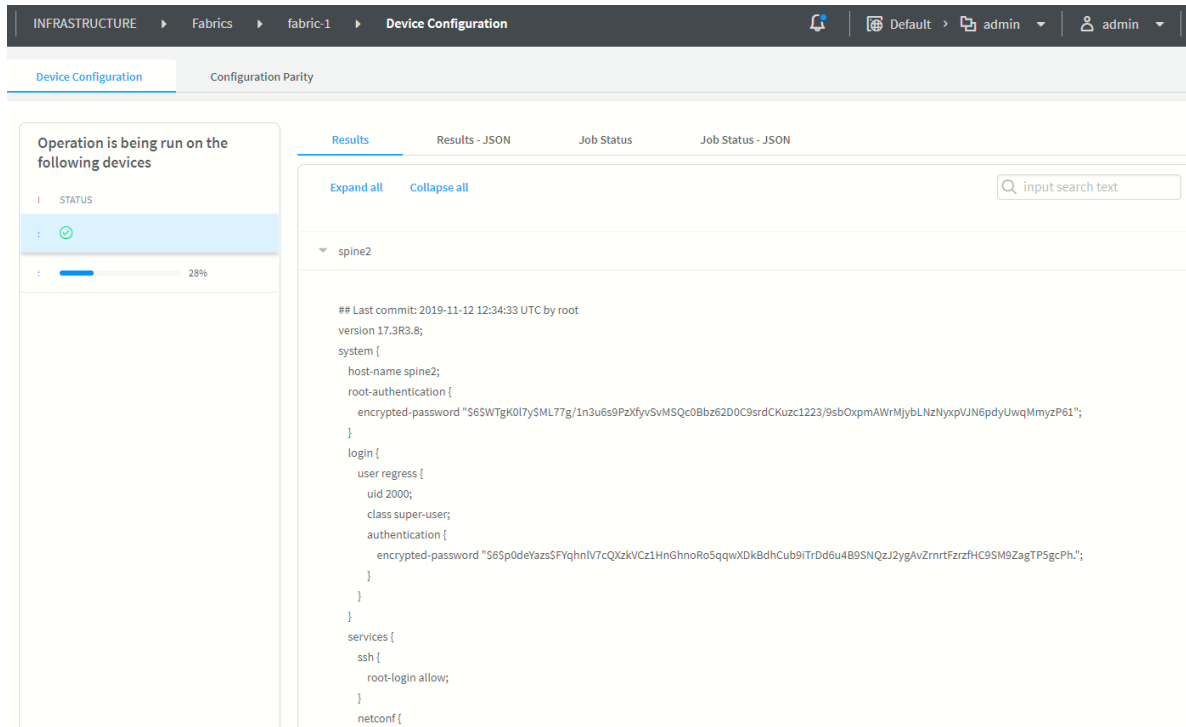
The screenshot shows the Contrail Command UI interface. At the top, there is a breadcrumb navigation: **INFRASTRUCTURE > Fabrics > fabric-1**. A notification bell icon is on the right. Below the breadcrumb, there is a toolbar with various icons and a blue **Action** dropdown menu. A tooltip labeled **View Configuration** is positioned over the **Action** menu. The main content area displays a table titled **Fabric devices**. The table has columns: **STATUS**, **NAME**, **MANAGEMENT**, **LOOPBACK IF**, **VENDOR NAME**, **PRODUCT NAME**, **ROLE**, **ROUTING BRIDGE**, and **INTERFACES**. Two rows are visible, both with a green **ACTIVE** status. The first row is for **spine2** (Juniper vqfx-10000, spine role, Route-Reflect bridge, 13 interfaces). The second row is for **spine1** (Juniper vqfx-10000, spine role, DC-Gateway bridge, 13 interfaces). Each row has a blue checkmark in the **STATUS** column and a three-dot menu icon in the **INTERFACES** column.

STATUS	NAME	MANAGEMENT	LOOPBACK IF	VENDOR NAME	PRODUCT NAME	ROLE	ROUTING BRIDGE	INTERFACES
ACTIVE	spine2	10.XX.XX.XX		Juniper	vqfx-10000	spine	Route-Reflect	13
ACTIVE	spine1	10.XX.XX.XX		Juniper	vqfx-10000	spine	DC-Gateway	13

3. The **Device Configuration** page is displayed where you can see two panels. The **Operation is being run on the following devices** panel on the left side displays the device names and the **Results** panel on the right side displays the configuration result of each device. The **View Configuration** operation

takes some time to complete. You have to wait until the operation is completed. On the left panel, a progress bar next to each device displays the progress of the operation on each device. See [Figure 113 on page 306](#).

Figure 113: View Configuration Operation In Progress



4. You can view the configuration details of each device in the **Results** panel once the **View Configuration** operation is successfully completed.

In the **Results** panel, click **Expand all** to view the entire configuration results generated for a device. Click **Collapse all** to hide the displayed results.

You can also perform a search operation on the expanded results by entering a specific text or keyword in the **input search text** field. Information about the results matching your text or keyword is displayed in the following places:

- In the left panel, next to the device names.
- In the **Results** panel, next to the device names.
- In the **Results** panel, next to the **input search text** field.
- In the expanded **Results** panel, highlighting the text or keyword.

See [Figure 114 on page 307](#).

Figure 114: Results of the View Configuration Operation Is Displayed

INFRASTRUCTURE > Fabrics > fabric-1 > Device Configuration

Configuration Parity

Operation is being run on the following devices

STATUS	Count
✓	2
✓	2

Results - JSON

Expand all Collapse all 4 RESULTS IN 2 DEVICES

spine2 2 RESULTS

```
## Last commit: 2019-11-12 12:34:33 UTC by root
version 17.3R3.8;
system {
  host-name spine2;
  root-authentication {
    encrypted-password "$6$WTgK0l7ySML77g/1n3u6s9PzXfyvSvMSQc0Bbz62D0C9srdCKuzc1223/9sbOxpmAWrMjybLNzNyxpUJN6pdyUwqMmyzP61";
  }
  login {
    user regress {
      uid 2000;
      class super-user;
      authentication {
        encrypted-password "$6$P0deYazsSFYqhnIV7cQXzKVCz1HnGhnoRo5qqwXdkBdhCub9ITrDd6u4B9SNQzrJ2ygvAvZrmtFzrzfHC9SM9ZagTP5gcPh.";
      }
    }
  }
  services {
    ssh {
      root-login allow;
    }
  }
}
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2003	Starting with Contrail Networking Release 2003, you can use the Contrail Command user interface (UI) to view the configuration information of each devices deployed in a fabric.

RELATED DOCUMENTATION

[Detecting and Managing Manual CLI Configuration Changes](#) | 308

Detecting and Managing Manual CLI Configuration Changes

IN THIS SECTION

- [Detecting a CLI Change | 308](#)
- [Accept, Ignore, or Reject a CLI Change | 312](#)

Contrail Networking Release 2003 supports the detection of manual CLI configuration changes. You can either accept these manual changes as part of the configuration, or you can reject the CLI changes and remove the manual change from the configuration. Starting in Contrail Networking Release 2008, you can also ignore manual CLI configuration changes.

For example, consider a scenario where a user logs in to the command line interface of a fabric device, makes changes to the existing configuration, and commits the configuration. Contrail Command detects such manual CLI configuration changes. You can then choose to accept, ignore, or reject the change.

NOTE: When you make configuration changes by using the command line interface, ensure that you do not use - ; { } [] , in the description field. Using these characters would cause the CLI detection workflow to fail.

You can view the following information before you accept, ignore, or reject the new CLI configuration change.

- The user who made the change.
- The time the change was made.
- The actual change in configuration.

These topics provide instructions to detect, accept, ignore, and reject manual CLI configuration changes.

Detecting a CLI Change

Contrail Command can detect a manual CLI change when

- You run the View Configuration job.

For steps to view device configuration, see "[Viewing Configuration of Devices Deployed in Contrail Fabric](#)" on page 305.

- You upgrade a device image.

For more information on upgrading a device image, see "[Image Management](#)" on page 54.

- You run a fabric job such as Reconfigure Roles.

For more information on reconfiguring device roles, see "[Reconfigure Roles](#)" on page 96.

As an example, consider that you have made a CLI configuration change on a fabric device. Follow these steps to detect the manual CLI configuration change when you run the View Configuration job.

1. Navigate to **Infrastructure>Fabrics**.


The Fabrics page is displayed.

2. Click the name of the fabric to view fabric devices.

The Fabric Devices page is displayed.

3. To view configuration of a device, select the check box next to the device name and click the **View Configuration** icon. See [Figure 115 on page 309](#).

Figure 115: View Configuration Icon



The Device Configuration tab is displayed.

The device that you have selected to view configuration for is listed in the **Operation is being run on the following devices** table. The status of the view configuration job for the device is displayed in the Status column next to the device name.

If there is a manual change in configuration, the job status is **Review Changes**. If there is no change, the status is a green tick mark. See [Figure 116 on page 310](#).

Figure 116: View Configuration Job Status

The screenshot displays the Configuration Parity interface. On the left, a table lists devices and their status:

DEVICE	STATUS
5d9-qfx3	Review Changes
5d9-qfx5	Active

The main area shows the configuration for device 5d9-qfx3, which is expanded. The configuration text is as follows:

```
## Last commit: 2020-08-20 23:39:35 UTC by root
version 20200730.082005.4_builder.r1127248;
system {
  host-name 5d9-qfx3;
  root-authentication {
    encrypted-password "$6$VJWt04nB$X5.cvcJ0J8egRRR6Q5zfxYTxep8t4j1WmYpErads77pa1k6AExjaYSvPRNLcqmJDS8g-UrzEVAZe9YFMhgVlw0";
  }
}
services {
  ssh {
    root-login allow;
  }
  telnet;
  netconf {
    ssh;
  }
}
syslog {
  user * {
    any emergency;
  }
}
file messages {
  any notice;
  authorization info;
}
file interactive-commands {
  interactive-commands any;
}
```

A notification at the bottom right states: "Job Template job for Show current or rollback configuration has finished".

After the view device configuration job is completed, navigate to the Fabric devices page to view the status of all devices. A device is in the **Changed** status when there are new manual CLI changes made. Devices with no manual CLI changes are in the **Active** status. See [Figure 117 on page 311](#).

Figure 117: Status of a Fabric Device

Fabric devices										Action
STATUS	NAME	MANAGEMENT I	LOOPBACK IP	VENDOR NAME	PRODUCT NAMI	ROLE	ROUTING BRIDI	INTERFACES		
▶ <input type="checkbox"/> ● ACTIVE	5d9-qfx5	XX.XX.XX	XX.XX.XX	Juniper	qfx5110-48s-4c	leaf	CRB-Access	6	...	
▶ <input type="checkbox"/> ● CHANGED	5d9-qfx3	XX.XX.XX	XX.XX.XX	Juniper	qfx5120-48t-6c	leaf	AR-Client CRB-Access 1 more	52	...	
▶ <input type="checkbox"/> ● ACTIVE	5d9-qfx4	XX.XX.XX	XX.XX.XX	Juniper	qfx10002-36q	spine	AR-Replicator CRB-Gateway 4 more	5	...	
▶ <input type="checkbox"/> ● ACTIVE	5d9-qfx2	XX.XX.XX	XX.XX.XX	Juniper	qfx5120-48y-8c	leaf	AR-Client CRB-Access	5	...	

No items selected

For more information on reviewing changes, see ["Accept, Ignore, or Reject a CLI Change"](#) on page 312.

NOTE: After you upgrade a device image or after you run the reconfigure roles fabric job, navigate to the Fabric Devices page to determine if the device is in the **Changed** status. If the device is in the **Changed** status, Contrail Command has detected that the device has a new CLI configuration change.

Accept, Ignore, or Reject a CLI Change

After Contrail Command has detected a CLI configuration change, you can accept, ignore, or reject CLI changes.

Follow these steps to accept, ignore, or reject a CLI change by using the Contrail Command UI.

1. Navigate to **Infrastructure>Fabrics**.

The Fabrics page is displayed.

2. Click the name of the fabric to view fabric devices.

The Fabric Devices page is displayed. The device with new CLI changes is listed in the Fabric Devices page and is in **Changed** status.

3. From the Fabric Devices page, select the device in the **Changed** status by selecting the check box next to the device name.

You can select more than one device.

4. Click the **View Configuration** icon as shown in [Figure 115 on page 309](#).

The Device Configuration tab is displayed.

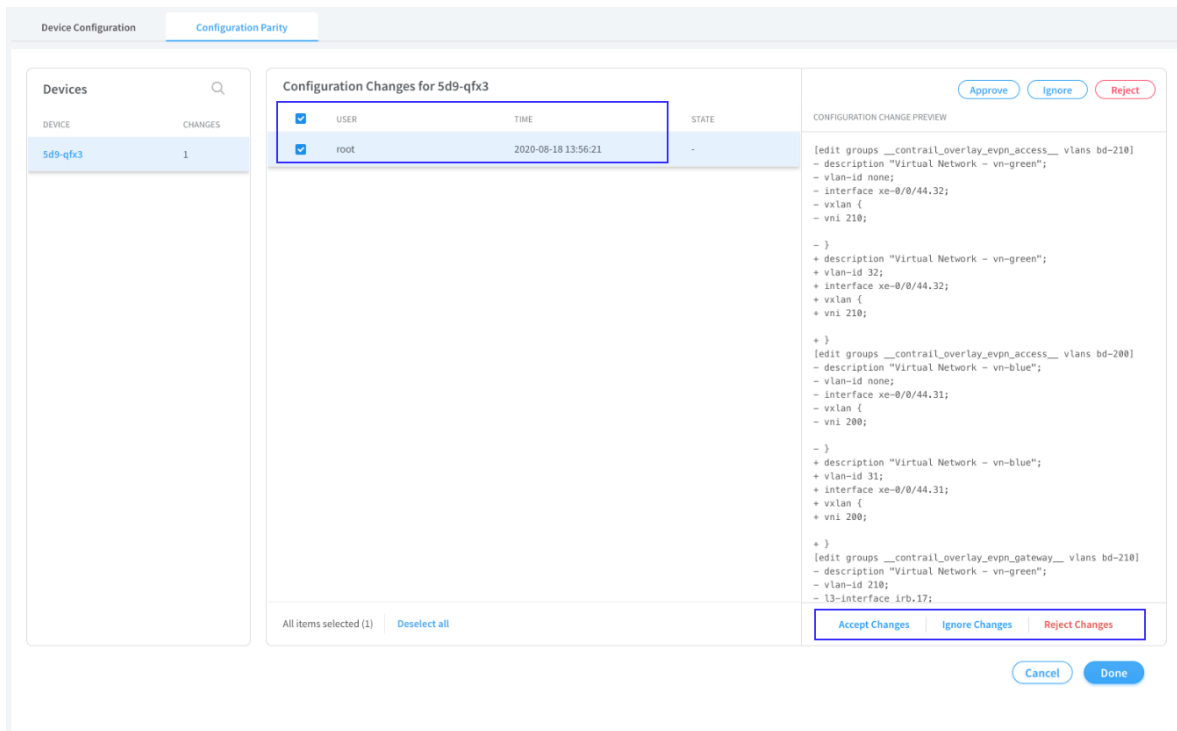
5. Click the **Configuration Parity** tab to view the device with new CLI configuration changes.

Alternatively, click **Review Changes** from the Device Configuration tab as shown in [Figure 116 on page 310](#), and you are directed to the **Configuration Parity** tab.

6. Select the check box next to the device name to view details of the user who made the change and the time the change was made.

This information is displayed in the Configuration Changes for *device* table as show in [Figure 118 on page 313](#).

Figure 118: Configuration Parity Tab



7. Select the check box next to the name of the user to view the Configuration Change Preview section. The Configuration Change Preview section displays the complete CLI configuration change. You can accept, ignore, or reject this configuration change from this section.

- Click **Accept Changes** to accept the manual CLI configuration change.

This new configuration is also saved in the Cassandra database when you accept changes.

In the event when the device is replaced (see "[Return Material Authorization](#)" on page 179) with a new device, you can retain the new configuration change or push the new change on to the new device from the Cassandra database.

- Click **Ignore Changes** to ignore the new CLI configuration change.

The new CLI configuration is still detected by Contrail Command and is available in the Cassandra database.

In the event when the device is replaced with a new device, the new configuration change that was ignored, will not be available on the new device.

- Click **Reject Changes** to remove the new CLI configuration from the device.

The new CLI configuration changes are not saved in the Cassandra database.

8. Click **Done** to confirm changes.

The Fabric Devices page is displayed. The device is now in the **Active** status.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2008	Starting in Contrail Networking Release 2008, you can also ignore manual CLI configuration changes.
2003	Contrail Networking Release 2003 supports the detection of manual CLI configuration changes.

RELATED DOCUMENTATION

| [Viewing Configuration of Devices Deployed in Contrail Fabric](#) | 305

Certificate Lifecycle Management Using Red Hat Identity Management

IN THIS SECTION

- [Fully Qualified Domain Names](#) | 315
- [Performing Lifecycle Management of Certificates using Identity Management](#) | 315

Contrail Networking Release 5.1 supports using Transport Layer Security (TLS) with RHOSP to perform lifecycle management, including renewal, expiration, and revocation, of certificates using Red Hat Identity Management (IdM). Because IdM uses fully qualified domain names (FQDNs) to manage endpoints instead of IP addresses, Contrail Networking services are also enhanced to use FQDNs.

Prior to Contrail Networking Release 5.1, lifecycle management of certificates was done manually.

Fully Qualified Domain Names

Contrail Networking Release 5.1 is integrated with IdM to perform lifecycle management of certificates. Contrail Networking services are also enhanced to use FQDNs in the following scenarios:

- Establishing connections between Contrail Networking components
- Input parameters for Contrail Docker container instead of IP addresses
- Contrail TripleO Heat Templates pass FQDNs instead of IP addresses for configuration of Contrail Networking containers using only TLS. You can configure TripleO Heat Templates to pass FQDNs without TLS by setting the `contrail_nodes_param_suffix`: 'node_names' option.
- Certificates are issued for every Contrail Networking node and stored in the `/etc/contrail/ssl` folder which is mounted on all Docker containers

Performing Lifecycle Management of Certificates using Identity Management

Perform the following steps to install the IdM server and manage certificates.

1. Deploy and configure IdM server.
For information on installing an IdM server, see [Installing an IdM Server: Introduction](#).
2. Before deploying the undercloud, set up the **novajoin** plugin on the undercloud node.

```
$ sudo yum install python-novajoin
$ sudo /usr/libexec/novajoin-ipa-setup \
    --principal admin \
    --password <IdM admin password> \
    --server <IdM server hostname> \
    --realm <overcloud cloud domain (in upper case)> \
    --domain <overcloud cloud domain> \
    --hostname <undercloud hostname> \
    --precreate
```

3. Prepare the undercloud configuration.

```
[DEFAULT]
enable_novajoin = true
ipa_otp = <otp> # is returned at previous step
```

```

undercloud_hostname = <undercloud FQDN>
undercloud_nameservers = <IdM IP>
overcloud_domain_name = <domain>
...

```

4. Check if firewalld is enabled on the IPA (Identity, Policy, Audit) server and the required ports are allowed.

```
rpm -qa | grep firewalld
```

If firewalld is not installed, the undercloud installation will fail. To install firewalld, use the following command:

```

yum install firewalld
firewall-cmd --permanent --add-port={80/tcp,443/tcp,389/tcp,636/tcp,88/tcp,88/udp,464/
tcp,464/udp,53/tcp,53/udp,123/udp}

firewall-cmd --permanent --add-service={freeipa-ldap,freeipa-ldaps,dns}

```

5. Deploy the undercloud.

```

$ openstack undercloud install
$ source stack rc

```

6. (Optional) Check the following services:

```

(undercloud) [stack@queensa ~]$ systemctl |grep nova
  novajoin-notify.service                                loaded active
running   OpenStack Nova IPA Notification Service
  novajoin-server.service                               loaded active
running   OpenStack Nova IPA Join Service
  openstack-nova-api.service                           loaded active
running   OpenStack Nova API Server
  openstack-nova-compute.service                       loaded active
running   OpenStack Nova Compute Server
  openstack-nova-conductor.service                    loaded active
running   OpenStack Nova Conductor Server
  openstack-nova-scheduler.service                    loaded active
running   OpenStack Nova Scheduler Server

```

7. Configure overcloud DNS and overcloud domain names.

```
$ openstack subnet set ctlplane-subnet --dns-nameserver <idm_server_address>
```

8. Add overcloud domain names to the **contrail-net.yaml** environment file.

```
DnsServers: ["<idm_server_address>"]
CloudDomain: lab.local
CloudName: overcloud.lab.local
CloudNameInternal: overcloud.internalapi.lab.local
CloudNameStorage: overcloud.storage.lab.local
CloudNameStorageManagement: overcloud.storagemgmt.lab.local
CloudNameCtlplane: overcloud.ctlplane.lab.local
```

9. Deploy overcloud with the following environment files.

```
$ openstack overcloud deploy --templates ~/tripleo-heat-templates \
  -e ~/overcloud_images.yaml \
  -e ~/tripleo-heat-templates/environments/network-isolation.yaml \
  -e ~/tripleo-heat-templates/environments/contrail/contrail-plugins.yaml \
  -e ~/tripleo-heat-templates/environments/contrail/contrail-services.yaml \
  -e ~/tripleo-heat-templates/environments/contrail/contrail-net.yaml \
  -e ~/tripleo-heat-templates/environments/contrail/contrail-tls.yaml \
  -e ~/tripleo-heat-templates/environments/ssl/enable-internal-tls.yaml \
  -e ~/tripleo-heat-templates/environments/ssl/tls-everywhere-endpoints-dns.yaml \
  -e ~/tripleo-heat-templates/environments/services/haproxy-internal-tls-certmonger.yaml \
  -e ~/tripleo-heat-templates/environments/services/haproxy-public-tls-certmonger.yaml \
  --roles-file ~/tripleo-heat-templates/roles_data_contrail_aio.yaml
```

The **contrail-net.yaml**, **enable-internal-tls.yaml**, **tls-everywhere-endpoints-dns.yaml**, **haproxy-internal-tls-certmonger.yaml**, and **haproxy-public-tls-certmonger.yaml** files enable TLS.

10. Check that the host is added to the IPA server.

```
# login to IPA
(undercloud) [stack@undercloud ~]$ kinit admin
(undercloud) [stack@undercloud ~]$ ipa host-find undercloud.my3domain
----- 1 host matched -----
Host name: undercloud.my3domain Description:
Undercloud host Principal name: host/undercloud.my3domain@MY3DOMAIN
```



```
Principal alias: host/undercloud.my3domain@MY3DOMAIN
SSH public key fingerprint: SHA256:GAMClAFagNN709Kb9AcFWfUG30Y06pcR0EdJBWXWIak (ssh-rsa),
SHA256:KqTDFkQEoKki7FMzuhBwc0+Y/09t4rHXQcqPKglJPMI (ecdsa-sha2-nistp256),
SHA256:QSIBCIiRW03eR6+PPyvDWiWEHXC1MewREAt8hMTU0gU (ssh-ed25519)
```

11. View the list of monitored certificates on an overcloud node.

```
[heat-admin@overcloud-novacompute-1 ~]$ sudo getcert list
Number of certificates and requests being tracked: 4.
Request ID 'contrail': status: MONITORING
stuck: no key pair storage: type=FILE,location='/etc/contrail/ssl/private/server-privkey.pem'
certificate: type=FILE,location='/etc/contrail/ssl/certs/server.pem'
CA: IPA
issuer: CN=Certificate Authority,O=MY3DOMAIN
subject: CN=overcloud-novacompute-1.my3domain,O=MY3DOMAIN
expires: 2021-04-20 14:18:21 UTC
dns: overcloud-novacompute-1.ctlplane.my3domain,overcloud-novacompute-1.internalapi.my3domain,overcloud-novacompute-1.tenant.my3domain,overcloud-novacompute-1.my3domain
principal name: contrail/overcloud-novacompute-1.my3domain@MY3DOMAIN
key usage: digitalSignature,nonRepudiation,keyEncipherment,dataEncipherment
eku: id-kp-serverAuth,id-kp-clientAuth
pre-save command:
post-save command: "sudo docker ps -q --filter=name="contrail*" | xargs -i sudo docker restart {}"
track: yes
auto-renew: yes
```

Collapsed Spine Architecture

IN THIS SECTION

- [Benefits | 320](#)

Starting from Release 2011, Contrail Networking supports collapsed spine, which is an architecture, in which there is no defined leaf layer. In collapsed spine architecture, the Layer 3 IP-based underlay and the EVPN-VXLAN overlay, which are usually run on the leaf, are built on the spine switches. These spine devices also act as the border gateway.

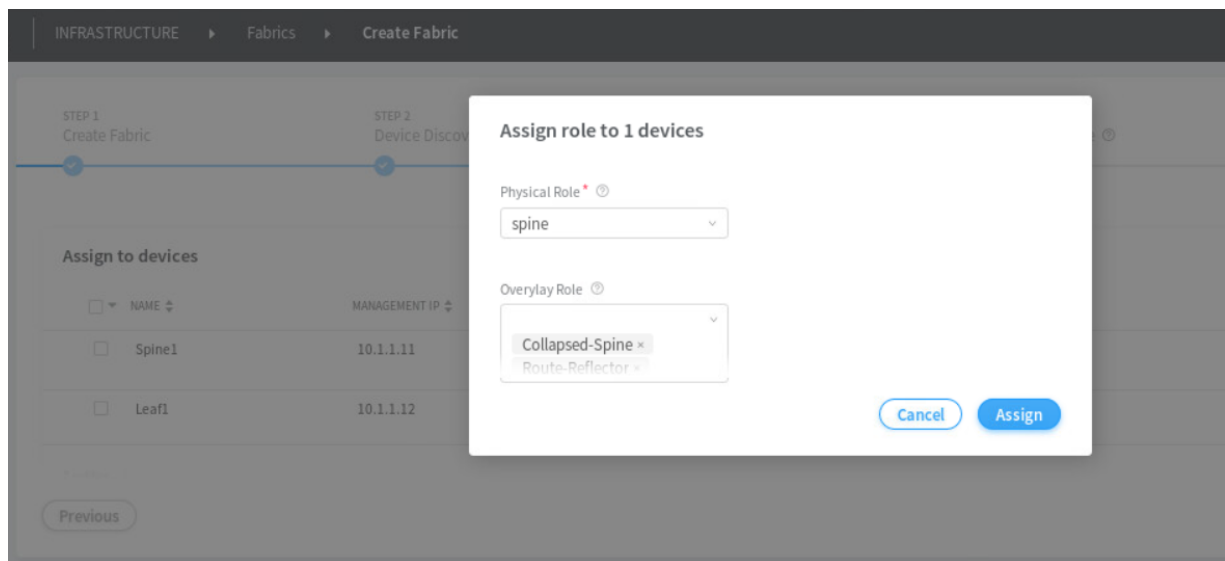
Collapsed spine architecture is supported on the following devices running Junos OS release 20.2R2 or later:

- QFX5120-32C—Can function as an L2 or L3 gateway
- QFX10K—Can function as L2, L3, or DC gateway

Supported QFX10K devices:

- QFX10002-72Q/36Q/60C
- QFX10002-36Q
- QFX10002-60C
- QFX10008
- QFX10016

Collapsed spine topology is supported in both greenfield and brownfield deployments. To configure a device as collapsed spine, you need to assign the physical role spine and overlay role Collapsed-Spine.



Benefits

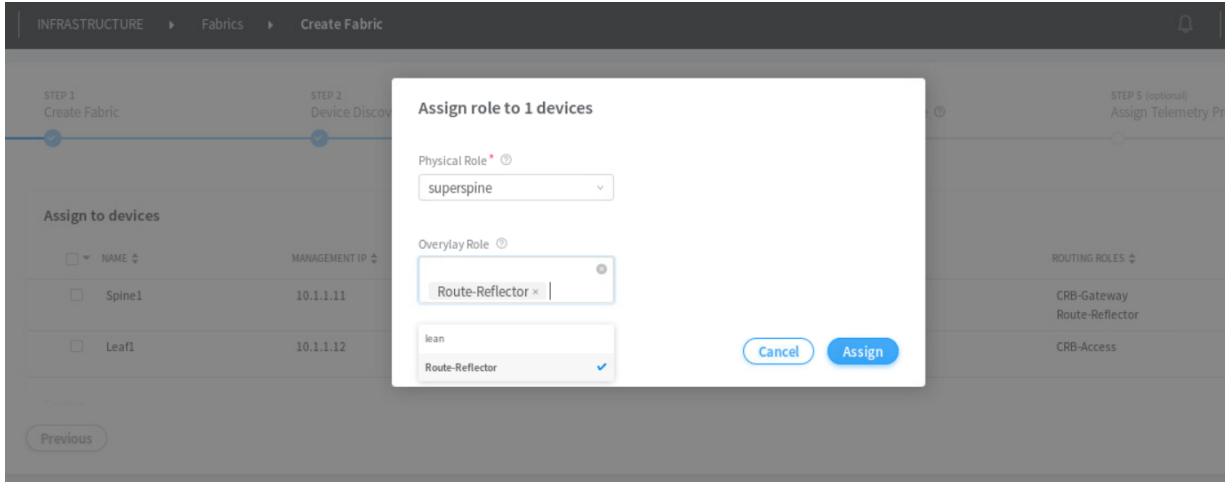
Collapsed spine architecture is useful for both enterprises and service providers that

- want to move to an IP fabric-based architecture with an EVPN-VXLAN overlay
- want to add two or more spine switches to ensure adequate bandwidth during maintenance window or to function as a backup in case of a spine failure.
- have small datacenters with mostly northbound and southbound traffic movement
- want to extend Layer 2 traffic across data centers.
- have multi-vendor legacy ToR switches that do not support EVPN-VXLAN.

Support for Superspine Role

Starting with Release 2011, Contrail Networking supports the superspine role. Devices with the role superspine can function as either a lean device that provides only IP forwarding or as a route reflector that can establish iBGP sessions with the peers. Superspines can be connected to spines in a CRB topology or to border-leaves in ERB topology. If the superspine is a lean device, then the spines should function as route reflector. If the superspine is functioning as a route reflector then the spines and leafs can peer with the superspine through iBGP sessions. Superspine role is supported in both greenfield and brownfield deployments.

To configure a device as superspine, you need to assign the Physical Role as superspine and overlay role as lean or Route-Reflector. If you assign the role lean, only IP forwarding will be supported and no iBGP session will be established with the peers. If you assign the Route-Reflector role, then the spine and leaf devices will establish iBGP sessions.



NOTE: Hierarchical route reflector feature is not supported in Release 2011.

5

CHAPTER

High Availability in Contrail Networking

[Using HA Cluster to Manage Fabric | 323](#)

[Hitless Software Upgrade of Data Center Devices Overview | 325](#)

[Performing Hitless Software Upgrade on Data Center Devices | 326](#)

[Fast Routing Convergence with Contrail Networking | 337](#)

[Configuring Fast Convergence from Contrail Command | 342](#)

Using HA Cluster to Manage Fabric

IN THIS SECTION

- [Topology Information | 324](#)

Contrail Networking Release 1911 supports High Availability (HA) cluster to manage fabrics.

In earlier releases, the All-in-One (AIO) cluster that contains the key components to run Contrail Networking on a single server, was used. After the fabric device discovery process begins, the AIO server becomes the DHCP server.

With the introduction of this high availability scenario, the DHCP server (dnsmasq) runs only during the zero-touch-provisioning (greenfield onboarding) process. After the fabric onboarding process is complete, the config files that are generated by the device manager and applied to dnsmasq, are deleted. This ensures that dnsmasq and device manager are active only on one node. After the files are deleted, the dnsmasq does not serve any more clients on the ZTP network.

In earlier releases, lease file records are maintained in `/var/lib/dnsmasq/dnsmasq.leases`. Starting in Contrail Networking Release 1911, lease file records are maintained in an external storage called Cassandra database.

NOTE: You cannot use both AIO cluster and HA cluster at the same time to manage the same fabric.

Topology Information

Figure 119: HA Topology for ZTP Subnet

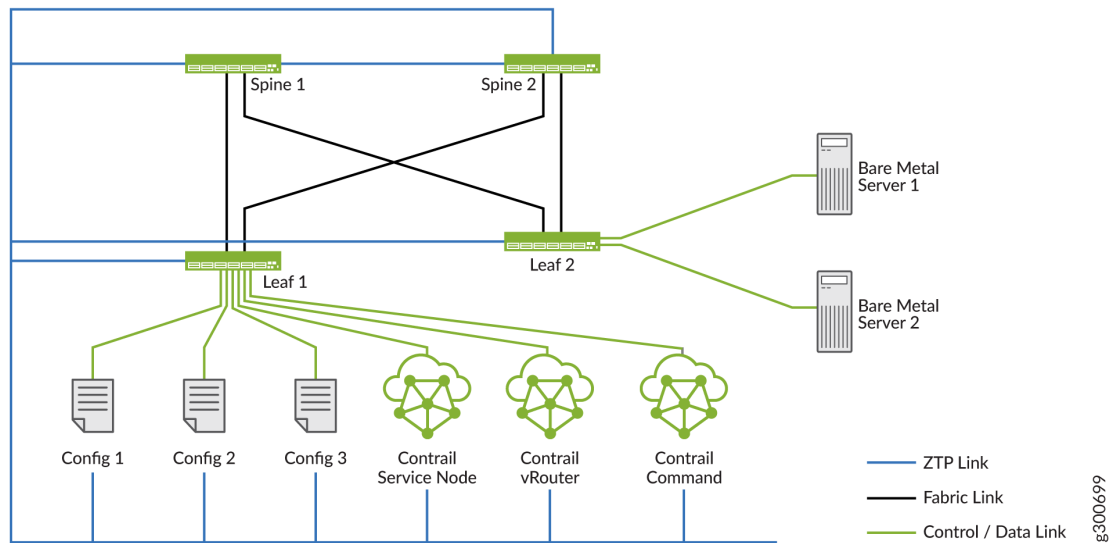


Figure 119 on page 324 shows an HA topology for ZTP. All spine and leaf switches, controllers, and nodes that are part of the deployment, are on the same ZTP subnet. All controllers are connected to a single ToR switch. The HA cluster is connected to a leaf switch.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1911	Contrail Networking Release 1911 supports High Availability (HA) cluster to manage fabrics.
1911	Starting in Contrail Networking Release 1911, lease file records are maintained in an external storage called Cassandra database.

RELATED DOCUMENTATION

[Fabric Overview](#) | 4

[Discover a Device](#) | 21

[Provision Fabric Devices Using End-to-End ZTP](#) | 32

Hitless Software Upgrade of Data Center Devices

Overview

IN THIS SECTION

- [Benefits of Hitless Software Upgrade | 326](#)

Contrail Networking Controller supports the automation of basic device management functions such as software image upgrade on the devices in the data center fabric. You can perform Contrail Networking Controller-assisted maintenance activities such as a hitless software image upgrade on the leafs and spines of the data center fabric devices managed by Contrail Networking, with zero packet loss.

Software image upgrade on a networking device in a data center is a time consuming task and might include rebooting the device. During upgrade, if user traffic is being routed through the device then the packets are lost which adversely affects the data center fabric performance.

During hitless upgrade, the devices are placed in a new mode called maintenance mode for the duration of the maintenance activity. The following sequence of steps are performed during hitless upgrade.

- **Initial Verification**

- Verifying that traffic can be routed from the selected device to another equally capable device. If no such device is present, then hitless upgrade cannot be performed because there will be traffic loss.
- Verifying that the selected upgrade image is compatible with the devices.
- Performing health checks on devices. Health checks are pre-configured parameters against which the devices are checked. If the health checks for devices fail, then the upgrade process for that device is terminated by default. However, you can change the default setting to not terminate upgrade upon health check failure.

If all the checks in the initial verification are cleared, Contrail Networking Controller places the device in the maintenance mode and performs the software upgrade.

- **Maintenance Mode**

- Before the device or devices are placed in maintenance mode, Contrail Networking Controller captures a snapshot of the existing state of the device. This snapshot is used to verify the operational state of the device when the maintenance activity or software upgrade is completed.

- The traffic flowing through the device is rerouted through another equally capable device and the Contrail Networking Controller verifies that there is no traffic flowing through the device.
- The device is then taken offline and placed in the maintenance mode.
- The Contrail Networking Controller upgrades the software image to the required version on the device.
- **Final Verification**
 - The device is taken out of the maintenance mode and traffic is routed through it again.
 - Contrail Networking Controller captures a snapshot of the operational state of the device to verify against the snapshot taken previously.

NOTE: For hitless software upgrade to work as per design and for zero packet loss, all devices must be redundantly connected. If any device is not redundantly connected, then you will have connectivity and packet loss when the device reboots.

Benefits of Hitless Software Upgrade

- Maintenance activities can be performed on devices in a data center without a maintenance window.
- No user traffic is lost during image upgrade on devices in the data center.

RELATED DOCUMENTATION

[Performing Hitless Software Upgrade on Data Center Devices](#) | 326

Performing Hitless Software Upgrade on Data Center Devices

Perform the following steps to upgrade the software image on the devices in a data center fabric with no loss of user traffic.

To perform hitless software upgrade on data center devices.

1. Upload the software images to which you want to upgrade your devices.
 - a. Navigate to the **Infrastructure > Fabrics** page in Contrail Command. A list of fabrics is displayed in the **Fabrics** tab.
 - b. Click the **Upload** button in the **Images** tab. The **Upload Image** page appears.
 - c. Enter the required software image details and click **Upload**. [Table 53 on page 329](#) lists all the mandatory parameters that must be entered to upload a software image.

Figure 120: Upload Image

Upload Image

Device Image Tags Permissions

Name*

Pick a File* ?

Vendor Name* ?

Device Family* ?

Supported Platforms* ?

Os Version* ?

Image MD5 ?

Table 53: Upload Image Fields

Field	Description
Name	Enter a name for the software Image. This name cannot be changed once the image has been uploaded.
Pick a File	Select the actual image file to be uploaded.
Vendor name	Enter the image vendor name. For example, Juniper, Arista, and so on.
Device Family	Enter the device family. For example, junos, junos-qfx, and so on.
Supported Platforms	Enter all the device platforms that the image is compatible on. For more information on supported platforms, see "Contrail Supported Platforms" on page 185 .
Os Version	Enter the OS version of the image. For example, 18.1R1.10.

NOTE: In the Os Version field, you should specify an exact OS version of the image. For example, if the OS version is '18.1R1.10 ' then you should specify it as '18.1R1.10' only. If you specify it as '18.1R1', the image upgrade will be successful. However, the upgrade job (running in the backend) will not be able to bring the device out of the maintenance mode (because of version mismatch) with no user feedback.

- d. Upon successful image upload, the **Images** tab appears listing the newly uploaded software image. Apart from the image name, you can edit image details at any time.

The same list of device images is available for image upgrade in [3](#).

Figure 121: Device Images

NAME	VENDOR	OS VERSION	FAMILY	URI	SUPPORTED PLATFORMS
jinstall-14.2R7.5-domestic-si...	juniper	14.2R7.5	junos	/contrail_container/jinstall-1...	-
jinstall-14.2R8.4-domestic-si...	juniper	14.2R8.4	junos	/contrail_container/jinstall-1...	-
jinstall-18.1R1.10-domestic-s...	juniper	18.1R1.10	junos-qfx	/contrail_container/jinstall-1...	-
jinstall-18.1R1.9-1-domestic-...	juniper	18.1R1.9	junos-qfx	/contrail_container/jinstall-1...	-
jinstall-18.1R1.9-domestic-si...	juniper	18.1R1.8	junos-qfx	/contrail_container/jinstall-1...	-
qfx_5h_try2	juniper	17.3R3.10	junos-qfx	/v1/AUTH_f6174b38527b443...	juniper-qfx10002-36q

2. Click the **Fabrics** tab and select a data center fabric.
The list of devices connected in a spine and leaf topology and corresponding details of each device in the selected fabric is displayed. The roles assigned to the devices are also displayed.
3. Click **Action > Image Upgrade**. The **Select Device** page appears. The list of images available to be upgraded to is displayed.
4. Select the image and the compatible devices to be upgraded to that image in the **Assign Devices** tab.
You can select one or more devices in the fabric. You can also select multiple images.

Figure 122: Select Device > Assign Devices

INFRASTRUCTURE > Fabrics > fab01 > Upgrade

STEP 1 Select Device | STEP 2 Testing | STEP 3 Upgrade

Assign Devices | Parameters

Images			
NAME	VENDOR	DEVICE FAMILY	DEVICES
▶ jinstall-14.2R7...	juniper	junos-qfx	vqfx1, vqfx2 1 more
▶ jinstall-14.2R8...	juniper	junos-qfx	vqfx5, vqfx3 1 more
▶ jinstall-host-qf...	juniper	junos-qfx	

Devices for image			
<input type="checkbox"/>	DEVICE NAME	VENDOR	DEVICE FAMILY
No data to display			

Cancel Next

5. Select the health check parameters for each device in the **Parameters** tab.

The health check parameters confirm that the devices and the network as a whole are stable to perform hitless image upgrade. By default, if health check fails for a particular device, then image upgrade is terminated. You can deselect the **Abort on health check failure** check box to continue upgrade on a device even if the health check fails.

Figure 123: Select Device > Parameters

INFRASTRUCTURE > Fabrics > fab01 > Upgrade

STEP 1 Select Device STEP 2 Testing STEP 3 Upgrade

Assign Devices **Parameters**

Abort on health check failure

Devices to upgrade simultaneously
4

BGP

Flaps allowed for BGP neighbors: 4 Down peers allowed: 0 Check: Flap count Down peer count Peer state

Alarm

Check: System alarm Chassis alarm

Interface

Check: Error Drop Carrier transition

Routing Engine

Cancel Next

6. Click **Next**. The **Testing** page appears.

The **Test Run** tab checks that the devices selected for upgrade are not already running the selected software version. The **Test Run** tab also displays the result of the health check on the devices for the parameters selected previously in the **Parameters** tab. If health check fails for the selected parameters, then you can go back to the previous page by clicking **Previous** and either changing the value of the health check parameter or disabling the parameter altogether. You can perform this step multiple times until health check passes for the device or you are able to determine that upgrade on the devices is feasible. Alternatively, you can click **Previous** and deselect the **Abort on health check failure** check box in the **Parameters** tab to continue upgrade on a device even if health check fails.

Figure 124: Testing > Test Run

INFRASTRUCTURE > Fabrics > fab01 > Upgrade

STEP 1 Select Device

STEP 2 Testing

STEP 3 Upgrade

Test Run Strategy Plan

Testing Progress

```

Tue Jan 22 2019 11:59:08 GMT-0800 (Pacific Standard Time)
Completed health check successfully. Putting device vqfx3 into maintenance mode.

Tue Jan 22 2019 11:59:10 GMT-0800 (Pacific Standard Time)
Completed health check successfully. Putting device vqfx2 into maintenance mode.

Tue Jan 22 2019 11:59:11 GMT-0800 (Pacific Standard Time)
Job summary: Job execution completed successfully.Detailed job results:
Successfully completed job for 6 devices.
Completed playbook execution for job template "hitless_upgrade_strategy_template" with execution id "1548187106366_2bda31a5-70b4-429c-8f6d-143cb45d4bd0"

```

Previous Cancel Next

7. Click the **Strategy Plan** tab. The **Strategy Plan** tab displays the strategy used to upgrade the images on the selected devices. Image upgrades occurs in batches, where multiple devices are upgraded at one go. The default maximum size of a batch is four devices.

The leafs are upgraded first and in a separate batch from their corresponding spines. If multihoming is configured on a BMS, the corresponding leafs are upgraded in different batches. No more than half the total number of spines will be upgraded in a batch. The batches are formed so as to have backup devices in a separate batch to the devices being upgraded in order to make the upgrade hitless. You can view the summary of the strategy used to upgrade the devices at the top and you can scroll down to view complete details of the devices. The estimated time for image upgrade per batch is also displayed.

Figure 125: Testing > Strategy Plan

INFRASTRUCTURE > Fabrics > fab01 > Upgrade

STEP 1 Select Device STEP 2 Testing STEP 3 Upgrade

Test Run Strategy Plan

Summary

Total estimated duration is 1:30:00.
 Note that this time estimate may vary depending on network speeds and system capabilities.
 The following batches of devices will be upgraded in the order listed:

Batch 1:
 vqfx3 18.1R1.9 --> 18.1R1.10
 vqfx5 18.1R1.9 --> 18.1R1.10

Batch 2:
 vqfx6 18.1R1.9 --> 18.1R1.10
 vqfx2 18.1R1.9 --> 18.1R1.10

Batch 3:
 vqfx4 18.1R1.9 --> 18.1R1.10

The following devices will not be upgraded for the reasons listed:
 vqfx1 (Upgrade image version matches current image version)

NOTE:

Details

Detailed information for the devices to be upgraded is listed below:

- vqfx2
 uuid : 24c49243-e0f4-4d9a-876e-683d44b37383
 vendor : Juniper
 family : junos-qfx
 product : vqfx-10000
 serial number : None
 management ip : 17x.xxx.x.1
 username : root
 password : ** hidden **
 new image version: 18.1R1.10
 current image version: 18.1R1.9
 image family : junos-qfx

Previous Cancel Next

8. Click **Next**. A confirmation page requesting confirmation of the image upgrade process is displayed.
9. Click **Yes** to confirm that you want to continue with the image upgrade. The **Upgrade** page appears displaying the status of the image upgrade progress for each device. The cumulative list of devices is displayed and the upgrade process happens according the batches determined in the strategy plan. The overall progress of all the devices is also displayed.
 Alternatively, click **No** to go back to the previous page.

Figure 126: Testing > Strategy Plan Confirmation

 Upgrade process will start after confirmation

Please get acquainted with the strategy plan in the test run log.
Upon confirmation, the strategy plan will get executed.
Are you sure?

No

Yes

10. Click on each device to view the image upgrade progress for that device. Click the device again to toggle back to display the overall image upgrade progress of all devices. [Table 54 on page 335](#) displays the states displayed during the course of the upgrade.

Table 54: Image Upgrade Progress States

State	Description
Loading Validating	The devices are prepping for the upgrade by running health checks.
Health Check Failed	Health check on the device has failed. You can click Previous and go back the Parameters page to either change the health check parameter value or disable the parameter.
Activating Maintenance Mode	The device has passed health check and the device is being placed under maintenance mode.
Deactivating Maintenance Mode	Removing maintenance mode configuration from device and exiting maintenance mode.
Maintenance Mode Activated	Maintenance mode is active on the device.

Maintenance Mode Deactivated	Deactivating maintenance mode is complete and maintenance mode configuration is successfully removed from the device.
Maintenance Mode Failure	Internal error detected during maintenance mode activation or deactivation.
Hitless Image Upgrade Successful	Device image is successfully upgraded.
Hitless Image Upgrade Failed	Device image is not upgraded.
Skipped	Attempted to upgrade to the same image version or the device family does not support hitless upgrade.

Figure 127: Upgrade

The screenshot shows the 'Upgrade' process in the 'Infrastructure > Fabrics > fab01 > Upgrade' page. The process is in Step 3: Upgrade. The progress bar shows Step 1 (Select Device) and Step 2 (Testing) are completed, and Step 3 (Upgrade) is in progress.

Progress for upgrading devices

DEVICE	IMAGE	STATUS
vqfx1	jinstall-14...	Loading...
vqfx2	jinstall-14...	MAINTENANCE_MODE_DEACTIVATED
vqfx4	jinstall-14...	Loading...
vqfx5	jinstall-14...	MAINTENANCE_MODE_DEACTIVATED
vqfx3	jinstall-14...	DEACTIVATING_MAINTENANCE_MODE
vqfx6	jinstall-14...	MAINTENANCE_MODE_DEACTIVATED

Upgrade progress for vqfx2

```

set protocols bgp group CLOS export MAINTENANCE-MODE-underlay
set protocols bgp group CLOS export export-bgp

Tue Jan 22 2019 12:03:24 GMT-0800 (Pacific Standard Time)
Configuration pushed down to activate maintenance mode on the device
vqfx2.

Tue Jan 22 2019 12:03:43 GMT-0800 (Pacific Standard Time)
Deploying following config to device 'vqfx2' (it may take a while)
delete groups __contrail_overlay_bgp__ protocols bgp group
_contrail_asn-64512 export MAINTENANCE-MODE

delete protocols bgp group CLOS export MAINTENANCE-MODE-underlay

```

Buttons: Previous, Cancel, Finish

11. Click **Finish** when all the devices have been upgraded.

Alternatively, to cancel the upgrade process, click **Cancel**. The **Infrastructure > Fabrics** page is displayed.

NOTE: You can re-enter the upgrade workflow if you exit at any point in the process. Also, in case of any failure, the reason is available in the device logs.

RELATED DOCUMENTATION

[Hitless Software Upgrade of Data Center Devices Overview | 325](#)

[Terminating Ongoing Fabric Jobs | 113](#)

Fast Routing Convergence with Contrail Networking

IN THIS SECTION

- [What is Convergence | 337](#)
- [Fast Network Convergence in a Network Managed by Contrail Networking | 338](#)

What is Convergence

Convergence or routing convergence is a state in which a set of routers in a network share the same topological information. The routers in the network collect the topology information from one another through the routing protocol. The state of convergence is achieved when all routers send routing information to all routers in the network. In other words, in a converged network, all routers are aware of the network topology and the optimal route to send a packet. Any change — for example, the failure of a device — in the network affects convergence until information about the change is propagated to all routers and convergence is achieved again. The time taken by the routers in the network to reach convergence after a change in topology is termed convergence time.

Network convergence and fast failover in case of failures in the network are critical for high performance service provider networks that run sensitive applications. The speed of achieving convergence in a network depends on the following actions:

- **Detection**—A device detects a failure in the route. Corrective action, that is, identifying a new forwarding path, can be taken only after identifying the device that failed. Unlike a physical network

where device availability or failure is identified through events, in virtual network, device reachability is established through keepalive messages. To achieve fast network convergence, the detection time – the time taken to detect failure – is of high importance and must be kept within acceptable limits.

- **Local Repair**—As soon as a device failure is detected in the primary route, traffic is diverted to backup route. At this point, the failure or the change in topology is not propagated to all the devices in the network.
- **Global Repair**—Global repair or network convergence is said to have achieved when the change in topology is propagated to all devices in the network through routing protocols.

The availability of services depends on the time taken for failure detection and correction.

Fast Network Convergence in a Network Managed by Contrail Networking

Contrail Networking provides software defined networking solution that offers network virtualization at the compute node-level through overlay networking. In a software-defined network, failures might occur in the overlay or in the underlay. A failure in the overlay can be the failure of a virtual machine or a pod. The vRouter can detect, rectify, and propagate any such failure in the overlay to the gateways by using the health check mechanism. Of the several possible types of failure in the underlay, the most critical ones are the SDN gateway failure and compute node failure.

Fast convergence feature improves the convergence time in case of failures in a cluster managed by Contrail networking. In a typical Contrail Networking-managed network, the customer end points connect to the vRouter through MPLSoUDP, GRE, or VXLAN overlay tunnels from the gateway device. The vRouters connect to the MPLS gateway through the fabric endpoints.

The fabric underlay routing has evolved to the extent that most fabric underlay failures such as those in the leaf or spine, in the links connecting the leaf and spine or the gateway and spine, vRouter and leaf, and so on can be efficiently addressed within a very short time without impacting the flow of traffic. However, there are two types of failures that can lead to silent packet drop. They are failures in the vRouter and failures in the gateway device, both of which are referred to as tunnel end point failures.

The overlay tunnels are maintained by the control node, which exchanges routes between the tunnel end points. The control node uses BGP to exchange routes with the gateway devices and XMPP to exchange routes with the vRouter. As there are no keepalive messages running on the overlay tunnels, the control node depends on the BGP hold time expiration to diagnose failures in the northbound connectivity to the gateway, and on the XMPP timeout for failures for southbound connectivity to the vRouter. Any failure in gateway can lead to a silent packet drop for up to 90 seconds, which is the default value for BGP hold time and expiration. This is because the control node can detect the gateway

failure only after the BGP hold time expires. Similarly, a failure in the vRouter can lead to a silent packet drop for up to 15 seconds, which is the default value for XMPP timeout.

Figure 128 on page 339 shows what happens when the gateway router fails. BGP hold time expires after 90 seconds as the destination is unreachable and the same is propagated to the control node, which recognizes that the gateway router has failed. This leads to a silent packet drop for 90 seconds until the routing table in the control node is updated and convergence achieved.

Figure 128: Tunnel Endpoint failure: SDN Gateway


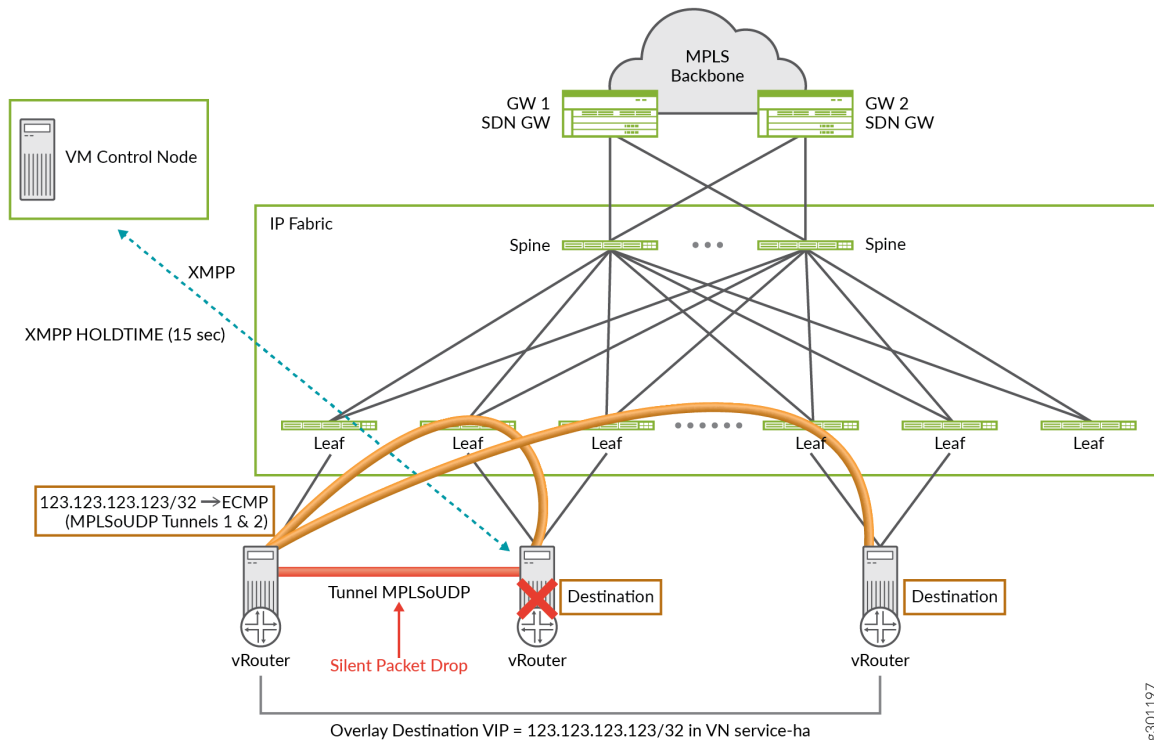


Figure 129 on page 340 shows a scenario where the vRouter fails. The control node comes to know about the vRouter failure only after 15 seconds when the XMPP hold time expires. Traffic is dropped during these 15 seconds until convergence is achieved.


Figure 129: Tunnel Endpoint Failure: vRouter



Starting with Release 2008, Contrail Networking supports fast convergence. The fast convergence feature reduces the convergence time in case of an overlay end point failure. The Contrail control plane responds to the changes in the underlay network and then takes action to achieve convergence quickly, reducing convergence time that would have taken in a typical scenario where control plane depends on BGP hold time expiration. Typically, the spines come to know of any tunnel end point failures through the BFD or the link down protocols. With the fast convergence feature, the spine propagates this information to the Contrail Controller and removes the tunnel end point from the control node through a routing table update. The control node recognizes this as a tunnel end point failure and initiates routing convergence. To respond to northbound failures (gateway failure), the control node performs a next-hop reachability check, and as soon as a failure is detected, the control plane initiates convergence. To achieve fast convergence in case of a southbound failure (vRouter failure), you can set the XMPP hold time to a value as low as one (1) second. Whenever the XMPP hold time expires, the control node recognizes it as a failure in the vRouter and initiates convergence. Though you can set a low value of one, the recommended timeout value is nine (9) seconds. A lower value is recommended only for smaller clusters.

Figure 130 on page 341 shows how Contrail Networking achieves fast convergence by using the destination reachability information that the spine gathers through BFD or link down protocols, and by using the XMPP timeout information sent to the control node to detect a failure in the vRouter.

Figure 130: Fast Convergence in a Contrail-managed Network



Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2008	Starting with Release 2008, Contrail Networking supports fast convergence.
2008	To achieve fast convergence in case of a southbound failure (vRouter failure), you can set the XMPP hold time to a value as low as one (1) second. Whenever the XMPP hold time expires, the control node recognizes it as a failure in the vRouter and initiates convergence.

RELATED DOCUMENTATION

| [Configuring Fast Convergence from Contrail Command](#) | 342

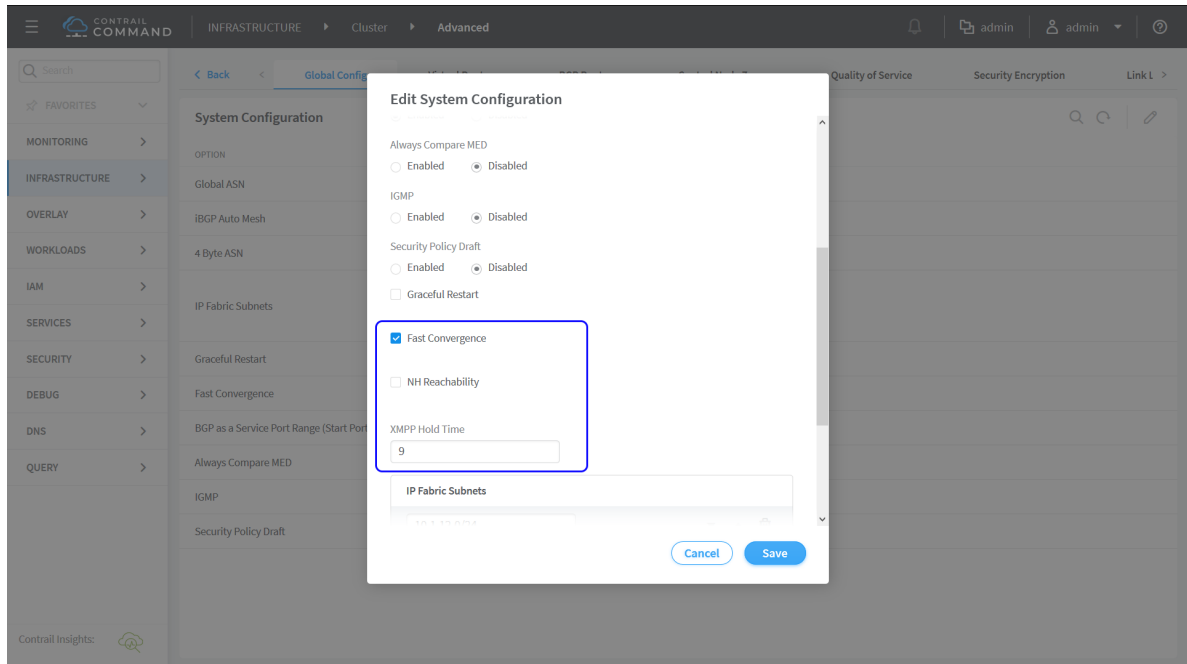
Configuring Fast Convergence from Contrail Command

This topic describes how to configure fast convergence.

To configure fast convergence:

1. Navigate to the **Infrastructure > Cluster >** and click **Advanced Options**.
2. From the **Global Config** tab, click **Edit System Configuration**.

Figure 131: Configure Fast Convergence Parameters



- **Fast Convergence**—Select this check box to enable convergence.
- **NH Reachability**—Select this check box to preform next-hop reachability check for gateway and vRouter. This field is displayed only if the **Fast Convergence** check box is selected.
- **XMPP Hold Time**—Enter a value in the range 1–90 in seconds. The XMPP hold time is used to detect vRouter failures. The timer expires when there is a vRouter failure, and the same is propagated to the control node, which in turn initiates convergence.

NOTE: Before you enable next-hop reachability check for fast convergence, make sure that BGP peering between the spines and control nodes is enabled for family inet/inet6 to exchange tunnel endpoint IP addresses. Also, enable redistribution of /32 vRouter host addresses from the underlay routing table to BGP.

3. Click **Save** to complete the configuration.

RELATED DOCUMENTATION

| [Fast Routing Convergence with Contrail Networking](#) | 337



CHAPTER

Integrating VMware with Contrail Networking Fabric

[Understanding VMware-Contrail Networking Fabric Integration | 345](#)

[Deploying Contrail vCenter Fabric Manager Plug-in | 348](#)

[Fabric Discovery and ESXi Discovery by Using Contrail Command | 351](#)

[Adding Distributed Port Groups | 358](#)

[Updating vCenter Credentials on Contrail Command | 360](#)

Understanding VMware-Contrail Networking Fabric Integration

IN THIS SECTION

- [Benefits of CVFM plug-in | 346](#)
- [CVFM Design Overview | 346](#)
- [Getting Started with CVFM Plug-in | 347](#)
- [Limitations of the CVFM Plug-in | 347](#)

Contrail Networking Release 1910 supports integrating VMware with Contrail Networking fabric. A dedicated Contrail vCenter Fabric Manager (CVFM) plug-in is deployed for this integration. This plug-in connects various ESXi hosts and helps manage VMware underlay networks. The CVFM plug-in is installed when you install the Contrail Command user interface (UI). After the plug-in is installed, the plug-in runs as a service in a container on the control node. You can enable this plug-in when you provision the Contrail Command UI. However, if you do not enable this plug-in during provisioning, you can enable the plug-in from the **Infrastructure>Cluster** page of the Contrail Command UI. For Contrail Networking Release 2008 and later, you can enable the plug-in from the **Infrastructure>External Systems** page of the Contrail Command UI.

In earlier releases, VMware provides a standard vCenter solution called vSphere ESX Agent Manager (EAM) to deploy, monitor, and manage Contrail VMs on ESXi hosts. Enterprise customers generally have a large number of VMware ESXi hypervisors and use EAM to manage tasks on virtualized platforms. Customers also use other VMware features such as creating Distributed Virtual Switches (DVS), creating Distributed Port Groups (DPG) on DVS, adding virtual machines on port groups, removing virtual machines from port groups, and moving virtual machines between port groups and hosts. However, EAM lacks the ability to automate the data center infrastructure.

With Contrail Networking Release 1910, the CVFM plug-in helps synchronize the configuration of VMware Distributed Port Groups (DPG) with the configuration on TOR (leaf) switches. After the CVFM plug-in is deployed, Contrail Networking will act as an automation tool that extends the management of ESXi hosts through VMware vCenter, to the data center infrastructure. For more information, see the **Design Overview** section of this topic.

Benefits of CVFM plug-in

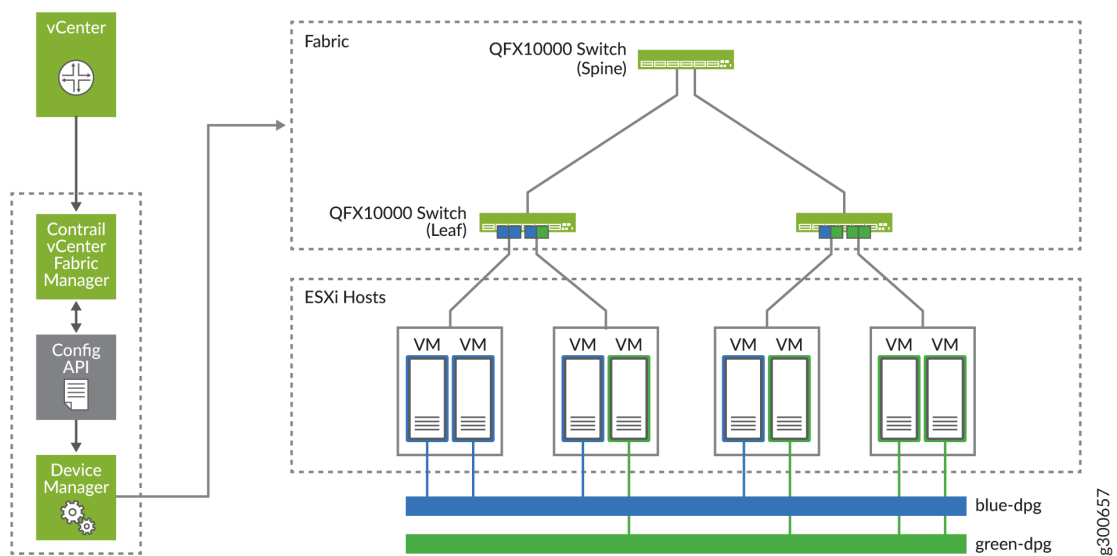
The following are the benefits of CVFM plug-in:

- Helps in integrating VMware with Contrail Networking fabric
- Synchronizes the configuration of VMware Distributed Port Groups (DPG) with the configuration on TOR (leaf) switches
- Enables Contrail Networking to act as an automation tool that extends the management of ESXi hosts through VMware vCenter, to the data center infrastructure
- Detects and communicates changes in the vCenter environment to the Contrail Device Manager

CVFM Design Overview

Figure 132 on page 346 depicts the CVFM plug-in installed on the Contrail Networking control node. The CVFM plug-in detects changes in the vCenter environment and pushes the new configurations to the Contrail Device Manager. The Contrail Device Manager then pushes these configurations to fabric devices such as QFX series switches.

Figure 132: Integrating VMware with Contrail Networking Fabric



The leaf and spine switches (QFX series) are connected to virtual machines in the ESXi host environment. VLANs are configured on the DPG of these QFX series switches. The CVFM plug-in

automatically adds and removes configurations of the VLANs. For more information on deploying the CVFM plug-in, see ["Deploying Contrail vCenter Fabric Manager Plug-in" on page 348](#).

Getting Started with CVFM Plug-in

The CVFM plug-in is installed when you install the Contrail Command UI.

1. You can then enable the CVFM plug-in while provisioning Contrail Command.

For more information, see the **Deploying CVFM Plug-in while Provisioning Contrail Command** section of the ["Deploying Contrail vCenter Fabric Manager Plug-in" on page 348](#) topic.

2. You can also enable the plug-in after provisioning Contrail Command.

For more information, see the **Deploying CVFM Plug-in after Provisioning Contrail Command** section of the ["Deploying Contrail vCenter Fabric Manager Plug-in" on page 348](#) topic.

3. After you have enabled the plug-in, you can update vCenter credentials or override configuration information from Contrail Command.

For more information, see ["Updating vCenter Credentials on Contrail Command" on page 360](#).

4. After you have enabled the plug-in, you must run the ESXi discovery process from Contrail Command. For more information, see ["Fabric Discovery and ESXi Discovery by Using Contrail Command" on page 351](#).

5. You can also add DPG. For more information, see ["Adding Distributed Port Groups" on page 358](#).

Limitations of the CVFM Plug-in

The following are the limitations of the CVFM plug-in.

- Supports DPG with standard VLAN. It does not support trunk (virtual machine to the DVS)/private VLAN.
- Supports network devices that are supported by Contrail Device Manager.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1910	Contrail Networking Release 1910 supports integrating VMware with Contrail Networking fabric. A dedicated Contrail vCenter Fabric Manager (CVFM) plug-in is deployed for this integration.

RELATED DOCUMENTATION

[Deploying Contrail vCenter Fabric Manager Plug-in | 348](#)

[Updating vCenter Credentials on Contrail Command | 360](#)

[Fabric Discovery and ESXi Discovery by Using Contrail Command | 351](#)

Deploying Contrail vCenter Fabric Manager Plug-in

IN THIS SECTION

- [Prerequisites | 349](#)
- [Deploying CVFM Plug-in while Provisioning Contrail Command | 349](#)
- [Deploying CVFM Plug-in after Provisioning Contrail Command | 349](#)
- [Troubleshooting Information | 350](#)

Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plug-in. With this release, the CVFM plug-in is installed when you install the Contrail Command user interface (UI). You can then enable this plug-in when you provision Contrail Command. However, if you have not enabled this plug-in during provisioning, you can enable the plug-in from the **Infrastructure>Cluster** page of the Contrail Command UI. For Contrail Networking Release 2008 and later, you can enable the plug-in from the **Infrastructure>External Systems** page of the Contrail Command UI.

For more information on CVFM plug-in, see "[Understanding VMware-Contrail Networking Fabric Integration](#)" on page 345.

These topics provide instructions on how to deploy the CVFM plug-in.

Prerequisites

Before you deploy the CVFM plug-in, ensure that you have:

- Installed vCenter version 6.5 or later.
- Installed ESX version 6.5 or later.
- A vCenter license with Distributed Virtual Switch (DVS) support.
- Login credentials for vCenter.
- Installed Contrail Command Release 1910 or later. For more information, see [Installing Contrail Command](#).

Deploying CVFM Plug-in while Provisioning Contrail Command

You can enable the CVFM plug-in while provisioning Contrail Command.

For steps to enable the CVFM plug-in while provisioning Contrail Command, refer to the *How to Install Contrail Command and Provision Your Contrail Cluster* topic.

Deploying CVFM Plug-in after Provisioning Contrail Command

Follow these steps to deploy the CVFM plug-in after provisioning Contrail Command.

1. For Contrail Networking Release 2008 and later,
 - a. Navigate to **Infrastructure>External Systems**.
The External Systems page is displayed.
 - b. Click **Add Orchestrator** and select **Manage vCenter** from the list.
The Deploy vCenter page is displayed.

For releases prior to Contrail Networking Release 2008,

- a. Navigate to **Infrastructure>Cluster** page.
The Overview tab is displayed.
- b. Click **Manage vCenter** in the Control Nodes widget.
The Deploy vCenter page is displayed.

2. Enter the following information:

Table 55: Enter vCenter Information

Field	Description
vCenter Server (For Contrail Networking Release 2008 and later.) vCenter IP Address (For releases prior to Contrail Networking Release 2008.)	Enter the vCenter IP address.
Data Center Name	Enter the name of the data center under vCenter that CVFM will work on.
Username	Enter the vCenter user name.
Password	Enter the vCenter password.

3. Select control node(s) from the Affected Nodes list.

The Affected Nodes list displays the control nodes that you can select to deploy the CVFM plug-in on.

4. Click **Deploy**.

The CVFM plug-in is deployed.

Troubleshooting Information

1. CVFM container continuously restarts

Check the following:

- a. Name of the CVFM container: `vcenter_fabric_manager_vcenter-fabric-manager_1`

NOTE: The `vcenter_fabric_manager_vcenter-fabric-manager_1` container runs on A-S-S Contrail Networking Controller.

- b. Standard CVFM log file: `/var/log/contrail/contrail-vcenter-fabric-manager.log`

- c. vCenter details in configuration file inside the CVFM container: `/etc/contrail/contrail-vcenter-fabric-manager/cvfm.conf`

Fix—reprovision CVFM with correct parameters.

2. LLDP disabled on ESXi

Issues:

- a. Configuration is not pushed to the network devices.
- b. No connection between ports and physical interface objects in config API

Fix—Enable LLDP on DVS by using VMware vSphere UI, delete ESXi servers in Contrail Command, and rerun the ESXi discovery job.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1910	Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plug-in.

RELATED DOCUMENTATION

[Understanding VMware-Contrail Networking Fabric Integration | 345](#)

[Updating vCenter Credentials on Contrail Command | 360](#)

[Fabric Discovery and ESXi Discovery by Using Contrail Command | 351](#)

How to Install Contrail Command and Provision Your Contrail Cluster

Fabric Discovery and ESXi Discovery by Using Contrail Command

IN THIS SECTION

● [Fabric Discovery | 352](#)

Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plug-in. After you deploy this plug-in, you must run the fabric discovery job and the ESXi discovery job from the Contrail Command UI.

These topics provide instructions to run the discovery job.

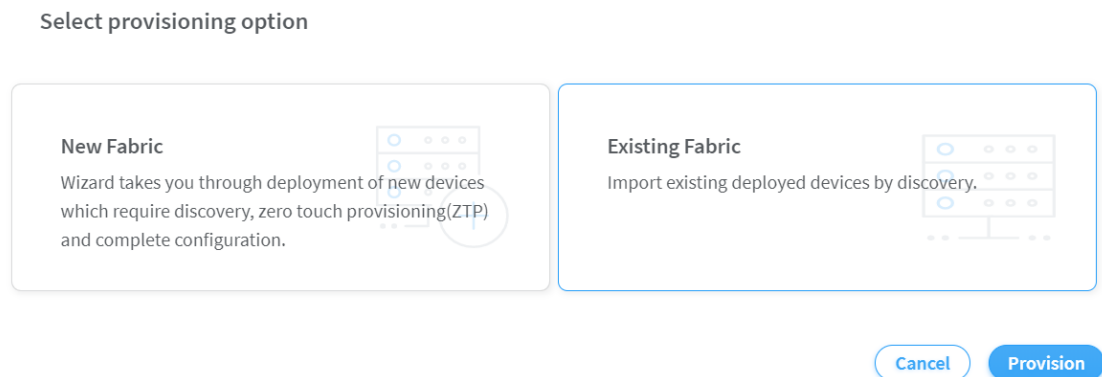
Fabric Discovery

Follow these steps to discover fabric devices.

1. Navigate to the **Infrastructure > Fabrics** page in Contrail Command.
2. Click **Create**.
You are prompted to select a provisioning option.
3. Click **Existing Fabric** to import existing (brownfield) devices by discovery.

NOTE: VMware-Contrail Networking Fabric integration supports greenfield device discovery and brownfield device discovery.

Figure 133: Select Provisioning Option



4. Click **Provision**.

The Create Fabric page is displayed.

5. Enter the fabric provisioning information as listed in [Table 56 on page 353](#).

Table 56: Provision Existing Fabric

Field	Action
Name	Enter a name for the fabric.
Overlay ASN (iBGP)	<p>Enter autonomous system (AS) number in the range of 1-65,535.</p> <p>If you enable 4 Byte ASN in Global Config, you can enter 4-byte AS number in the range of 1-4,294,967,295.</p>
Node profiles	<p>Add node profiles.</p> <p>You can add more than one node profile.</p> <p>All preloaded node profiles are added to the fabric by default. You can remove a node profile by clicking X on the node profile.</p>

Table 56: Provision Existing Fabric (Continued)

Field	Action
VLAN-ID Fabric Wide Significance	<p>Select the check box to enable enterprise style of configuration for the CRB-Access role on QFX devices. De-select the check box to enable service provider style of configuration for the CRB-Access role. The check box is selected by default since enterprise style is the default setting.</p> <p>You can modify the enterprise style setting to service provider style once configured. However, you cannot modify the service provider style to enterprise style.</p> <p>NOTE: Contrail Networking Release 1909 supports QFX10002-60C device running Junos OS Release 19.1R2 and later. QFX10002-60C device works only if enterprise style of configuration is enabled. To enable enterprise style of configuration, select the VLAN-ID Fabric Wide Significance check box when onboarding the QFX10002-60C device. For more information on enterprise style of configuration, see "Configuring EVPN VXLAN Fabric with Multitenant Networking Services" on page 289.</p> <p>For more information on supported hardware platforms and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185.</p>
Username	Enter a username for the device.
Password	Enter a password for the device.
Management subnets	<p>Enter the following information:</p> <p>CIDR—Enter CIDR network address.</p> <p>Gateway—Enter gateway address.</p> <p>NOTE: You enter the CIDR address range in the Management subnets field to search for devices. Any device that has a previously configured management IP on the subnet is discovered.</p>

Table 56: Provision Existing Fabric (Continued)

Field	Action
Loopback subnets	<p>Click Loopback subnets and enter loopback address in the CIDR field.</p> <p>NOTE: Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.</p>
Underlay ASNs (eBGP)	<p>Click Additional Configuration click +Add under Underlay ASNs (eBGP).</p> <p>Enter autonomous system (AS) number in the range of 1-65,535 in the CIDR field.</p> <p>If you enable 4 Byte ASN in Global Config, you can enter 4-byte AS number in the range of 1-4,294,967,295.</p> <ul style="list-style-type: none"> • Enter minimum value in ASN From field. • Enter maximum value in ASN To field.
Fabric subnets (CIDR)	<p>Click +Add under Fabric subnets (CIDR).</p> <p>Enter fabric CIDR address in the CIDR field.</p> <p>NOTE: Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.</p>
Advanced interface filters	<p>Click Advanced interface filters and select the Import configured interfaces check box.</p>

6. Click **Next**.

The Discovered devices page is displayed. The **Device discovery progress** bar on the Discovered devices page displays the progress of the device discovery job. The list of devices discovered is listed in the Discovered devices section.

7. Select the device you want to add to the fabric and then click **Add**.

The device is added to the fabric.

8. Click **Next** to assign roles.

The Assign to devices page is displayed.

9. Assign physical roles and routing bridging roles.

- To configure centrally-routed bridging (CRB):

For Spine Devices:

- Select **spine** from the Physical Role list.
- Select **CRB-Gateway** from the Routing Bridging Role list.

For Leaf Devices:

- Select **leaf** from the Physical Role list.
- Select **CRB-Access** from the Routing Bridging Role list.
- To configure edge-routed bridging (ERB):

For Spine Devices:

- Select **spine** from the Physical Role list.
- Select **CRB-MCAST-Gateway** from the Routing Bridging Role list.

For Leaf Devices:

- Select **leaf** from the Physical Role list.
- Select **ERB-UCAST-Gateway** from the Routing Bridging Role list.

NOTE: Contrail Networking Release 19XX supports CRB-Access, CRB-Gateway, DC-Gateway, ERB-UCAST-Gateway, and CRB-MCAST-Gateway roles overlay roles. For more information, see [Centrally-Routed Bridging Overlay Design and Implementation](#).

Assign a DC-Gateway Role to the spine device.

- Select **spine** from the Physical Role list.
- Select **DC-Gateway** from the Routing Bridging Role list.

For more information on supported hardware platforms, associated node profiles and roles, see ["Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 185](#).

10. Click **Assign** to confirm selection and then click **Autoconfigure** to initiate the auto-configuration job.
The Autoconfigure page is displayed.
11. After the autoconfigure process is completed, click **Proceed to Servers Discovery**.
You are redirected to the **Infrastructure>Servers>Servers Discovery** page.

ESXi Discovery

Follow these steps to discover ESXi servers by using the Contrail Command UI.

After you discover fabric devices, click **Proceed to Servers Discovery**. You are redirected to the **Infrastructure>Servers>Servers Discovery** page.

1. Click **ESXi** option button as shown in [Figure 134 on page 357](#).

Figure 134: Infrastructure > Servers > Servers Discovery

STEP 1
Configure

STEP 2
Servers Discovery

Choose server type to discover Physical/Virtual ESXi

vCenter Credentials

vCenter Server*

Data Center Name*

Username*

Password*

The vCenter Credentials section is displayed.

2. Enter the vCenter username in the **Username** field.
3. Enter the vCenter password in the **Password** field.

NOTE: The **vCenter IP Address** and **Data Center Name** fields are populated with vCenter credentials that were entered before deploying the CVFM plug-in.

4. Click **Next**.

The Servers Discovery Page is displayed. The **Device discovery progress** bar on the Discovered devices page displays the progress of the device discovery job. The list of devices discovered is listed in the Discovered Servers section.

5. After the servers discovery job is completed, click **Finish**.

The **Infrastructure>Servers** page is displayed. The list of ESXi servers are displayed in the Servers page.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1910	Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plug-in.
1909	Contrail Networking Release 1909 supports QFX10002-60C device running Junos OS Release 19.1R2 and later.

RELATED DOCUMENTATION

[Understanding VMware-Contrail Networking Fabric Integration | 345](#)

[Deploying Contrail vCenter Fabric Manager Plug-in | 348](#)

[Updating vCenter Credentials on Contrail Command | 360](#)

Adding Distributed Port Groups

You can add Distributed Port Groups (DPG) by using the VMware vSphere Web Client.

You add a DPG after you complete fabric discovery and ESXi discovery. For more information, see ["Fabric Discovery and ESXi Discovery by Using Contrail Command" on page 351](#).

Prerequisites

1. Ensure that Link Layer Discovery Protocol (LLDP) is enabled on leaf switches, spine switches, and ESXi hypervisors.

Follow these steps to enable LLDP by using the VMware vSphere Web Client.

- a. Navigate to **DSwitch**.

- b. Click **Actions**.
 - c. Select **Settings > Edit Settings**.
The Edit Settings page is displayed.
 - d. Click **Advanced**.
 - e. From the Discovery Protocol section, select **Link Layer Discovery Protocol** from the **Type** list.
 - f. Select an operational mode from the **Operation** list.
 - g. Click **OK** to confirm.
2. Ensure that the maximum transmission unit (MTU) configured on the leaf switches matches the MTU of the ESXi switches.

Follow these steps to add a DPG.

1. Select **DSwitch>Configure**.
2. Enter a name for the DPG in the **Name** field. Select location from the **Location** list.
3. Click **Next**.
The Configure Settings page is displayed.
4. Select **Static Binding** from the **Port Binding** list.
5. Select **Elastic** from the **Port allocation** list.
6. Enter number of ports in the **Number of ports** field.
7. Select **(default)** from the **Network resource pool** list.
8. Select **VLAN** as the VLAN type.
9. Enter VLAN ID in the **VLAN ID** field.
10. Click **Next**.

After you add a DPG, assign the DPG to the virtual machines. The configuration then gets pushed to the leaf switches that were discovered in the fabric discovery process.

RELATED DOCUMENTATION

| [Understanding VMware-Contrail Networking Fabric Integration](#) | 345

Updating vCenter Credentials on Contrail Command

Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plug-in. With this release, the CVFM plug-in is installed when you install the Contrail Command user interface (UI).

You can enable this plug-in:

- While provisioning Contrail Command
- After you provision Contrail Command

For more information, see "[Deploying Contrail vCenter Fabric Manager Plug-in](#)" on page 348.

You can also update vCenter credentials or override the configuration of the CVFM plug-in after you have enabled the plug-in. These steps provide instructions to update vCenter credentials from Contrail Command.

1. For Contrail Networking Release 2008 and later,
 - a. Navigate to **Infrastructure>External Systems**.
The External Systems page is displayed.
 - b. Click **Add Orchestrator** and select **Manage vCenter** from the list.
The Deploy vCenter page is displayed.

For releases prior to Contrail Networking Release 2008,

- a. Navigate to **Infrastructure>Cluster** page.
The Overview tab is displayed.
 - b. Click **Manage vCenter** in the Control Nodes widget.
The Deploy vCenter page is displayed.
2. You can update the following fields:

Table 57: Update vCenter Information

Field	Description
vCenter Server (For Contrail Networking Release 2008 and later.)	Edit the vCenter IP address.
vCenter IP Address (For releases prior to Contrail Networking Release 2008.)	

Table 57: Update vCenter Information (*Continued*)

Field	Description
Data Center Name	Edit the name of the data center under vCenter that CVFM will work on.
Username	Edit the vCenter user name.
Password	Update the vCenter password.
Select Control Nodes	Select control node(s) from the Affected Nodes list.

3. Click **Update**.

The credentials are updated.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
1910	Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plug-in.

7

CHAPTER

Integrating OpenStack with Contrail Networking Fabric

[Understanding OpenStack-Contrail Networking Fabric Integration](#) | 363

[Deploying ML2 Plug-in with Red Hat OpenStack](#) | 366

Understanding OpenStack-Contrail Networking Fabric Integration

IN THIS SECTION

- [Modular Layer 2 \(ML2\) Neutron Plug-in | 363](#)
- [Benefits of ML2 Plug-in | 364](#)
- [Design Overview | 364](#)

Contrail Networking Release 2011 supports integrating OpenStack with Contrail Networking Fabric. A Modular Layer 2 (ML2) Neutron plug-in is deployed for this integration. With this integration, you can manage underlay networks for OpenStack compute nodes.

Modular Layer 2 (ML2) Neutron Plug-in

Starting in Contrail Networking Release 2011, the ML2 Neutron plug-in is used to integrate OpenStack with Contrail Networking Fabric. Neutron is an OpenStack project that manages networking between interface devices. The ML2 plug-in enables OpenStack to use various layer 2 networking technologies that are found in complex real-world data centers.

In order to facilitate multi-vendor solutions, Neutron offers the following plug-ins:

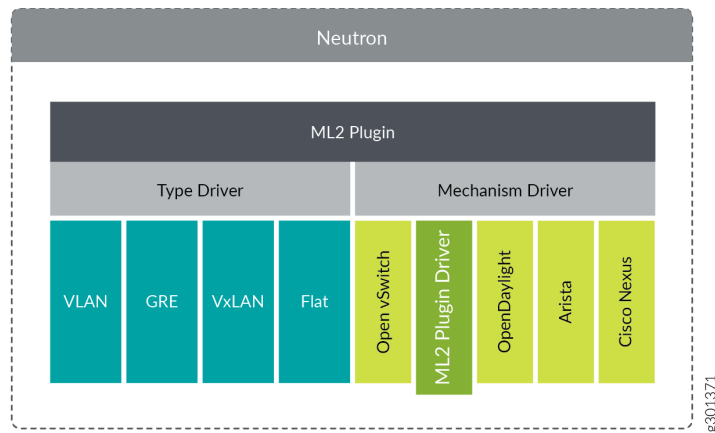
- Monolithic Plug-in
- Modular Layer 2 (ML2) Plug-in

The Monolithic plug-in is no longer supported and is removed from OpenStack. With this release, the only supported plug-in is ML2. For more information on the plug-in, see [ML2 Plug-in](#).

With Contrail Networking Release 2011, the ML2 plug-in is used with Contrail Command UI to facilitate multi-vendor solutions. ML2 plug-in works with Open vSwitch, Linux Bridge, and HyperV layer 2 agents. The ML2 plug-in simplifies adding support for layer 2 networking technologies. This plug-in requires less initial effort to deploy than what would be required to add a new monolithic plug-in.

The ML2 framework offers the following drivers:

Figure 135: ML2 Framework



Type Drivers Describes the type of underlying technology that is deployed.

Mechanism Drivers Facilitates the multi-vendor solutions built upon the technology specified by type driver.

Benefits of ML2 Plug-in

The following are the benefits of ML2 plug-in:

- Helps in integrating OpenStack with Contrail Networking Fabric.
- Used with Contrail Command to facilitate multi-vendor solutions.
- Simplifies adding support for layer 2 networking technologies.

Design Overview

This section describes the topology.

Figure 136: Topology

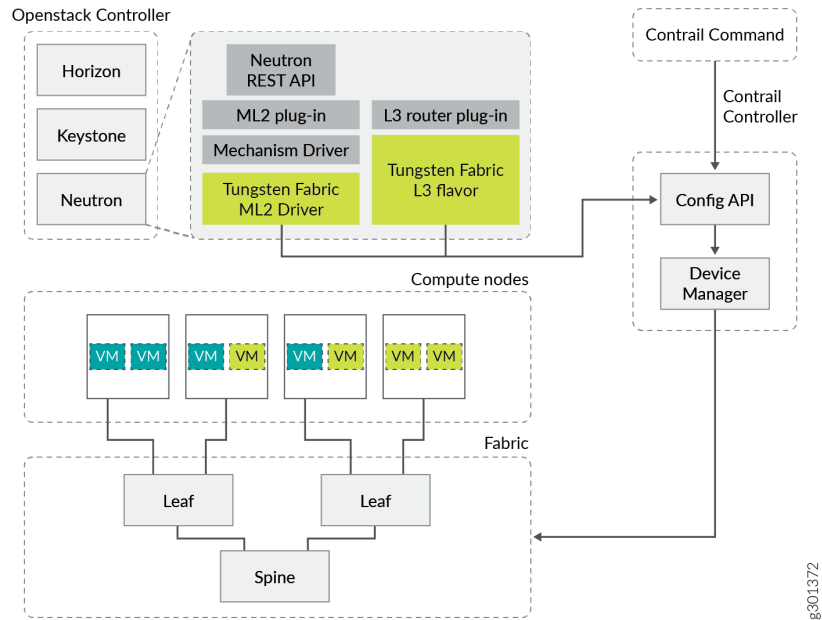


Figure 136 on page 365 depicts the OpenStack Controller, Contrail Networking Controller and the fabric. The Neutron Rest API, ML2 plug-in for layer 2 networking technologies, L3 router plug-in, and the mechanism drivers rests within the Neutron server. The ML2 plug-in mechanism driver (Tungsten Fabric ML2 Driver) and the Tungsten Fabric L3 driver connects the server to the Config API. The connection between the Config API and the fabric is established through the device manager.

ML2 Neutron Plug-in

This plug-in is also called the ML2 plug-in. The ML2 Neutron plug-in is used to integrate OpenStack with Contrail Networking Fabric. This plug-in enables OpenStack to use various layer 2 networking technologies that are found in complex real-world data centers. For more information, see "[Modular Layer 2 \(ML2\) Neutron Plug-in](#)" on page 363.

Tungsten Fabric Mechanism Driver

This driver is also called the Tungsten Fabric ML2 Driver or ML2 Driver. This driver is located in the OpenStack repository. The driver communicates with the Contrail Controller and manages Juniper devices. This driver is triggered by the ML2 plug-in in response to networking events such as creating a virtual network or modifying a port, that is handled by Neutron.

Tungsten Fabric Router Service Plug-in

This is also called the L3 Router Plug-in. The Tungsten Fabric Router Service Plug-in supports routing functionality on fabric. This plug-in also supports vendor-specific routing functionality. This plug-in pushes the ERB or CRB routing configuration on leaf or spine devices respectively.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2011	Starting in Contrail Networking Release 2011, the ML2 Neutron plug-in is used to integrate OpenStack with Contrail Networking Fabric.

RELATED DOCUMENTATION

[Deploying ML2 Plug-in with Red Hat OpenStack](#) | 366

Deploying ML2 Plug-in with Red Hat OpenStack

IN THIS SECTION

- [Deploy Contrail Command and CFM without Orchestrator](#) | 367
- [Configure Fabric by using Contrail Command](#) | 368
- [Deploy RHOSP13 with ML2 Plug-in](#) | 373
- [Configure Connectivity between RHOSP Internal API Network and Contrail Command Virtual Machines](#) | 378
- [Add Red Hat OpenStack Orchestrator](#) | 379
- [Create Swift Containers in OpenStack](#) | 380
- [\(Optional\) Deploy AppFormix and sFlows](#) | 381
- [Sample Network Files](#) | 384

Starting in Contrail Networking Release 2011, the ML2 Neutron plug-in is used to integrate OpenStack with Contrail Networking Fabric. Follow these steps to deploy ML2 plugin with Red Hat OpenStack 13 (RHOSP 13).

Deploy Contrail Command and CFM without Orchestrator

Follow these steps to deploy Contrail Command with intermediate OpenStack Keystone.

1. Deploy Contrail Command.
 - a. Prepare input data for Contrail Command deployer.

```

cat >command_servers.yml <<EOF
---
command_servers:
  server1:
    ip: 192.xx.xx.5
    connection: ssh
    ssh_user: cloud-user

    registry_insecure: false
    container_registry: svl-artifactory.juniper.net/contrail-nightly
    container_tag: master-latest
    config_dir: /etc/contrail

  contrail_config:
    database:
      type: postgres
      dialect: postgres
      password: contrail123
    keystone:
      assignment:
        data:
          users:
            admin:
              password: contrail123
    insecure: true
    client:
      password: contrail123
EOF

```

- b. Deploy Contrail Command deployer.

```

sudo docker run -ti --rm --net host --privileged --name contrail_command_deployer \
-v $(pwd)/command_servers.yml:/command_servers.yml \
-v $(pwd)/instances.yml:/instances.yml \

```

```
-v $(pwd)/.ssh/id_rsa:/root/.ssh/id_rsa \
svl-artifactory.juniper.net/contrail-nightly/contrail-command-deployer:master-latest
```

- From the Contrail Command UI, deploy Contrail Control nodes and select **None** when selecting orchestrator.

Configure Fabric by using Contrail Command

Follow these steps to configure fabric by using Contrail Command

Ensure that the following requirements are met.

- Switches are configured to provide connectivity for RHOSP networking.
- Names of servers are used as host names in RHOSP deployments.
- Virtual port groups that are created for deployment on ports and that are used by the fabric, conforms to ML2 naming convention.
 - For OVS ports, there is one virtual port group for every control node.

ML2 naming convention: `vpg#{base64(nodename)}`.

These virtual port groups are not used for SRIOV ports.

- For SRIOV ports, there is one virtual port group for every pair of compute node and physical network.

ML2 naming convention: `vpg#{base64(nodename)}#{base64(physnet)}`

All SRIOV ports need are tagged with the name of the physical network they are associated to. For example, `label=tenant1`.

- UUID of the virtual port group is set as a result of the `uuid.uuid3(uuid.NAMESPACE_DNS, str(name))` Python function.
- Created virtual port groups for all networks used in OOO provisioning.

The following script for creating virtual port groups is provided in the config-api container. This script is used for creating the infrastructure ports that are needed for RHOSP deployment.

```
python /opt/contrail/utils/provision_infra_nw.py -connections <connection.yaml> -fabric
<fabricname>
Sample Connection.yaml:
rhop-provisioning1:
```

```

cidr: 192.XX.XX.0/24
gateway: 192.XX.XX.254
vlan: 801
servers:
  5c7s5-node1.localdomain:
    5c7-qfx6:
      - xe-0/0/54_0
  5c7s5-node2.localdomain:
    5c7-qfx6:
      - xe-0/0/54_1
rhosp-int-api1:
cidr: 10.XX.XX.0/24
gateway: 10.XX.XX.254
vlan: 811
data: True
servers:
  5c7s5-node1.localdomain:
    5c7-qfx5:
      - xe-0/0/50_0
  5c7s5-node2.localdomain:
    5c7-qfx5:
      - xe-0/0/50_1

```

- Add servers to Contrail Command.

The server name should match the name that the node will inherit once OOO provisioning is complete. `hostname_map.yaml` is used here.

Follow these steps to import servers by using the Contrail Command UI.

1. Navigate to **Infrastructure>Servers** and click **Import**.

The Import Server pop-up is displayed.

2. To import a server, click **Browse** and navigate to the local directory and select the `.json` file.

Alternatively, you can drag and drop the `.json` file in the **Drag a file here, or browse** pane.

3. Click **Import** to import the server.

- Import Node (Server) Profiles.

Follow these steps to import node profiles by using the Contrail Command UI.

1. Navigate to **Infrastructure>Servers** and click the **Server Profiles** tab.

The Import Server Profile pop-up is displayed.

2. To import a server profile, click **Browse** and navigate to the local directory and select the .json file.

Alternatively, you can drag and drop the .json file in the **Drag a file here, or browse** pane.

3. Click **Import** to import the server profile.

- Associate node profiles (server profiles) and assign tags to SRIOV port only.

Follow these steps to associate node profiles to servers by using the Contrail Command UI.

1. Navigate to **Infrastructure>Servers**.

The Servers page is displayed.

2. Select the server you want to assign a server profile to by selecting the check box next to the name of the server.

3. Click **Assign** to server profile.

The Assign Server Profile pop-up is displayed.

4. Select the server profile from the Server Profile list and click **Assign**.

The profile is now assigned.

Sample Server Profile

```
{
  "nodes": [
    {
      "name": "5c7s5-node2.localdomain",
      "type": "baremetal",
      "ports": [{
        "name": "enp94s0f0",
        "mac_address": "90:e2:ba:4c:65:c9",
        "switch_name": "5c7-qfx6",
        "port_name": "xe-0/0/54:1",
        "switch_id": "10:0e:7e:bd:94:72"
      }],
    },
    {
      "name": "enp94s0f1",
      "mac_address": "90:e2:ba:4c:65:c9",
      "switch_name": "5c7-qfx5",
      "port_name": "xe-0/0/50:1",
      "switch_id": "10:0e:7e:bd:94:72"
    }
  ]
}
```

```

    },
    {
      "name": "enp94s0f2",
      "mac_address": "90:e2:ba:4c:65:c9",
      "switch_name": "5c7-qfx5",
      "port_name": "xe-0/0/2",
      "switch_id": "10:0e:7e:bd:94:72"
    },
    {
      "name": "enp94s0f3",
      "mac_address": "90:e2:ba:4c:65:c9",
      "switch_name": "5c7-qfx6",
      "port_name": "xe-0/0/2",
      "switch_id": "10:0e:7e:bd:94:72"
    }
  ]
}

```

Sample Node profile:

```

{
  "resources": [
    {
      "kind": "card",
      "data": {
        "name": "card1",
        "fq_name": ["card1"],
        "interface_map": {
          "port_info": [
            {
              "name": "enp94s0f2",
              "labels": ["physnet1"]
            },
            {
              "name": "enp94s0f3",
              "labels": ["physnet2"]
            }
          ]
        }
      }
    }
  ]
}

```

```

    }
  }
},
{
  "kind": "hardware",
  "data": {
    "name": "sriov-server1",
    "fq_name": ["sriov-server1"],
    "card_refs": [
      {
        "to": ["card1"]
      }
    ]
  }
},
{
  "kind": "tag",
  "data": {
    "tag_type_name": "label",
    "tag_value": "phynet1",
    "fq_name": ["label=physnet1"]
  }
},
{
  "kind": "tag",
  "data": {
    "tag_type_name": "label",
    "tag_value": "physnet2",
    "fq_name": ["label=physnet2"]
  }
},
{
  "kind": "node_profile",
  "data": {
    "hardware_refs": [
      {
        "to": ["sriov-server1"]
      }
    ],
    "parent_type": "global-system-config",
    "name": "sriov_1",
    "fq_name": ["default-global-system-config", "sriov_1"],
    "node_profile_vendor": "Sriov-server",

```

```

        "node_profile_type": "end-system"
    }
}
]
}

```

Deploy RHOSP13 with ML2 Plug-in

Follow these steps to deploy RHOSP13 with ML2 plug-in.

For detailed instructions on deployment, see [RHOSP13 DIRECTOR INSTALLATION AND USAGE](#).

1. Prepare Heat templates working folder.

```

# make copy of heat templates
cp -r /usr/share/openstack-tripleo-heat-templates/ tripleo-heat-templates
# get latest TF heat templates
git clone https://github.com/tungstenfabric/tf-tripleo-heat-templates -b stable/queens
# copy TF templates into working folder
cp -r contrail-tripleo-heat-templates/* tripleo-heat-templates/

```

2. If you use Nova Scheduler Hints for node placement,
 - a. Set appropriate capabilities properties for baremetal nodes.

Example Output

```

openstack baremetal node set --property capabilities='node:overcloud-
novacompute-0,boot_option:local' <id1>
openstack baremetal node set --property capabilities='node:overcloud-
controller-0,boot_option:local' <id2>
(see details https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/13/
html/advanced_overcloud_customization/sect-controlling_node_placement)

```

- b. Prepare scheduler hints environment (scheduler_hints.yaml) file.

```

parameter_defaults:
  ComputeSchedulerHints:
    'capabilities:node': 'overcloud-novacompute-%index%'

```



```

ComputeSriovSchedulerHints:
  'capabilities:node': 'overcloud-computesriov-%index%'
ControllerSchedulerHints:
  'capabilities:node': 'overcloud-controller-%index%'

```

3. If you use custom hostnames during server onboarding in Contrail Command, prepare `hostname_map.yaml`.

```

parameter_defaults:
  HostnameMap:
    overcloud-novacompute-0: b5s3
    overcloud-novacompute-1: b5s4
    overcloud-computesriov-0: b5s1

```

4. Modify compute role file to include `OS::TripleO::Services::NeutronDhcpAgent` service.
5. Prepare network parameters (`params.yaml`).

```

parameter_defaults:
  # admin user password
  AdminPassword: qwe123QWE
  # Customize all these values to match the local environment
  TenantNetCidr: 10.x.x.0/24
  InternalApiNetCidr: 10.x.x.0/24
  ExternalNetCidr: 10.x.x.0/24
  StorageNetCidr: 10.x.x.0/24
  StorageMgmtNetCidr: 10.x.x.0/24
  # CIDR subnet mask length for provisioning network
  ControlPlaneSubnetCidr: '24'
  # Allocation pools
  TenantAllocationPools: [{'start': '10.x.x.10', 'end': '10.x.x.200'}]
  InternalApiAllocationPools: [{'start': '10.x.x.10', 'end': '10.x.x.200'}]
  ExternalAllocationPools: [{'start': '10.x.x.10', 'end': '10.x.x.200'}]
  StorageAllocationPools: [{'start': '10.x.x.10', 'end': '10.x.x.200'}]
  StorageMgmtAllocationPools: [{'start': '10.x.x.10', 'end': '10.x.x.200'}]
  # Routes
  ControlPlaneDefaultRoute: 192.x.x.1
  InternalApiDefaultRoute: 10.x.x.1
  ExternalInterfaceDefaultRoute: 10.x.x.1
  # Vlans
  InternalApiNetworkVlanID: 710
  ExternalNetworkVlanID: 720

```

```

StorageNetworkVlanID: 730
StorageMgmtNetworkVlanID: 740
TenantNetworkVlanID: 3211
# Services
EC2MetadataIp: 192.x.x.1 # Generally the IP of the Undercloud
DnsServers: ["8.x.x.8"]
NtpServer: 3.europe.pool.ntp.org

```

6. Adjust options in the setup in the tripleo-heat-templates/environments/contrail/contrail-plugins-ml2.yaml file.

```

# Counts of nodes:
ControllerCount: 1
ComputeCount: 1
ComputeSriovCount: 1

# ml2/openvswitch_agent.ini: bridge_mappings
NeutronBridgeMappings:
  - datacentre:br-ex
  - tenant:br-vlans

# ml2/ml2_conf.ini: network_vlan_ranges
NeutronNetworkVLANRanges:
  - tenant:1:1000

# Sriov role specific options
ComputeSriovParameters:
  KernelArgs: "iommu=pt intel_iommu=on"
  TunedProfileName: "virtual-host"
# ml2/openvswitch_agent.ini: bridge_mappings
NeutronBridgeMappings:
  - datacentre:br-ex
  - tenant:br-vlans
  - tenant1:br-link1
  - tenant2:br-link2
# ml2/ml2_conf.ini: network_vlan_ranges
NeutronNetworkVLANRanges:
  - tenant:1:1000
  - tenant1:1:1000
  - tenant2:1:1000
# ml2/sriov_agent.ini: physical_device_mappings

```

```

NeutronPhysicalDevMappings:
  - tenant1:eth4
  - tenant2:eth5
NeutronSriovNumVFs:
  - eth4:8
  - eth5:8
# nova.conf: passthrough_whitelist
NovaPCIPassthrough:
  - devname: "eth4"
    physical_network: "tenant1"
  - devname: "eth5"
    physical_network: "tenant2"

# Adjust registry where Contrail containers are
# (contrail-node-init and contrail-openstack-neutron-ml2-init)
ContrailRegistry: '192.xxx.xx.10:8787'
ContrailImageTag: 'latest'
ContrailRegistryInsecure: true

# Address of Contrail Config API in format: ip1,ip2
# !!! Setup to correct IPs to point to existing Contrail cluster
# These IPs should be accessible from overcloud nodes, they are IPs where Config API
listen on.
ExternalContrailConfigIPs: <Config API IPs>

# Tags for ML2 plugin to differentiate ports in RHOSP networks
# (should be same as used in servers discovery in Contrail Command)
# !!! Adjust if other values are used in Contrail Command during servers discovery
ContrailManagementPortTags:
  - 'rhaps-provisioning'
  - 'rhaps-external'
  - 'rhaps-storage'
  - 'rhaps-internal'
  - 'rhaps-storage-mgmt'
ContrailDataPortTags:
  - 'rhaps-data'

```

This configuration will ensure that the ML2 Plugin ignores all TungstenFabric ports that contains any one of the following tags.

```
'rhaps-provisioning'
  'rhaps-external'
  'rhaps-storage'
  'rhaps-internal'
  'rhaps-storage-mgmt'
```

7. Prepare NIC files corresponding to the setup network layout.

Example for Compute tenant network to use VLANs without tunneling.

```
- type: ovs_bridge
  name: br-vlans
  members:
    - type: interface
      name: nic2
      primary: true
```

Example for SRIOV.

```
- type: ovs_bridge
  name: br-vlans
  members:
    - type: interface
      name: nic2
      primary: true
- type: ovs_bridge
  name: br-link0
  members:
    - type: interface
      name: nic3
      primary: true
- type: ovs_bridge
```

8. Upload Contrail containers to undercloud registry.

```
docker pull hub.juniper.net/contrail-node-init:2011.xx
```

```
docker pull hub.juniper.net/contrail-openstack-neutron-ml2-init:2011.xx
```

```
docker push 192.xxx.xx.1:8787/hub.juniper.net/contrail-node-init:2011.xx
```

```
docker push 192.xxx.xx.1:8787/hub.juniper.net/contrail-openstack-neutron-ml2-init:2011.xx
```

9. Deploy OpenStack.

```
source stackrc
openstack overcloud deploy --templates tripleo-heat-templates \
  --roles-file tripleo-heat-templates/roles_data_contrail_ml2.yaml \
  -e ~/overcloud_images.yaml \
  -e ~/hostname_map.yaml \
  -e ~/scheduler_hints.yaml \
  -e tripleo-heat-templates/environments/network-isolation.yaml \
  -e tripleo-heat-templates/environments/net-single-nic-with-vlans.yaml \
  -e tripleo-heat-templates/environments/contrail/contrail-plugins-ml2.yaml \
  -e params.yaml
```

10. After OpenStack is deployed, save internal virtual API.

```
Get RHOSP overcloud internal vip, .e.g
# ssh to one of openstack overcloud controller nodes
sudo hiera -c /etc/puppet/hiera.yaml internal_api_virtual_ip
```

Configure Connectivity between RHOSP Internal API Network and Contrail Command Virtual Machines

To configure connectivity between RHOSP internal API network and Contrail Command virtual machines, assign an IP from the network to an interface of the virtual machine.

```
[stack@command ~]$ cat /etc/sysconfig/network-scripts/ifcfg-eth2
# This file is autogenerated by os-net-config
DEVICE=eth2
ONBOOT=yes
HOTPLUG=no
NM_CONTROLLED=no
BOOTPROTO=none
MTU=1500
```

```
[stack@command ~]$ cat /etc/sysconfig/network-scripts/ifcfg-eth2.710
# This file is autogenerated by os-net-config
TYPE=vlan
VLAN=yes
DEVICE=eth2.710
ONBOOT=yes
HOTPLUG=no
NM_CONTROLLED=no
BOOTPROTO=none
MTU=1500
IPADDR=10.1.0.9
NETMASK=255.255.xxx.x
```

Add Red Hat OpenStack Orchestrator

You can add Red Hat OpenStack Orchestrator by using the Contrail Command user interface.

Follow these steps to add Red Hat OpenStack Orchestrator.

1. Navigate to **Infrastructure**>**External Systems**.

The External Systems page is displayed.

2. Click **Add Orchestrator** and select **RedHat OpenStack** from the list.

The Add OpenStack page is displayed.

3. Enter OpenStack Keystone endpoint IP address in the IP address field.

4. Enter OpenStack Keystone auth user name in the Username field.

5. Enter OpenStack Keystone auth password.

6. Click **Additional Configuration** and enter the information as given in [Table 58 on page 379](#).

Table 58: Additional Configuration

Field	Action
Domain Name	Specify the name of the OpenStack project domain. The default value is Default .
Protocol	Select the Keystone protocol you want to use for this configuration.

Table 58: Additional Configuration (Continued)

Field	Action
URL Version	The URL version for keystone authentication is /v3 by default.
Tenant	Specify the tenant for Keystone authentication. The default value is admin .
Region Name	Enter the name of the OpenStack-managed region within the data center. The default value is RegionOne .
Public port	Enter the port number to connect to Keystone authentication server. The default value is 5000 .

7. Click **Add** to add the orchestrator.

Create Swift Containers in OpenStack

Create a swift container and name it "contrail_container" with public read and list permissions. You can create a swift container from the Openstack UI.

Follow these steps to create a swift container by using the OpenStack UI.

1. Navigate to **Project>Object Store>Containers**.
The Containers page is displayed.
2. Click **+Container** to create a container.
The Create Container pop-up is displayed.
3. Enter a name for the container in the Container Name field.
4. Select **Public** from the Container Access options to enable anyone with the public URL to gain access to objects in the container.
5. Click **Submit** to create container.

(Optional) Deploy AppFormix and sFlows

Install AppFormix and xFlows by using `appformix-ansible-deployer`. Ensure that `instance.yml` has information on OOO and Keystone.

After you have installed Contrail and Red Hat OpenStack, follow these steps to install AppFormix HA and xFlows HA.

Follow these steps to install AppFormix HA.

Before you begin, ensure that Python3 is installed on xFlow nodes.

1. Navigate to `contrail_command` container.

```
/var/tmp/contrail_cluster/<uuid1>/instances.yml
```

2. Make a copy of the `/var/tmp/contrail_cluster/<uuid1>/instances.yml` file.
3. Edit `instances.yml` to include `appformix_controller` role and `appformix_bare_host` role in all nodes that are monitored. Include `appformix_openstack_controller` role in OpenStack node.
4. Log into the Contrail Command container.

```
cd /usr/share/contrail/appformix-ansible-deployer/venv
. venv/bin/activate
ansible-playbook -e config_file=<instance_file_path> --skip-tags=install_docker playbooks/
install_appformix_ansible.yml
```

The ansible files are now downloaded and the inventory file in `/opt/software/appformix/inventory` directory is generated.

5. Navigate to <https://ssd-git.juniper.net/appformix/AppFormix/wikis/appformix-installation-for-openstack-in-ha> and add the following to the

`/opt/software/appformix/inventory/hosts` file.

```
[appformix_controller]
<appformix_controller_ip1> keepalived_vrrp_interface=<if>
<appformix_controller_ip2> keepalived_vrrp_interface=<if>
<appformix_controller_ip3> keepalived_vrrp_interface=<if>
```


/opt/software/appformix/inventory/group_vars/all file.

```

appformix_vip: <appformix_vip_address>
openstack_auth_url: http://<keystone_auth_host>:5000//v3
openstack_project_domain_name: Default
openstack_user_domain_name: Default
openstack_username: admin
openstack_password: password
openstack_project_name: admin
openstack_tenant_name: admin
openstack_identity_api_version: 3
openstack_interface: internal
openstack_baremetal_api_version: 1.29
#docker exec -it contrail_command bash
#cd /usr/share/contrail/appformix-ansible-deployer/appformix
#source venv/bin/activate
(venv)# cd /opt/software/appformix/; ansible-playbook -i
inventory --skip-tags=install_docker
contrail-insights-ansible/appformix_openstack_ha.yml

```

Follow these steps to install xFlow HA.

1. Identify the Contrail Cluster ID from the /contrail-clusters API by using a debugger.
2. Add `appformix_flows` role to the node in the `instances.yml` file, where you want to install xFlows.

```

#docker exec -it contrail_command bash
#cd /usr/share/contrail/appformix-ansible-deployer/xflow
#source venv/bin/activate
#bash deploy_insights_flow.sh <instances.yml path> --cluster-id <contrail_cluster_id>

```

Sample `instances.yml` file snippets.

in-band installation of xFlows.

```

instances:
  host1:
    ip: 10.XX.XX.137
    provider: bms

```

```

roles:
  config:
    analytics:
    openstack:
    appformix_openstack_controller:
host2:
  ip: 10.XX.XX.136
  provider: bms
  roles:
    appformix_bare_host:
host3:
  ip: 10.XX.XX.135
  provider: bms
  roles:
    appformix_bare_host:
    appformix_flows:
.....
contrail_configuration:
  AUTH_MODE: keystone
  KEYSTONE_AUTH_HOST: 10.XX.XX.137
  KEYSTONE_AUTH_URL_VERSION: /v3
.....
xflow_configuration:
  telemetry_in_band_cidr: 1.XX.XX.1/24
  loadbalancer_management_vip: 10.XX.XX.166
  loadbalancer_collector_vip: 1.XX.XX.3
  telemetry_in_band_vlan_id: 51

```

xflow_configuration for **out-of-band** installation of xFlows.

```

xflow_configuration:
  loadbalancer_collector_vip: 10.XX.XX.166

```

3. After AppFormix and xFlows installation is completed, add endpoints.

Navigate to **Infrastructure>Cluster>Advanced Options>Endpoints** page in the Contrail Command UI and click **Create** to add endpoints.

Sample Network Files

- `tripleo-heat-templates/network/config/single-nic-vlans/role.role.j2.yaml`

```

heat_template_version: queens
description: >
  Software Config to drive os-net-config to configure VLANs for the {{role.name}} role.
parameters:
  ControlPlaneIp:
    default: ''
    description: IP address/subnet on the ctlplane network
    type: string
  {%- for network in networks %}
  {{network.name}}IpSubnet:
    default: ''
    description: IP address/subnet on the {{network.name_lower}} network
    type: string
  {%- endfor %}
  {%- for network in networks %}
  {{network.name}}NetworkVlanID:
    default: {{network.vlan}}
    description: Vlan ID for the {{network.name_lower}} network traffic.
    type: number
  {%- endfor %}
  ControlPlaneSubnetCidr: # Override this via parameter_defaults
    default: '24'
    description: The subnet CIDR of the control plane network.
    type: string
  ControlPlaneDefaultRoute: # Override this via parameter_defaults
    description: The default route of the control plane network.
    type: string
  {%- for network in networks %}
  {%- if network.ipv6|default(false) and network.gateway_ipv6|default(false) %}
  {{network.name}}InterfaceDefaultRoute:
    default: '{{network.gateway_ipv6}}'
    description: default route for the {{network.name_lower}} network
    type: string
  {%- elif network.gateway_ip|default(false) %}
  {{network.name}}InterfaceDefaultRoute:
    default: '{{network.gateway_ip}}'
    description: default route for the {{network.name_lower}} network

```

```

    type: string
{%- endif %}
{%- endfor %}
    DnsServers: # Override this via parameter_defaults
        default: []
        description: A list of DNS servers (2 max for some implementations) that will be added to
        resolv.conf.
        type: comma_delimited_list
    EC2MetadataIp: # Override this via parameter_defaults
        description: The IP address of the EC2 metadata server.
        type: string
    DnsSearchDomains: # Override this via parameter_defaults
        default: []
        description: A list of DNS search domains to be added (in order) to resolv.conf.
        type: comma_delimited_list
resources:
    OsNetConfigImpl:
        type: OS::Heat::SoftwareConfig
        properties:
            group: script
            config:
                str_replace:
                    template:
                        get_file: ../../scripts/run-os-net-config.sh
                    params:
                        $network_config:
                            network_config:
                                - type: ovs_bridge
{%- if role.name.startswith('CephStorage') or role.name.startswith('ObjectStorage') or
role.name.startswith('BlockStorage') %}
                            name: br-storage
{%- else %}
                            name: bridge_name
{%- endif %}
                        use_dhcp: false
                        dns_servers:
                            get_param: DnsServers
                        domain:
                            get_param: DnsSearchDomains
                        addresses:
                            - ip_netmask:
                                list_join:
                                    - /

```

```

        - - get_param: ControlPlaneIp
          - get_param: ControlPlaneSubnetCidr
    routes:
    - ip_netmask: 169.254.xxx.xxx/32
      next_hop:
        get_param: EC2MetadataIp
    - default: true
      next_hop:
        get_param: ControlPlaneDefaultRoute
    members:
    - type: interface
      name: nic1
      # force the MAC address of the bridge to this interface
      primary: true
{%- for network in networks if network.enabled|default(true) and network.name in
role.networks %}
{%- if network.name not in ["Tenant"] %}
    - type: vlan
      vlan_id:
        get_param: {{network.name}}NetworkVlanID
      addresses:
    - ip_netmask:
        get_param: {{network.name}}IpSubnet
{%- endif %}
{%- endfor %}
    - type: ovs_bridge
      name: br-vlans
      members:
    - type: interface
      name: nic2
      primary: true

outputs:
  OS::stack_id:
    description: The OsNetConfigImpl resource.
    value:
      get_resource: OsNetConfigImpl

```

- **tripleo-heat-templates/network/config/single-nic-vlans/compute-sriov.yaml**

```

heat_template_version: queens
description: >
  Software Config to drive os-net-config to configure VLANs for the {{role.name}} role.

```

```

parameters:
  ControlPlaneIp:
    default: ''
    description: IP address/subnet on the ctlplane network
    type: string
  {%- for network in networks %}
  {{network.name}}IpSubnet:
    default: ''
    description: IP address/subnet on the {{network.name_lower}} network
    type: string
  {%- endfor %}
  {%- for network in networks %}
  {{network.name}}NetworkVlanID:
    default: {{network.vlan}}
    description: Vlan ID for the {{network.name_lower}} network traffic.
    type: number
  {%- endfor %}
  ControlPlaneSubnetCidr: # Override this via parameter_defaults
    default: '24'
    description: The subnet CIDR of the control plane network.
    type: string
  ControlPlaneDefaultRoute: # Override this via parameter_defaults
    description: The default route of the control plane network.
    type: string
  {%- for network in networks %}
  {%- if network.ipv6|default(false) and network.gateway_ipv6|default(false) %}
  {{network.name}}InterfaceDefaultRoute:
    default: '{{network.gateway_ipv6}}'
    description: default route for the {{network.name_lower}} network
    type: string
  {%- elif network.gateway_ip|default(false) %}
  {{network.name}}InterfaceDefaultRoute:
    default: '{{network.gateway_ip}}'
    description: default route for the {{network.name_lower}} network
    type: string
  {%- endif %}
  {%- endfor %}
  DnsServers: # Override this via parameter_defaults
    default: []
    description: A list of DNS servers (2 max for some implementations) that will be added to
    resolv.conf.
    type: comma_delimited_list
  EC2MetadataIp: # Override this via parameter_defaults

```



```

        members:
          - type: interface
            name: nic1
            # force the MAC address of the bridge to this interface
            primary: true
    {%- for network in networks if network.enabled|default(true) and network.name in
role.networks %}
    {%- if network.name not in ["Tenant"] %}
        - type: vlan
          vlan_id:
            get_param: {{network.name}}NetworkVlanID
          addresses:
            - ip_netmask:
                get_param: {{network.name}}IpSubnet
    {%- endif %}
    {%- endfor %}
        - type: ovs_bridge
          name: br-vlans
          members:
            - type: interface
              name: nic2
              primary: true
        - type: ovs_bridge
          name: br-link0
          members:
            - type: interface
              name: nic3
              primary: true
        - type: ovs_bridge
          name: br-link1
          members:
            - type: interface
              name: nic4
              primary: true

outputs:
  OS::stack_id:
    description: The OsNetConfigImpl resource.
    value:
      get_resource: OsNetConfigImpl

```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2011	Starting in Contrail Networking Release 2011, the ML2 Neutron plug-in is used to integrate OpenStack with Contrail Networking Fabric.

RELATED DOCUMENTATION

| [Understanding OpenStack-Contrail Networking Fabric Integration](#) | 363

8

CHAPTER

Extending Contrail Networking to Bare Metal Servers

[Bare Metal Server Management | 392](#)

[How Bare Metal Server Management Works | 396](#)

[LAG and Multihoming Support | 398](#)

[Adding Bare Metal Server to Inventory | 400](#)

[Launching a Bare Metal Server | 402](#)

[Onboarding and Discovery of Bare Metal Servers | 404](#)

[Launching and Deleting a Greenfield Bare Metal Server | 405](#)

[Destination Network Address Translation for Bare Metal Servers | 406](#)

[Troubleshooting Bare Metal Servers | 411](#)

Bare Metal Server Management

IN THIS SECTION

- [Understanding Bare Metal Server Management | 392](#)
- [Features of the Bare Metal Server Management Framework | 394](#)

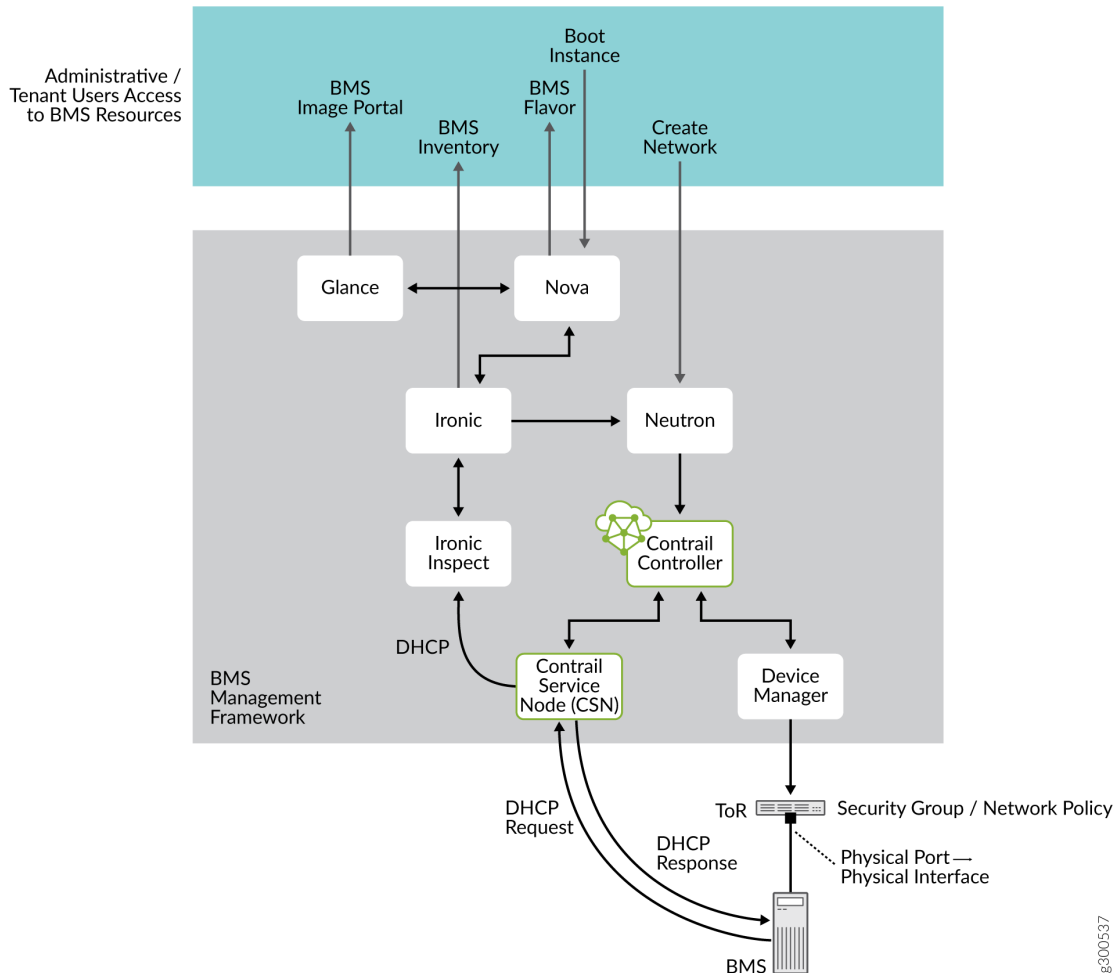
A bare metal server or a bare metal machine is a physical server that is dedicated to a specific customer, unlike a virtual machine. You can deploy bare metal machines in the same way as you deploy virtual machines by using Contrail UI.

Understanding Bare Metal Server Management

In Contrail Networking, you can manage the life cycle of bare metal servers (BMS) by using a backend framework, which acts as a bare metal server (BMS) manager. The BMS management framework in Contrail uses the functionality provided by the following OpenStack services: Ironic, Nova, and Glance. The BMS Management framework or the BMS framework manages the bare metal workload within a fabric. It includes BMS server life cycle management, onboarding of bare metal servers, bare metal image management, flavor management, inventory management, IP address management, security management, monitoring and reporting of life cycle management events, and discovery of bare metal servers.

An administrative user can configure the BMS framework and a tenant user can avail the services provided by the BMS framework. [Figure 137 on page 393](#) shows an architectural view of the BMS Management framework.

Figure 137: BMS Management- Detailed Architecture View



NOTE: In single-tenant environments, administrative and tenant workflows are performed by the same user.

To avail the functionalities of the BMS framework, you must first deploy a Contrail cluster with OpenStack. After this, the administrative user needs to specify the details of the server or node to be added to the BMS available in nodes database, from the Contrail Command UI. The BMS framework then creates a record of this new node and adds them to the available nodes database.

The administrative user creates images, nodes, and flavors, which the tenant users use to deploy bare metal servers in their network. A tenant user selects one of these flavors and images that suit their need to deploy a bare metal server. The BMS framework monitors the state of the deployed servers and provides this information to analytics DB by using Sandesh, which is an XML-based protocol for reporting analytics information. All the nodes onboarded or registered with BMS manager are in

Available state. After the tenant user has completed using the bare metal server and remove it, the server is then unprovisioned by the BMS framework and moved to the list of available nodes.

Alternatively, the tenant user can remove the BMS instance from the tenant's network. For example, if you want to rent a BMS from a service provider, the service provider deploys a BMS instance and gives you an IP address of the BMS instance, which you can use to access the BMS. Once you have completed using the BMS, you can delete the instance and the service provider reclaims the BMS. After reclaiming the BMS the service provider cleans it and rents it to the next client. The BMS framework in Contrail Networking manages all these tasks. If the service provider wants to remove the BMS instance from the service, they can delete it from the available servers and the next tenant will get a new BMS instance from a server.

The BMS framework can install tenant user-specific software images on BMS and attach them to the tenant user network in a multi-tenant cloud. It provides a single-click solution for the tenant users to manage the bare metal servers in their network.

Features of the Bare Metal Server Management Framework

The BMS management framework provides the following features:

- **BMS Image management**—Provides a list of available bootable images available to the tenant users to boot their server instances or BMS. The BMS framework uses Glance, which is an OpenStack service used for Image Management.
- **BMS Flavor management**—Provides a list of available flavors of the BMS available in the inventory. The flavors represent the capacity or class of the BMS, such as disk size, memory size, number of cores or the manufacturer of BMS. The BMS framework creates pools of BMS based on their capability, class, or both, and then makes them available to the tenant users. The BMS framework uses Nova, which is an OpenStack service used to provision computing instances or virtual servers. Nova can be used to create virtual machines and bare metal servers using Ironic. Flavors are used in OpenStack to define the compute, memory, and storage capacity of the Nova computing instances.
- **BMS Life Cycle Management**—Includes the following:
 - **Bringing powered off servers online in a secure manner**—As soon as a BMS is powered off, it is disconnected from the tenant user network and connected to a cleaning network for clean up of the server. A server is connected to a cleaning network for cleaning operations when it is not being used. If the server is deployed, it is connected to the provisioning network.
 - **Reclaiming the provisioned servers and instances after they are decommissioned by the tenant users**—After cleaning up, the BMS is added to the pool of available server ready to be deployed as a new BMS. The boot up process is performed on a secure network to prevent the possibility of snooping in a multi-tenant cloud. The cleaning process ensures that the BMS is ready to be deployed with the same or different image when needed.

The BMS framework uses Ironic, which is an OpenStack service used to launch bare metal machines. Ironic integrates with the bare metal driver in Nova to maintain BMS lifecycle management efficiently.

- BMS Inventory Management— Maintains an inventory of all the servers under the BMS framework. The inventory includes the deployed instances and servers as well as those available for deployment.
- BMS IPAM management— Ensures that the IP address management is consistent between the virtual and physical instances. IPAM is managed by the Contrail controller.
- BMS Network Security management— The boot cycle and/or cleaning of bare metal servers are extensive and lengthy processes, which makes provisioning and cleaning phases susceptible for snooping by hackers in multi-tenant cloud environments. Hence, the BMS framework uses private networks for the provisioning and cleaning phases of the servers. Once the servers are ready for deployment, the BMS framework deploys the servers in the tenant user network.
- Tenant Network management— Manages connectivity between the bare metal servers and tenant user networks or provisioning and cleaning networks depending on the deployment state of the server.
- BMS discovery and onboarding— The BMS framework supports both the discovery of new servers as well as onboarding of the brownfield servers.

NOTE: A deployed server must be unprovisioned and made available before it can be deleted from BMS node list.

RELATED DOCUMENTATION

[How Bare Metal Server Management Works | 396](#)

[LAG and Multihoming Support | 398](#)

[Adding Bare Metal Server to Inventory | 400](#)

[Launching a Bare Metal Server | 402](#)

[Onboarding and Discovery of Bare Metal Servers | 404](#)

[Launching and Deleting a Greenfield Bare Metal Server | 405](#)

[Destination Network Address Translation for Bare Metal Servers | 406](#)

[Troubleshooting Bare Metal Servers | 411](#)

Ironic Support with Juju

How Bare Metal Server Management Works

IN THIS SECTION

- Administrative Workflow | 396
- Tenant Workflow | 397

The BMS management framework is configured by the administrative user. The administrative user follows a specific workflow to configure critical data objects, which are then made available to the tenant users.

Administrative Workflow

The administrative user must perform the following workflow to configure the BMS framework:

- Create two private networks from the Contrail Command user interface (UI), which is visible only to the administrative user. One network is used for provisioning the servers during deployment phase and the other network is used for cleaning up bare metal servers when they are decommissioned. These private networks provide security to these servers from hackers when they are being provisioned or being cleaned up after removing from the tenant network. From the Contrail networking point of view, the private networks are normal virtual networks, except that they are accessible only to the administrative user.

To create virtual networks follow the procedure in "[Create Virtual Network](#)" on page 82.

NOTE: Though it is recommended that you create two networks for provisioning and cleaning, alternatively, you can use the same network for both provisioning and cleaning.

- Create the BMS images that are available to the tenants through a catalogue—You use the `diskimage-builder`, a special utility in the OpenStack Ironic service to create BMS images. For more information, see <https://docs.openstack.org/diskimage-builder/latest/>.
- Register the BMS images with Glance service—After the images are registered, these images become available to the tenant users for deployment. For more information, see <https://docs.openstack.org/ironic/latest/install/configure-glance-images.html>.

- Create bare metal flavors and register with Nova service based on the classes or bare metal servers to be offered or managed—You can create multiple bare metal flavors. For example, baremetal-huge, baremetal-large, baremetal-small, and so on. These flavors are then mapped to the inventory of the available bare metal servers at the time of deployment. The tenant users can view the flavors in the Contrail Command UI and use the flavors according to their requirement.
- Create Ironic nodes—A BMS server is represented as an Ironic node. The collection of the nodes form the BMS inventory.

To add a bare metal server to Inventory from the Contrail Command UI, the administrative user must follow the procedure in ["Adding Bare Metal Server to Inventory" on page 400](#).

- Create Ironic ports—These ports represent the NICs in the bare metal servers. This includes the MAC address and the physical connectivity information.
- Set up PXE boot interface—You set up Preboot Execution Environment (PXE) as part of BMS onboarding (or registering) of bare metal servers.

Tenant Workflow

After the BMS service is instantiated, the tenant users are offered a catalog of available services. They select the type of server they want to instantiate and the image they want to run. The tenant users need to follow the given workflow to avail the services provided by bare metal servers:

- Create tenant user network—BMS connects to this network when it is ready for use.
- Select the BMS flavor and BMS Image that you want to instantiate and issue a boot command. The tenant user selects a BMS that is available for deployment using the flavor. They use the flavors that are created by the administrative user. If no BMS meets the criteria specified by flavor, the launch command is rejected with the error message No Valid Host found.

NOTE: Booting a bare metal server is very much similar to instantiation of a virtual machine; the only difference is that the tenant user can select the appropriate flavor for BMS depending on the requirement.

- View availability zone information— An availability zone typically applies to virtual machines and can also be applied to BMS. You can view virtual machine availability zone information and BMS availability zone information in two different zones on the user interface.
- Launch a BMS—A bare metal server is launched in the same way as you launch a virtual machine.

To launch a new bare metal server from the Contrail Command UI, follow the procedure in ["Launching a Bare Metal Server"](#) on page 402.

RELATED DOCUMENTATION

[Bare Metal Server Management](#) | 392

[LAG and Multihoming Support](#) | 398

[Adding Bare Metal Server to Inventory](#) | 400

[Launching a Bare Metal Server](#) | 402

[Onboarding and Discovery of Bare Metal Servers](#) | 404

[Launching and Deleting a Greenfield Bare Metal Server](#) | 405

[Destination Network Address Translation for Bare Metal Servers](#) | 406

[Troubleshooting Bare Metal Servers](#) | 411

LAG and Multihoming Support

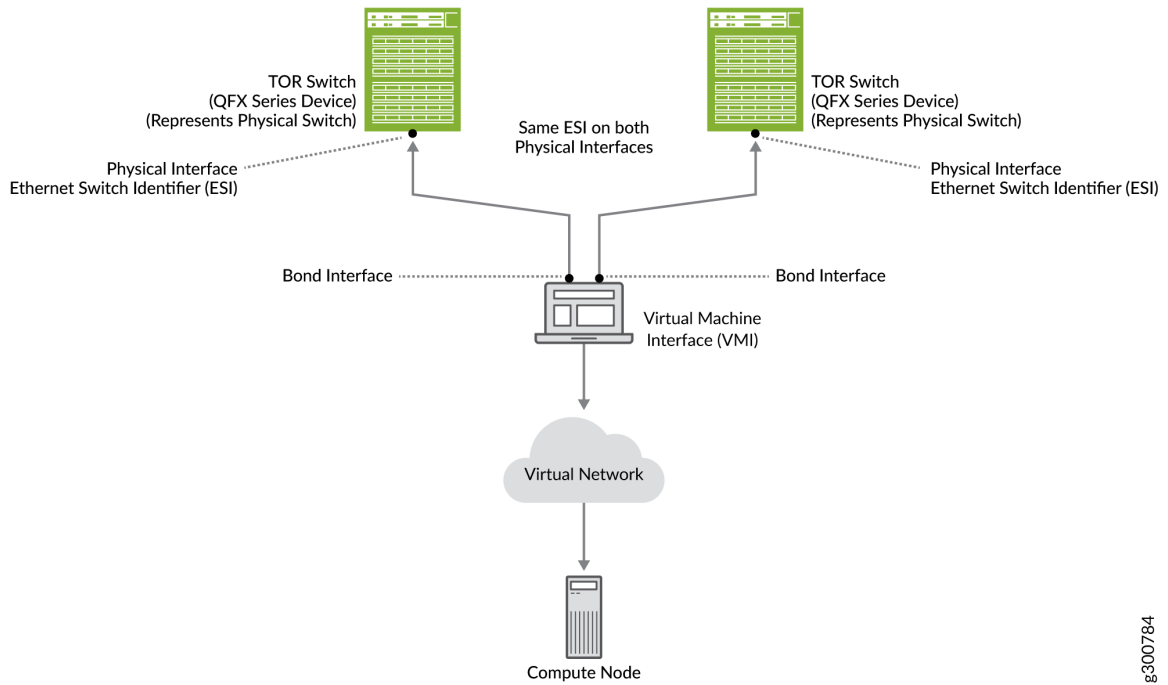
Bare metal servers connect to multiple TORs to establish redundancy (MLAG/multihomed configurations). Also, depending on the port bandwidth on the TOR and the NICs on the bare metal servers, multiple ports can be utilized to connect a bare metal server to the TOR (LAG configurations). These interfaces are also called as *bond* interfaces. On bonded interfaces, LACP protocol is enabled by default.

In LAG configuration, the two physical interfaces on the TOR switch (a QFX Series device) become members of a link aggregation group (LAG). The LAG connects to the aggregated Ethernet (AE) interface, which is again a physical interface that connects to the logical interface. The logical interface is connected to the virtual machine interface (VMI), which is connected to the virtual network (VN). The VN is connected to the node, which is the logical representation of a bare metal server.

In a multihomed configuration, a single port on a BMS connects to the physical interfaces on two QFX devices. The QFX devices have one physical interface each, both having the same Ethernet switch identifier (ESI). The physical interfaces are assigned the same ESI to enable the QFX device to recognize the interface as a multihomed interface. [Figure 138 on page 399](#) shows how BMS is connected to a TOR switch.

Contrail Networking Release 21.3 supports Layer 3 multihoming feature to provide high availability and load balancing of Contrail traffic without the MC-LAG or bond interfaces. For more information on Layer 3 multihoming, see [Layer 3 Multihoming](#).

Figure 138: Connectivity in Multihomed Configuration



6300784

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
21.3	Contrail Networking Release 21.3 supports Layer 3 multihoming feature to provide high availability and load balancing of Contrail traffic without the MC-LAG or bond interfaces.

RELATED DOCUMENTATION

[Bare Metal Server Management | 392](#)

[How Bare Metal Server Management Works | 396](#)

[Adding Bare Metal Server to Inventory | 400](#)

[Launching a Bare Metal Server | 402](#)

[Onboarding and Discovery of Bare Metal Servers | 404](#)

[Launching and Deleting a Greenfield Bare Metal Server | 405](#)

[Destination Network Address Translation for Bare Metal Servers | 406](#)

Adding Bare Metal Server to Inventory

The administrative user must follow these steps to add a bare metal server to Inventory from the Contrail Command user interface (UI):

1. Click **Infrastructure**> **Servers**.
A list of servers is displayed.
2. Click **Create** to create a server.
The Create Server page is displayed.
3. Select **Detailed** button from the Choose Mode options.
4. Select **Baremetal** button from the **Select workload type this server will be used for** options.
5. Enter a name for the host in the **Hostname** field.
6. Enter appropriate credentials for host in the **Credentials** field.
7. Select required kernel from the **Deploy Kernel** list.
8. Select required ram size from **Deploy Ramdisk** list.
9. Add the following values in the **Network Interfaces** fields:

Field	Action
Name	Assign a name to the Port.
MAC address	Enter the MAC address of the Port.
Device/ TOR- Interface	Select the Leaf/ TOR- Interface to which the Port is connected.
Enable PXE	Select this checkbox to enable PXE booting for only one Port.

10. Add the following values in the **Port Groups** field:

Field	Action
Name	Assign a name to the Port Group.

(Continued)

Field	Action
Member Interfaces	Add the ports that form the Port Group.

11. Add the following values in the **IPMI Info** fields:

Field	Action
IPMI Driver	Enter valid IPMI Driver name. The driver value for Openstack SKU: queens and ocata is <code>ispxe_ipmitool</code> . The driver value for Openstack SKU: rocky is <code>ipmi</code> . For more information, you can refer to Openstack document: Enabling drivers and hardware types .
IPMI Address	Enter the IPMI Address of the BMS Server.
IPMI Port	Enter port number on which IPMI is deployed. The default value is 623, as shown in the Contrail Command UI. You can update this to different IPMI port, according to your requirement.
IPMI Username	Enter IPMI Username.
IPMI Password	Enter IPMI Password.

12. Add the following values in the **Baremetal Properties** field based on the capacity of the server:

Field	Action
Memory mb	Enter RAM size of BMS Server in megabytes (Mb).
CPU's	Enter CPU count of BMS Server.
CPU Arch	Enter CPU Architecture of BMS Server. The default value is <code>x86_64</code> .
Local gb	Enter Disk Size of BMS Server in gigabytes (Gb).

(Continued)

Field	Action
Capabilities	This is the sets the capability of BMS Server. The default value is "boot_option:local".

13. Click **Create**. The **Servers** page is displayed with the list of servers created by the administrative user.

RELATED DOCUMENTATION

[Bare Metal Server Management | 392](#)

[How Bare Metal Server Management Works | 396](#)

[LAG and Multihoming Support | 398](#)

[Launching a Bare Metal Server | 402](#)

[Onboarding and Discovery of Bare Metal Servers | 404](#)

[Launching and Deleting a Greenfield Bare Metal Server | 405](#)

[Destination Network Address Translation for Bare Metal Servers | 406](#)

[Troubleshooting Bare Metal Servers | 411](#)

Launching a Bare Metal Server

The tenant user must follow these steps to launch a new bare metal server (BMS) from the Contrail Command UI:

1. Click **Workloads>Instances**.
The Instances page is displayed.
2. Click **Create** to create a new instance.
The Create Instance page is displayed.
3. Select **New Baremetal Server** as the Server Type.
4. Enter the following information in the **Create Instance** page:

Table 59: Add Existing Bare Metal Server Information

Field	Action
Instance Name	Enter a name for the BMS instance.
Select Boot Source	Select a Image or Instance Snapshot from the list.
Select Image	Select the BMS Image you created for the BMS from the list.
Select Flavor	Select the Flavor for the BMS from the list.
Select SSH Key	Select the SSH key for the BMS from the list, to login into SSH without password.
Availability Zone	Assign Availability Zone as nova-baremetal for BMS lifecycle management.
Count (1-10)	Assign values from 1 to 10, to spin the number of BMS instances.

5. Click **Create** to launch a new baremetal server.

RELATED DOCUMENTATION

[Bare Metal Server Management | 392](#)

[How Bare Metal Server Management Works | 396](#)

[LAG and Multihoming Support | 398](#)

[Adding Bare Metal Server to Inventory | 400](#)

[Onboarding and Discovery of Bare Metal Servers | 404](#)

[Launching and Deleting a Greenfield Bare Metal Server | 405](#)

[Destination Network Address Translation for Bare Metal Servers | 406](#)

[Troubleshooting Bare Metal Servers | 411](#)

Onboarding and Discovery of Bare Metal Servers

IN THIS SECTION

- [Onboarding of Bare Metal Servers | 404](#)
- [Discovery of Bare Metal Servers | 404](#)

BMS Manager supports onboarding and discovery of bare metal servers.

Onboarding of Bare Metal Servers

The BMS manager and Contrail Networking supports two types of bare metal servers deployments—greenfield deployments and brownfield deployments.

Greenfield deployments (LCM) are the bare metal servers that have not been deployed and requires to be managed by the BMS manager. These servers do not have an image installed on them. Greenfield servers do not have an IP address assigned.

Brownfield deployments (non-LCM) are the bare metal servers that are already deployed and are in active use by the tenants. These servers needs to be added to the Contrail Networking fabric management enrollment. These servers have IP addresses already assigned to them.

Discovery of Bare Metal Servers

The tenant user needs to onboard all bare metal servers that are already provisioned and configured. These bare metal servers are managed by the BMS management framework. The administrative users and the tenant users can onboard the servers by automatically discovering the servers or manually registering the servers.

Manual Discovery

Manual discovery is performed by registering all bare metal servers, their MAC addresses and their physical connectivity manually. This step is described in the *Administrative Workflow* section.

Auto Discovery

Auto Discovery of all servers can be achieved by utilizing the Ironic Inspector and the DHCP framework. When a server is powered on and physically connected to the TOR device, the DHCP frames are utilized to discover the MAC address as well as the connectivity information. Ironic Inspector uses the MAC address to match existing inventory. If a match is not found, an implicit registration of the server is performed, which is referred to as auto discovery.

RELATED DOCUMENTATION

[Bare Metal Server Management | 392](#)

[How Bare Metal Server Management Works | 396](#)

[LAG and Multihoming Support | 398](#)

[Adding Bare Metal Server to Inventory | 400](#)

[Launching a Bare Metal Server | 402](#)

[Launching and Deleting a Greenfield Bare Metal Server | 405](#)

[Destination Network Address Translation for Bare Metal Servers | 406](#)

[Troubleshooting Bare Metal Servers | 411](#)

[Terminating Ongoing Fabric Jobs | 113](#)

Launching and Deleting a Greenfield Bare Metal Server

This topic describes how to launch a greenfield bare metal server.

1. In the Contrail Web UI, select **Workloads > Instances > Create Instance**.
2. From the **Server Type** Field, select **New Bare Metal Server**.
3. Select the boot source, BMS image, and the BMS flavor available for the server type selected.
4. Click **Create**.

Following BMS launch, the BMS PXE boots from the ironic-provision network. The ironic-provision network is not visible to the tenant. BMS then connects to the provisioning network, connects to the TSN node, and gets a temporary IP address from the subnet of the provisioning network. This temporary IP address is not visible to the tenant. BMS downloads the boot image from the TFTP server and saves it locally for subsequent local boots. After the BMS is ready, it reboots. This time, the BMS boots from local image. During the second reboot, the BMS is disconnected from the ironic-

provision network and is connected to the tenant network. This process of transferring from the ironic-provisioning network to the tenant network is called *Network Flip*. Then, the TSN node provides the BMS an IP address from the tenant network. Once the BMS boots and is ready for use, it is connected to tenant network.

The tenant can delete a BMS when it is not needed in the network. When a BMS is disconnected from the tenant network, it is connected to the cleaning-network or the ironic-provisioning network. This network flip is done to prevent snooping of hackers when the BMS is being cleaned up. The ironic-provisioning network cleans up the server moves it back to the pool of available servers, to be ready for redeployment as a new BMS.

RELATED DOCUMENTATION

[Bare Metal Server Management | 392](#)

[How Bare Metal Server Management Works | 396](#)

[LAG and Multihoming Support | 398](#)

[Adding Bare Metal Server to Inventory | 400](#)

[Launching a Bare Metal Server | 402](#)

[Onboarding and Discovery of Bare Metal Servers | 404](#)

[Destination Network Address Translation for Bare Metal Servers | 406](#)

[Troubleshooting Bare Metal Servers | 411](#)

Destination Network Address Translation for Bare Metal Servers

IN THIS SECTION

- [Enabling DNAT in a Data Center Gateway | 407](#)
- [Extending a Public Virtual Network to the Data Center Gateway | 408](#)
- [Creating a Floating IP Address Pool | 408](#)
- [Mapping Floating IP Address to the Fixed IP address of the BMS Private Network | 409](#)

Contrail Networking Release 2005 supports Destination Network Address Translation (DNAT) for bare metal servers (BMS). DNAT enables traffic flow from a private network to the public network and also allows traffic flow from the public network to a private network. A private network can connect to a public network by routing traffic through a gateway device capable of performing DNAT.

In Contrail Networking, an MX Series device configured as a data center gateway (DC-GW) enables DNAT for a BMS deployed in a private network. The DC-GW device acts a bridge between a public network and a BMS by using a public IP address for the BMS. As part of DNAT, the DC-GW replaces the source IP address of the packet originating from the BMS with an IP address allocated from a public address pool configured on the MX Series device. The DC-GW then forwards the packet to the public network. Similarly, when the DC-GW also receives a packet from a public network, the DC-GW replaces the destination IP address of the packet with private IP address of the BMS and forwards the packet to the BMS.

Before you start using DNAT for BMS, you must enable DNAT in a DC-GW, create a public network and extend the network to the DNAT enabled DC-GW, create a floating IP address pool, and map a floating IP address to the BMS private network.

For more information on configuring an MX Series device as a DC-GW, see "[Configuring Data Center Gateway](#)" on page 220.

Enabling DNAT in a Data Center Gateway

An MX Series device is capable of DNAT for a BMS, when an MX Series device is configured as a data center gateway (DC-GW). You must perform the following steps to enable DNAT in a DC-GW device:

1. Navigate to **Infrastructure>Fabrics>Fabric Name**.

A list of fabric devices are displayed.

2. Select an MX Series from the list and click the **Options** icon.

Click **Edit** from the displayed list. The *Device Name Edit* page is displayed.

3. Expand the **Netconf Settings** section and enter the following values to add a service interface:

Field	Value
Username	Enter the username to add a service interface to the MX Series device.
Password	Enter the password to add a service interface to the MX Series device.
Junos Service Interface	Add a service interface created in Junos OS for an MX series device.

4. Click **Save** to enable DNAT in a DC-GW device.

Extending a Public Virtual Network to the Data Center Gateway

You must create a public virtual network that the DC-GW will use for DNAT. You must perform the following steps to create a public virtual network and extend the network to the DC-GW:

1. Navigate to **Overlay>Virtual Network**.

The **All networks** page is displayed.

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

The **Create Virtual Network** page is displayed.

3. Enter values in the fields as described in "[Create Virtual Network](#)" on page 82.

4. Expand the **Advanced** section.

Select **External** to configure the network as a public virtual network.

In the **Extend to Physical Router(s)** field, select the DC-GW device enabled with DNAT for BMS.

5. Click **Create** to create a public network extended to the DC-GW.

Creating a Floating IP Address Pool

You must create a floating IP address pool, which enables IP address mapping between the BMS deployed in a private virtual network and the DC-GW public virtual network. You must perform the following steps to create a floating IP address pool for the public virtual network:

1. Navigate to **Overlay>Floating IPs**.

The **All Floating IPs** tab is displayed.

2. Click the **Floating IP Pools** tab.

The **Floating IP Pools** page is displayed.

3. Click **Create** to create a floating IP pool for the public virtual network.

4. Enter the following values in the fields:

Field	Value
Name	Enter a name for the floating IP address pool.

(Continued)

Field	Value
Network	Select the public network you want to assign the floating IP address pool.
Description	Add a description for the floating IP address pool.

5. Click **Save** to create a floating IP address pool extended to the public network.

Mapping Floating IP Address to the Fixed IP address of the BMS Private Network

Mapping a floating IP address to the fixed IP address of the BMS enables the BMS to exchange data packets with a public network through a DC-GW. To map the floating IP address to the fixed IP address of the BMS you must perform the following steps:

NOTE: If a virtual port is not assigned to the BMS, follow the steps described in "[Configuring Virtual Port Groups](#)" on page 238 to create a virtual port for the BMS.

1. Navigate to **Overlay > Virtual Ports**.
A list of virtual port groups is displayed.
2. Click **Edit** icon of the virtual port assigned to the BMS.
The **Edit Virtual Port** page is displayed.
3. In the **Floating IPs** field, select the floating IP address, which is mapped to the public network.
4. Click **Save**. The **Virtual Ports** page is displayed.

Figure 139: Floating IP Address Mapped to the Fixed IP Address of the BMS

Virtual Ports								
	NAME	UUID	TAGS	NETWOI	FIXED IF	FLOATING IPS	DEVICE	VIRTUAL
<input type="checkbox"/>	fb7aa92a-882c-4a...	fb7aa92a-882c-4a...	-	vn2	XX.XX.XX	-	comp...	...
<input type="checkbox"/>	c05ee757-5135-4f...	c05ee757-5135-4f...	-	vn1	XX.XX.XX	-	comp...	...
<input type="checkbox"/>	a637f46e-c98e-4d...	a637f46e-c98e-4d...	-	vn1	10.20...	-	comp...	...
<input type="checkbox"/>	vpg111-1222-unta...	cfc3bab-1b58-4d...	-	vn1	1.1.1...	10. XX.XX.XX	vpg111	

The floating IP address is now mapped to the BMS private network.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
2005	Contrail Networking Release 2005 supports Destination Network Address Translation (DNAT) for bare metal servers (BMS). DNAT enables traffic flow from a private network to the public network and also allows traffic flow from the public network to a private network.

RELATED DOCUMENTATION

[Configuring Data Center Gateway | 220](#)

[Bare Metal Server Management | 392](#)

[How Bare Metal Server Management Works | 396](#)

[LAG and Multihoming Support | 398](#)

[Adding Bare Metal Server to Inventory | 400](#)

[Launching a Bare Metal Server | 402](#)

[Onboarding and Discovery of Bare Metal Servers | 404](#)

[Launching and Deleting a Greenfield Bare Metal Server | 405](#)

[Troubleshooting Bare Metal Servers | 411](#)

Troubleshooting Bare Metal Servers

This topic provides the steps to troubleshoot BMS.

- **Follow these steps to troubleshoot some of the common issues:**

- Verify that the following objects are created:
 - When the BMS is in provisioning state (when BMS is booting for the first time), there should be two neutron ports—one on provisioning network and another on the tenant network. Run the `openstack port list/show` command to view the list of ports.

The port connected to the provisioning network should have `local_link_information` displaying the name of the QFX or TOR and the port to which the bare metal server connected.

- After network flip, only one port should be present. The port connected to provisioning network should be deleted.
- Verify that the logical Interface(s) are created. Run the `curl http://localhost:8082/logical-interfaces` command to view the logical interfaces. The logical interface should point to the correct physical interface.
- **Follow these steps to troubleshoot LAG interfaces (AE interfaces):**
 - Ensure that an aggregated Ethernet physical interface is created. Run the `curl http://localhost:8082/physical-interfaces` command to verify. The AE interface name starts with `ae`.
 - Ensure that logical Interface is created. Run the `curl http://localhost:8082/logical-interfaces` command.

The logical interface should have parent reference pointing to the `ae` physical interface.

- Ensure that a link aggregation group (LAG) is created. Run the `curl http://localhost:8082/link-aggregation-group` command to verify.

- **Follow these steps to troubleshoot multihomed interfaces:**

- Ensure that two logical Interfaces are created. Run the `curl http://localhost:8082/logical-interfaces` command to verify.

Each logical interface should have a parent reference pointing to the physical interface. The Ethernet segment identifier (ESI) should be set to the same value for both physical Interfaces.

- **Follow these steps if you get the error message No Valid Host Found when you launch a BMS server.**

- Run the `openstack baremetal node list/show` command to verify that the nodes are registered on Ironic and are not in error state.
- Run the `openstack baremetal port list/show` command to verify that ports for the nodes are registered.
- Run the `openstack baremetal portgroup list/show` command to verify that the port groups (in case of LAG/MH deployments).
- Run the `openstack flavor list/show` command to verify the BMS flavors details to ensure that the flavor matches with the node specification.
- Review the `api-server` logs for errors. The log contains errors of there is a duplicate MAC address or the physical interface is not configured.
- Review the `ironic-conductor` logs for errors. For example, `PXE_ENABLED` port is not found.
- **Follow these steps if the server does not boot or if the server remains in boot state:**
 - Verify whether the server is assigned an IP address on the provisioning network.
 - If an IP address is not assigned, verify whether the TSN node is reachable.
 - If an IP address is assigned, check whether the TFTP boot server is reachable.

In either case, you can use the `tcpdump` tool to review the TCP packets to check whether the bare metal server can reach these servers.
 - Follow these steps if the server was assigned an IP address and is booted on provisioning network, but remains the same state. That is, network flip does not happen.
 - Verify the `ironic-conductor` logs to see whether Ironic Python Agent (IPA) on the bare metal server is able to communicate with Ironic Conductor.
 - Check whether the image was built correctly with the correct IPA.

RELATED DOCUMENTATION

[Bare Metal Server Management | 392](#)

[How Bare Metal Server Management Works | 396](#)

[LAG and Multihoming Support | 398](#)

[Adding Bare Metal Server to Inventory | 400](#)

[Launching a Bare Metal Server | 402](#)

[Onboarding and Discovery of Bare Metal Servers | 404](#)

Launching and Deleting a Greenfield Bare Metal Server | 405

Destination Network Address Translation for Bare Metal Servers | 406