JUNIPER
NETWORKS

Engineering
Simplicity

Contrail® Networking

Contrail Networking Fabric Lifecycle Management Guide

Published
2021-01-06

Release
1912

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

YEAR 2000 NOTICE

END USER LICENSE AGREEMENT

Table of Contents

5   **Managing Data Center Devices**

6

Integrating VMware with Contrail Networking Fabric

7 **Extending Contrail Networking to Bare Metal Servers**

About the Documentation

IN THIS SECTION

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
, 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 Deployment Guide	Provides information about the next steps to be taken after a successful installation of Contrail.
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 Analytics and Troubleshooting Guide	Provides information about AppFormix and Contrail analytics.

Documentation and Release Notes

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

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

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

Documentation Conventions

defines notice icons used in this guide.

Table 2: Notice Icons

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

defines the text and syntax conventions used in this guide.

Table 3: Text and Syntax Conventions

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

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

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

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		    route default {
		      nexthop *address*;
		      retain;
		    }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		

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

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

Documentation Feedback

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

• Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the Juniper Networks TechLibrary site, and do one of the following:
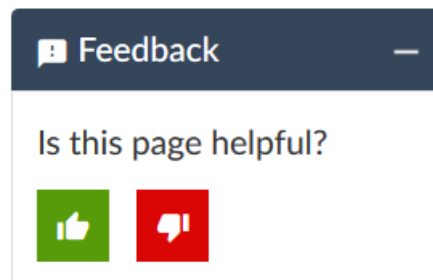
  • Click the thumbs-up icon if the information on the page was helpful to you.

  • Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

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

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.

- Product warranties—For product warranty information, visit https://www.juniper.net/support/warranty/.

- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

- Find CSC offerings: https://www.juniper.net/customers/support/

- Search for known bugs: https://prsearch.juniper.net/

- Find product documentation: https://www.juniper.net/documentation/

- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/

- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/

- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/

- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/

- Create a service request online: https://myjuniper.juniper.net

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

Creating a Service Request with JTAC

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

- Visit https://myjuniper.juniper.net.

- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

1
CHAPTER

Overview

Understanding Underlay Management

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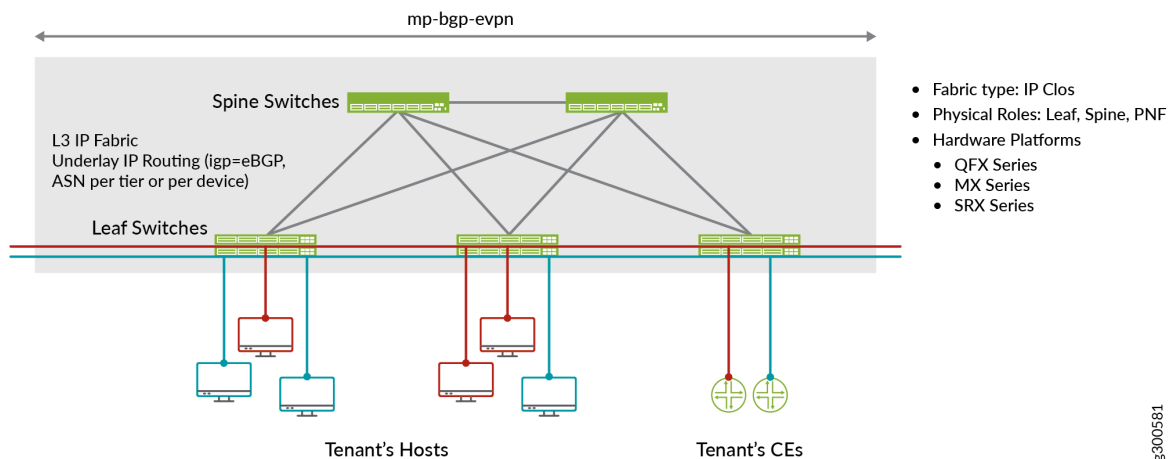
Contrail Networking helps you provision both greenfield and brownfield devices to form IP Clos networks. You can bring up all factory-default greenfield devices using zero-touch-provisioning to form an operational IP Clos network with underlay connectivity. However, unlike greenfield devices, brownfield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

RELATED DOCUMENTATION

Understanding Underlay Management | 15

Understanding Bare Metal Server Management | 218

Configuring QFX Series Devices as Data Center Gateway | 162

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:

2
CHAPTER

Fabric Administrative Tasks

Image Management

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 2: Device Images
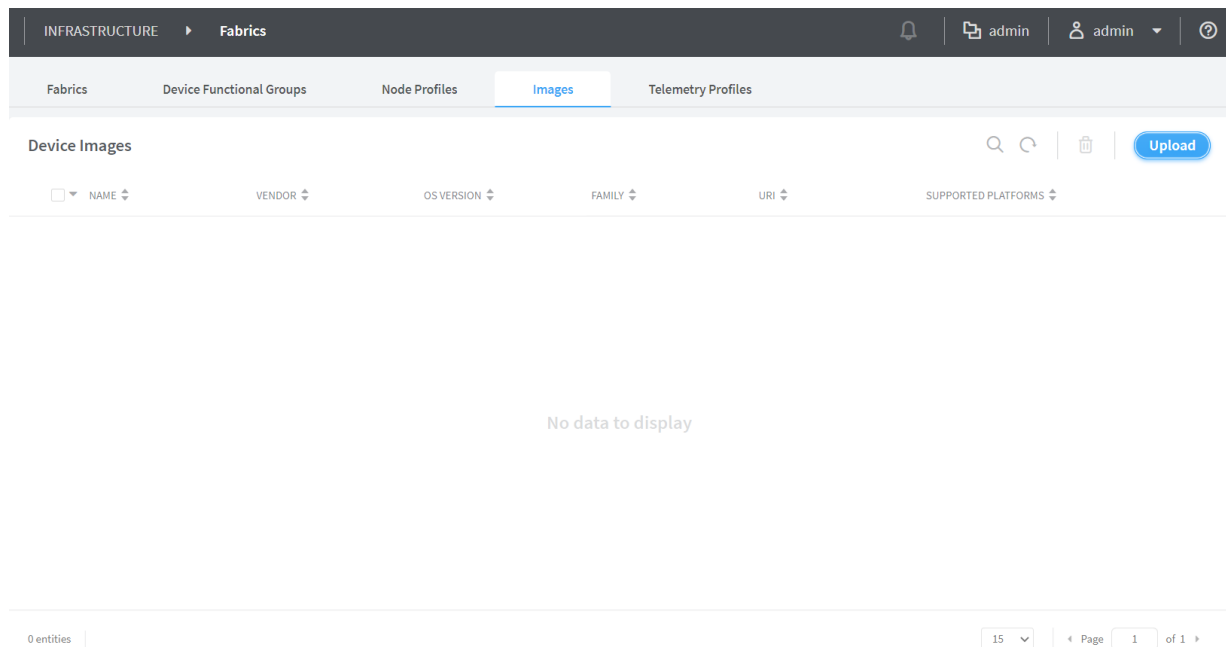
2. Click **Upload**.

   The Upload Image pop-up is displayed. See .

Figure 3: Upload Image

Upload Image

| Device Image | Tags | Permissions |

Name*

Pick a File* ⑦

Drag file here or browse

Vendor Name* ⑦

juniper

Device Family* ⑦

Supported Platforms* ⑦

Os Version* ⑦

Image MD5 ⑦

Image SHA1 ⑦

Cancel Upload

3. Enter the following information given in .

Table 4: Upload Image Details

Field	Action
Name	Enter a name for the device image.
Pick a file	Click **Upload** and navigate to the local directory and select the device image file. Click **Open** to confirm selection.

Table 4: Upload Image Details *(continued)*

Field	Action
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

Hitless Software Upgrade of Data Center Devices Overview

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 is lost during image upgrade on devices in the data center.

RELATED DOCUMENTATION

Performing Hitless Software Upgrade on Data Center Devices | 25

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**. lists all the mandatory parameters that must be entered to upload a software image.

Table 5: Upload Image Fields

Field	Description

Table 5: Upload Image Fields *(continued)*

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.
Os Version	Enter the OS version of the image. For example, 18.1R2.

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.

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 4: Select Device > Assign Devices

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 5: Select Device > Parameters

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 6: Testing > Test Run

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 devices are upgraded in different batches. 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 7: Testing > Strategy Plan

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 8: Testing > Strategy Plan Confirmation
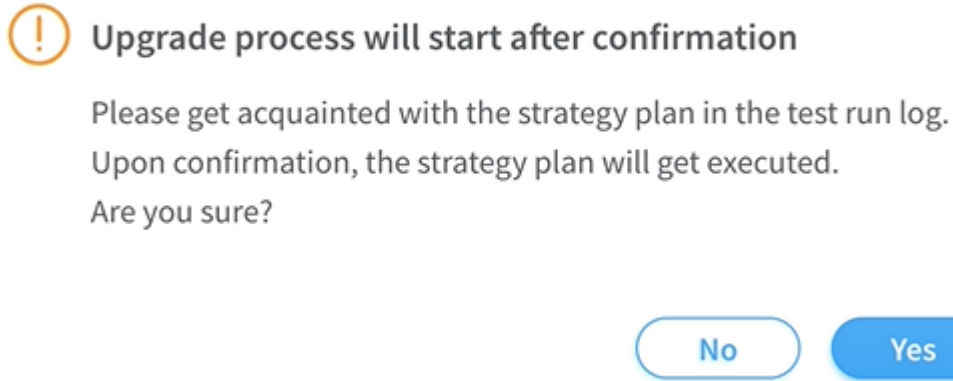
10. Click on each device to view the image upgrade progress for that device. Click the device again to toggle back to display the overall image upgrade progress of all devices. Table 6 on page 32 displays the states displayed during the course of the upgrade.

Table 6: 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.

Table 6: Image Upgrade Progress States *(continued)*

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 9: Upgrade
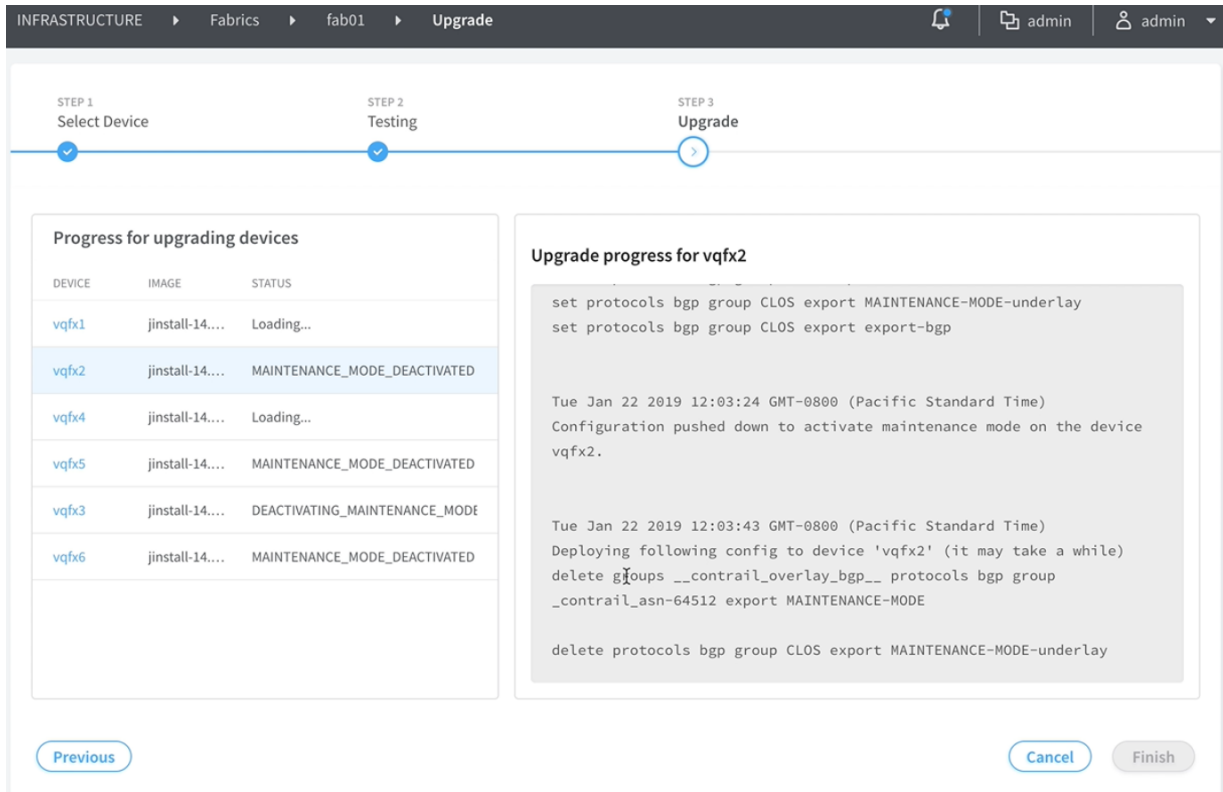
11. Click **Finish** when all the devices have been upgraded.

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

3
CHAPTER

Zero-Touch-Provisioning

Create a Fabric

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

Follow these steps to create a new fabric:

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

   The Fabrics page is displayed. See Figure 10 on page 36.

Figure 10: Fabrics Page
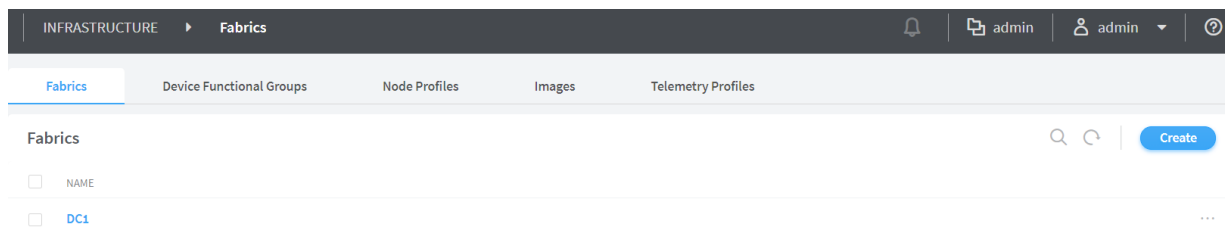
2. Click **Create**.

   You are prompted to select a provisioning option. See Figure 11 on page 37.

Figure 11: Select Provisioning Option

Select provisioning option

New Fabric

Wizard takes you through deployment of new devices which require discovery, zero touch provisioning(ZTP) and complete configuration.

Existing Fabric

Import existing deployed devices by discovery.

Cancel   Provision

- Click **New Fabric** to deploy new (greenfield) devices. See Figure 12 on page 41.

- Click **Existing Fabric** to import existing (brownfield) devices by discovery. See Figure 13 on page 45.

Click **Provision**.

The Create Fabric page is displayed.

If you select **New Fabric** as the provisioning option, see "Provisioning Option - New Fabric" on page 37.

If you select **Existing Fabric** as the provisioning option, see "Provisioning Option - Existing Fabric" on page 41.

Provisioning Option - New Fabric

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

Enter the information as given in Table 7 on page 37.

Table 7: Provisioning Option - New Fabric

Field	Action
Name	Enter a name for the fabric.
Device credentials	Enter root user password.

Table 7: Provisioning Option - New Fabric *(continued)*

Field	Action
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.
Device Info	Upload device information file. Navigate to the local directory and select the device information file. Click **Open** to confirm. For a sample YAML file, see Sample YAML File Snippet on page 40.
Node profiles	Add node profiles. You can add more than one node profile. 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 78. NOTE: The supported node profiles are **juniper-mx**, **juniper-qfx10k**, **juniper-qfx10k-lean**, **juniper-qfx5k**, **juniper-qfx5k-erb-only**, **juniper-qfx5k-lean**, **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 138.
Upgrade devices during the process?	Select the **Upgrade devices during the process?** check box as given in Figure 12 on page 41 to enable the OS Version list. Starting with Contrail Networking Release 1907, you can upgrade a device during the ZTP process.

Table 7: Provisioning Option - New Fabric *(continued)*

Field	Action
OS Version	Select the OS version you want to upgrade the device to, from the **OS Version** list. The **OS Version** list is enabled when you select the **Upgrade devices during the process?** check box. NOTE: The options in the **OS Version** list are the OS versions of the images that you uploaded.
VLAN-ID Fabric Wide Significance	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. 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 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 182. For more information on supported hardware platforms and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 138.

Table 7: Provisioning Option - New Fabric *(continued)*

Field	Action
Management subnets	Enter the following information to auto-assign management IP addresses to devices: **CIDR**—Enter CIDR address. **Gateway**—Enter gateway address.
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. • Enter minimum value in **ASN From** field. • Enter maximum value in **ASN To** field.
Fabric subnets (CIDR)	Enter fabric CIDR address. Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.
Loopback subnets (CIDR)	Enter loopback subnet address. Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.
PNF Servicechain subnets (CIDR)	Enter PNF device CIDR address. Starting in Contrail Networking Release 5.1, enter the subnet for allocating IP addresses in the **PNF Servicechain subnets** field to establish EBGP session between PNF device and SPINE switch.

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 12: Deploy Greenfield Devices

Click **Next**.

The Discovered devices page is displayed.

Provisioning Option - Existing Fabric

Enter the information as given in .

Table 8: Provisioning Option - Existing Fabric

Field	Action
Name	Enter a name for the fabric.
Username	Enter a username for the device.
Password	Enter a password for the device.
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.
Node profiles	Add node profiles. You can add more than one node profile. 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 78. **NOTE:** The supported node profiles are **juniper-mx**, **juniper-qfx10k**, **juniper-qfx10k-lean**, **juniper-qfx5k**, **juniper-qfx5k-lean**, **juniper-qfx5120**, **juniper-qfx5k-erb-only**, 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 138.

Table 8: Provisioning Option - Existing Fabric *(continued)*

Field	Action
VLAN-ID Fabric Wide Significance	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.

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.

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.

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 182.

For more information on supported hardware platforms and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 138 |

Table 8: Provisioning Option - Existing Fabric *(continued)*

Field	Action
Management subnets	Enter the following information: **CIDR**—Enter CIDR network address. **Gateway**—Enter gateway address. **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.
Loopback subnets (CIDR)	Enter loopback address. Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices. 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 48.
PNF Servicechain subnets (CIDR)	Enter PNF device CIDR address. Starting in Contrail Networking Release 5.1, enter the subnet for allocating IP addresses in the **PNF Servicechain subnets** field to establish EBGP session between PNF device and SPINE switch.

Figure 13: Import Brownfield Devices

Click **Next**.

The Discovered devices page is displayed.

For more information on device discovery, see .

Release History Table

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

Device discovery is initiated as soon as you click **Next** on the Fabrics page. For more information on using zero-touch-provisioning (ZTP) to create a new fabric, see "Create a Fabric" on page 36.

- If you have followed the steps provided in the **Provisioning Option - New Fabric** (greenfield) section of the "Create a Fabric" on page 36 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 36 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. The list of devices discovered is listed in the Discovered devices page.

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

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

   NOTE: You can select more than one device.

2. Click **Add**.

The device is added to the fabric.

Click **Next** to assign roles.

The Assign to devices page is displayed. For more information on assigning roles to devices, see .

RELATED DOCUMENTATION

Assign a Role to a Device

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 14: Sample Data Center Topology
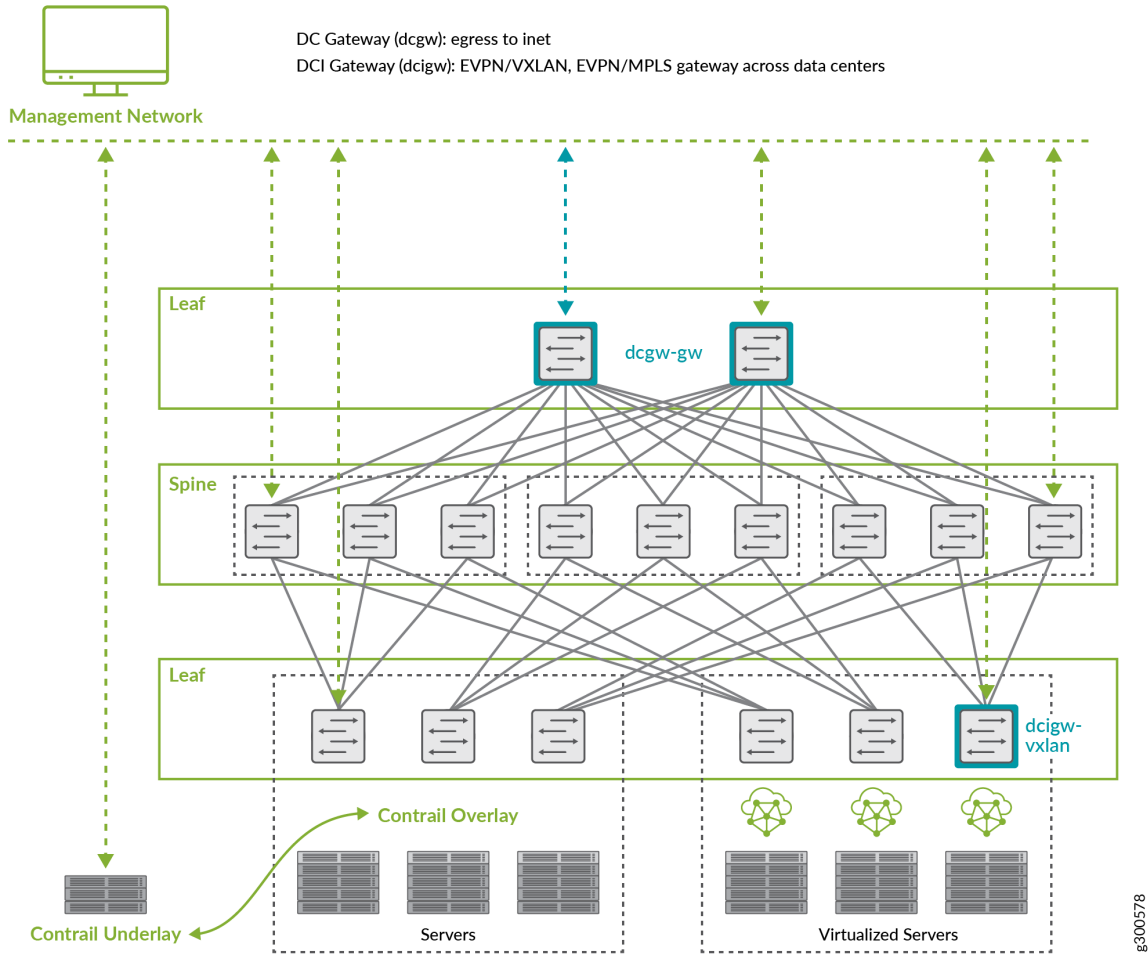
Management Network

DC Gateway (dcgw): egress to inet
DCI Gateway (dcigw): EVPN/VXLAN, EVPN/MPLS gateway across data centers

Leaf

dcgw-gw

Spine

Leaf

dcigw-vxlan

Contrail Overlay

Contrail Underlay

Servers

Virtualized Servers

g300578

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 138.

Figure 15: Supported Physical Roles and Routing Bridging Roles
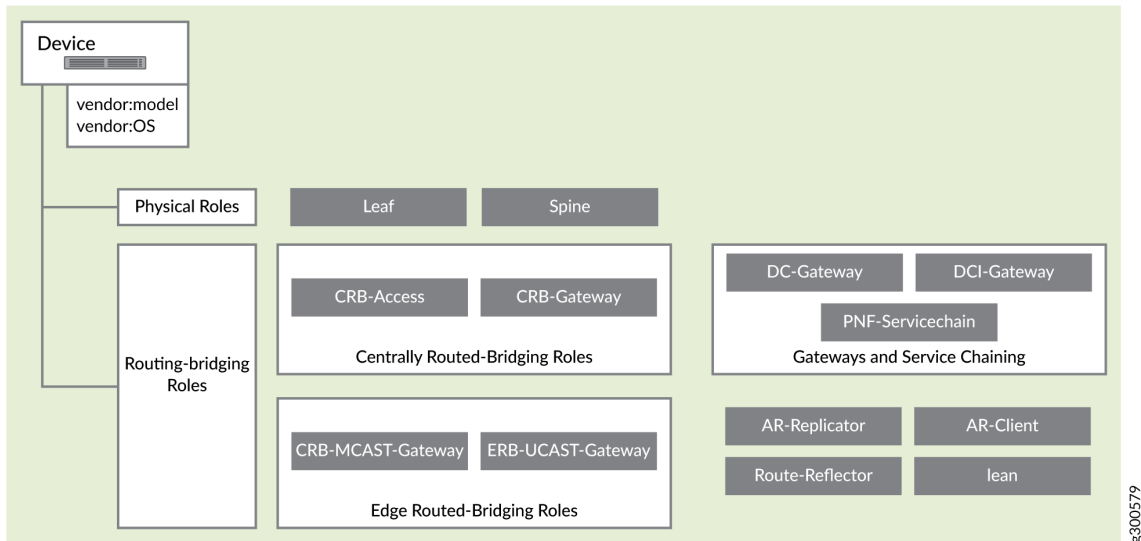
After you have completed the steps provided in the "Discover a Device" on page 46 topic, you can assign roles to the devices from the Assign to devices page.

Follow these steps to assign roles to devices:

1. Select the device you want to assign a role to by selecting the check box next to the device name.

2. Click the **Assign** icon at the end of the row to assign roles.

   NOTE: In Contrail Networking Release 1907, you can assign roles only to similar device types at the same time.

   To assign roles to similar device types, select the check box next to the device name and then click **Assign Role**.

   The Assign role to devices pop-up is displayed.

3. Select a physical role type from the **Physical Role** list.

4. Select a routing bridging role from the **Routing Bridging Role** list.

   These are the supported roles: CRB-Access, CRB-Gateway, DC-Gateway, Route-Reflector, ERB-UCAST-Gateway, DCI-Gateway, CRB-MCAST-Gateway, PNF-Servicechain, AR-Replicator, and AR-Client roles.

   For more information, see Centrally-Routed Bridging Overlay Design and Implementation.

Starting with Contrail Networking Release 1907, you can assign the AR-Replicator and AR-Client roles to enable assisted replication on the QFX10000 devices in a datacenter. Assisted replication feature optimizes replication of ingress broadcast, unknown unicast, and multicast (BUM) traffic received from the CE interfaces by replicating BUM traffic towards a single EVPN core Replicator PE (a QFX10000 device) rather than sending it to all the PE devices for replication.

5. Click **Assign** to confirm.

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

   The Autoconfigure page is displayed. For more information, see "Manage Device Configuration" on page 51.

Release History Table

Release	Description
1907	Starting with Contrail Networking Release 1907, you can assign the AR-Replicator and AR-Client roles to enable assisted replication on the QFX10000 devices in a datacenter.

RELATED DOCUMENTATION

Manage Device Configuration

After you assign device roles, you initiate the auto-configuration job by clicking **Autoconfigure** on the Assign to devices page. 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 **Finish**.

You can view node profile information from the Node Profiles page of the Contrail Command UI. For more information, see "View Node Profile Information" on page 78.

RELATED DOCUMENTATION

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 to delete a fabric.

    The Delete confirmation pop-up is displayed.

4. Click **Delete** to confirm.

RELATED DOCUMENTATION

Provisioning 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.

> **NOTE:** The minimum required version of Junos OS for QFX5000 and QFX10000 Series devices is 18.1R3-S5 or higher.

Sample Topology
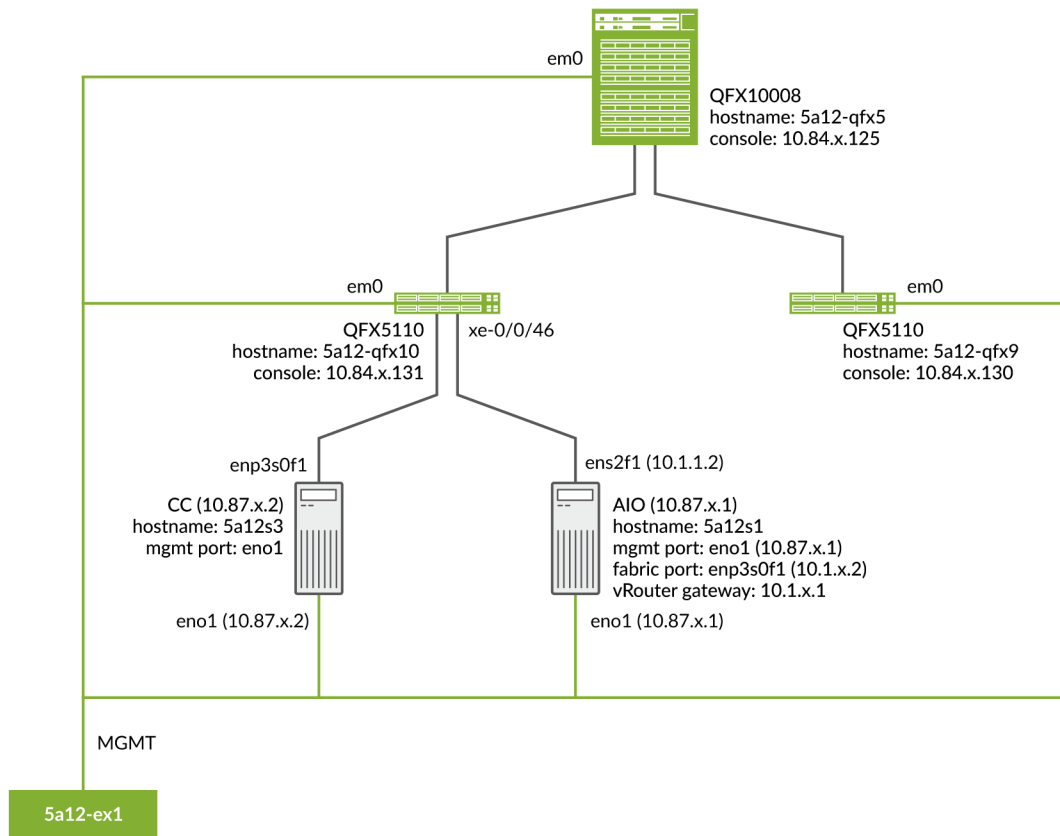
Prerequisites

These are example parameters. The interface name can be different based on your deployment.

- 5a12s3-node1:

  - Install CentOS 7.6.

  - Configure *eno1* port with the static IP **10.87.x.2/27**.

    ```
    HWADDR=ac:xx:xx:xx:xx:88
    NM_CONTROLLED=no
    BOOTPROTO=none
    DEVICE=enp2s0f0
    ONBOOT=yes
    IPADDR=10.87.x.2
    NETMASK=255.255.255.224
    GATEWAY=<GATEWAY_IP>
    ```

- 5a12s1-node1:

- 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=<GATEWAY_IP>
```

- 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=<GATEWAY_IP>
```

To provision fabric devices using ZTP via Contrail Command UI:

1. Install Contrail Command. For details, see *Installing Contrail Command*.

2. Install a Contrail cluster using Contrail Command. For details, see *Installing a Contrail Cluster Using Contrail Command*.

3. 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

4. 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 9: 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 9: 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*

root user password

Overlay ASN (iBGP)*

64512

Device Info*

Download Template: (*.yml)

⊕ Upload .yaml or .yml

Node profiles*

juniper-mx ×
juniper-qfx10k ×
juniper-qfx5k ×
juniper-qfx5k-lean ×
juniper-srx ×

Management subnets

CIDR*

Enter valid CIDR

Gateway*

Enter valid IPv4

Cancel   Next

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.
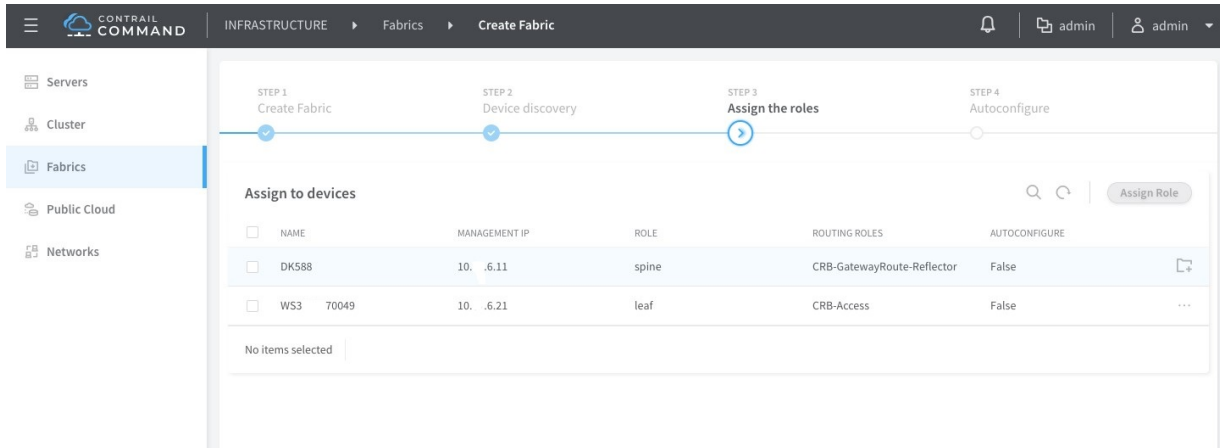
To configure underlay network via Contrail Command UI:

1. Create provisioning infrastructure network.

   a. Click **Networks**.

   b. Create a network by entering the required details.
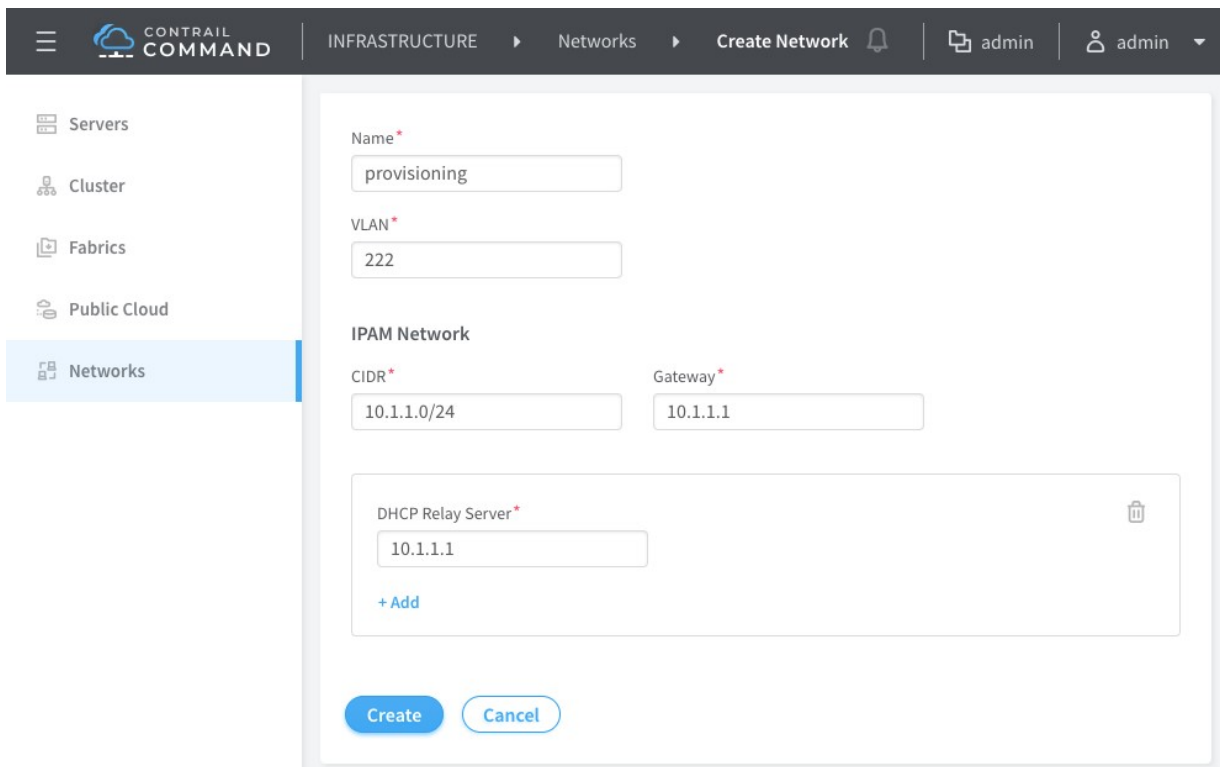
2. Import server topology.

   a. Click **Servers**.

b. Click **Import**.

c. Upload the **server topology** file.

Sample server topology yaml file:

```
nodes:
  - name: 5a12s1-node1
    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 10: Required Fields for server topology yaml file

Field	Details
name	Name of the infrastructure BMS node

Table 10: Required Fields for server topology yaml file *(continued)*

Field	Details
type	Type of the infrastructure BMS node. It must be "baremetal"
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 11: 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
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

Table 11: Required fields for Server Node Profile *(continued)*

Field	Details
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.

Release History Table

Release	Description
1907	Contrail Networking Release 1907 introduces the ability to configure hostnames to the devices being onboarded

4

CHAPTER

Fabric Configuration

Onboard Devices

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

1. Click **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.

Figure 16: Select Provisioning Option

4. Click **Provision**.

   The Create Fabric page is displayed.

5. Enter the following information:

   Table 12: Provision Existing Fabric

Field	Action
Name	Enter a name for the fabric.
Username	Enter a username for the device.
Password	Enter a password for the device.

Table 12: Provision Existing Fabric *(continued)*

Field	Action
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.
Node profiles	Add node profiles. You can add more than one node profile. All preloaded node profiles are added to the fabric by default. You can remove a node profile by clicking **X** on the node profile.
Management subnets	Enter the following information: **CIDR**—Enter CIDR network address. **Gateway**—Enter gateway address. 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.
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. • Enter minimum value in **ASN From** field. • Enter maximum value in **ASN To** field.
Fabric subnets (CIDR)	Enter fabric CIDR address. NOTE: Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.
Loopback subnets (CIDR)	Enter loopback address. NOTE: Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.

Table 12: Provision Existing Fabric *(continued)*

Field	Action
PNF Servicechain subnets (CIDR)	Enter PNF device CIDR address. NOTE: Starting in Contrail Networking Release 5.1, enter the subnet for allocating IP addresses in the PNF Servicechain subnets field to establish EBGP session between PNF device and SPINE switch.

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.

Figure 17: Device Discovery Progress Bar

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

7. Select the device(s) 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. Click the **Assign** icon at the end of the row to assign roles.

   The Assign role to devices pop-up is displayed.

10. Assign physical roles and routing bridging roles.

    For Spine Devices:

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

For Leaf Devices:

- Select **leaf** from the Physical Role list.

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

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.

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.

11. Click **Assign** to confirm selection and then 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.

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

    The onboarding job is now complete.

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.

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

    The All Networks page is displayed.

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

    The Create Virtual Network page is displayed.

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

4. 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 80.

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

5. 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.

6. Enter subnet information as given in .

Table 13: 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.

7. 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.

8. 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.

9. Enter fat flows information. See .

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 14: 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.
	Prefix	Enter source subnet prefix length. 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.

Table 14: Configure Fat Flow *(continued)*

Field		Action
Prefix Aggregation Destination	**Destination Subnet**	Enter the destination IP address.
		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.
		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.
	Prefix	Enter the destination subnet prefix length.
		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.

10. 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 15 on page 74.

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

Table 15: 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.

11. Enter advanced configuration information as given in Table 16 on page 74.

Table 16: 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.

Table 16: Advanced Configuration *(continued)*

Field	Action
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.
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 16: Advanced Configuration *(continued)*

Field	Action
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.
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.

12. 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. Click **Overlay**>**Logical Routers**.

   The Logical Routers page is displayed.

2. Click **Create**.

   The Create Logical Router page is displayed.

3. Enter the following information.

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.
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.
Logical Router Type	Select **SNAT Routing** or **VXLAN Routing** from the list.
Connected Networks	Select the networks that you want to connect this logical router to.
Public Logical Router	(Optional) Select this check box if you want the logical router to function as a public logical router.
VxLAN Network Identifier	Enter VXLAN network identifier in the range from 1 through 16,777,215. This field is disabled by default.

Field	Action
Route Target(s)	Click **+Add** to add route targets. Enter Autonomous System (AS) number in the **ASN** field. • 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 *L* or *l (lower-case L)* 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 *L* or *l (lower-case L)* 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. Enter route target in the **Target** field. • 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.

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

   The Logical Routers page is displayed.

5. Repeat Step 3 and Step 4 to create another logical router.

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 18: 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 138.

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 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 to define a policy rule.

   You can define more than one rule for a policy.

   Table 17: 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 **>**.

Table 17: Define Policy Rule *(continued)*

Field	Action
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.
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.

Monitoring Fabric Jobs

In Contrail Networking Release 1912, you can view detailed information, status, and logs of all active, failed, and completed fabric jobs for the past 24 hours in Contrail Command. You can also terminate ongoing jobs from the job status monitoring page.

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.

Release History Table

Release	Description
1912	In Contrail Networking Release 1912, you can view detailed information, status, and logs of all active, failed, and completed fabric jobs for the past 24 hours in Contrail Command.

RELATED DOCUMENTATION

Terminating Ongoing Fabric Jobs | 84

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 "Provisioning Fabric Devices Using End-to-End ZTP" on page 53.

- Importing existing brownfield devices and deploying new greenfield devices. For more information, see "Create a Fabric" on page 36.

- Initiating auto-configuration job. For more information, see "Manage Device Configuration" on page 51.

- Device image upgrade job. For more information, see "Performing Hitless Software Upgrade on Data Center Devices" on page 25.

- Activating or deactivating Maintenance Mode job. For more information, see "Activating Maintenance Mode on Data Center Devices" on page 184.

- Discovery of servers. For more information, see "Onboarding and Discovery of Bare Metal Servers" on page 229 and "Fabric Discovery and ESXi Discovery by Using Contrail Command" on page 207.

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 19 on page 85 displays that the image upgrade job is in progress.

Figure 19: Image Upgrade Job In Progress
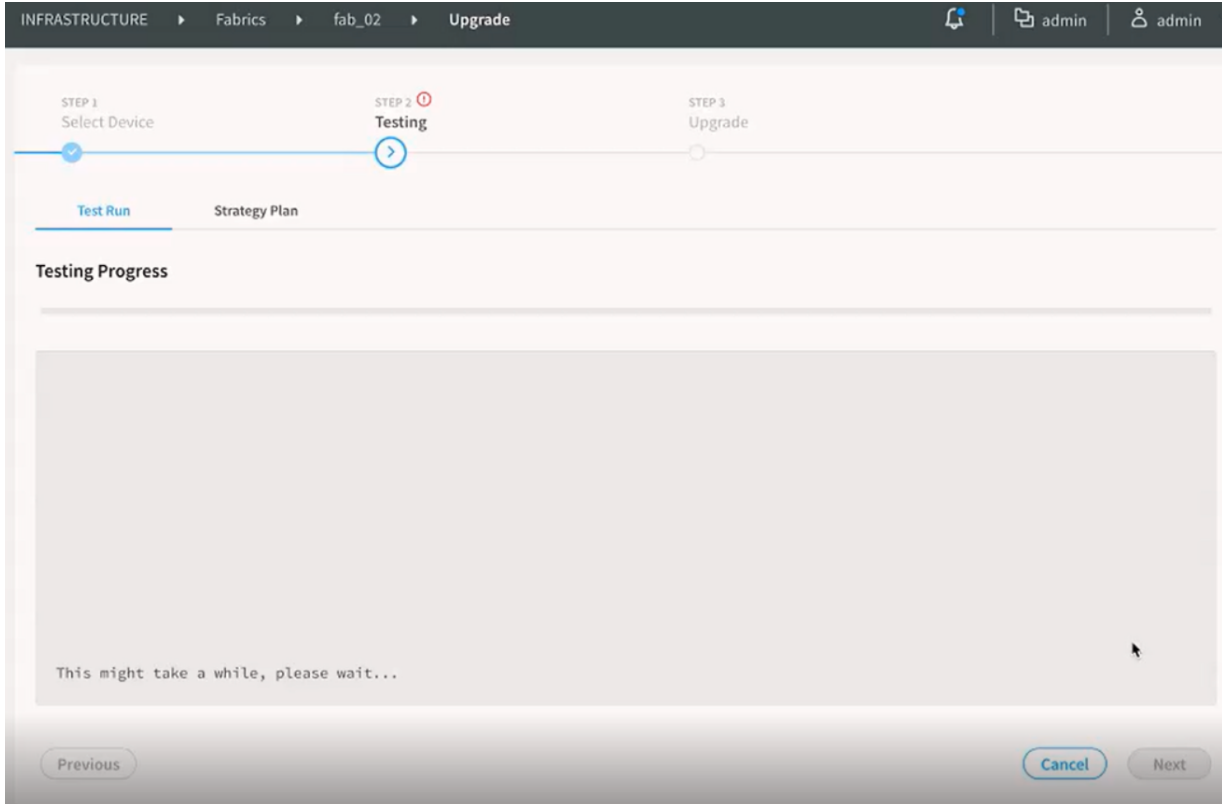
2. To terminate this image upgrade job, click **Cancel**.

3. In the pop-up that is displayed, click **Abort and Exit**.

The image upgrade job is terminated.

Release History Table

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 | 83

Using HA Cluster to Manage Fabric

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 20: HA Topology for ZTP Subnet
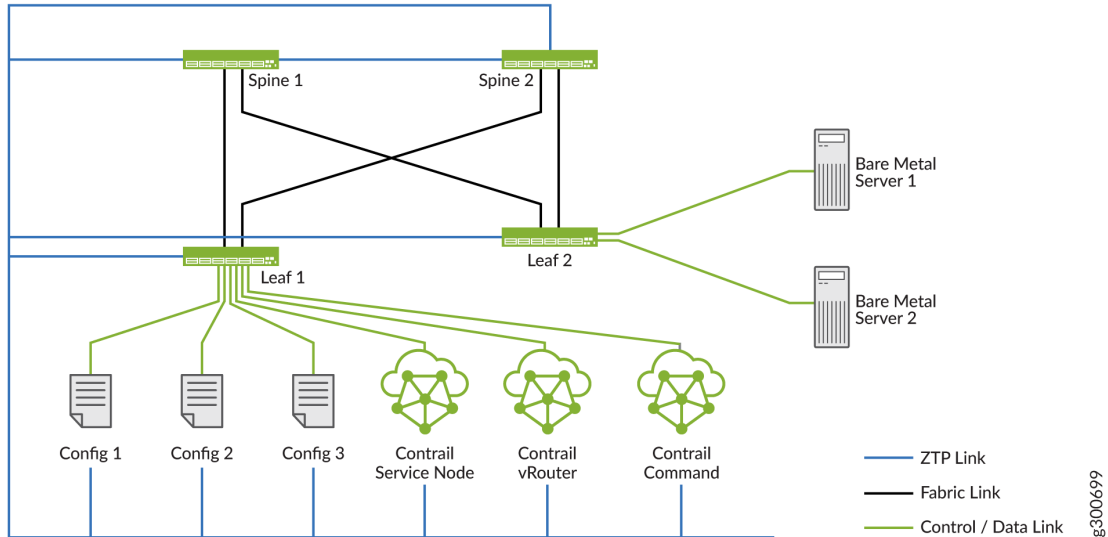
shows an HA topology for ZTP. All spine and leaf switches, controllers, and nodes that are part of the deployment, are on the same ZTP subnet. All controllers are connected to a single ToR switch. The HA cluster is connected to a leaf switch.

Release History Table

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 | **17**

Discover a Device | **46**

Provisioning Fabric Devices Using End-to-End ZTP | **53**

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 .yml 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 in the YAML file as shown below.

   Sample YAML File

   ```
   device_to_ztp:
     - serial_number: '74035760356'
     - serial_number: '55674325815'
     - serial_number: '11675330144'
     - serial_number: '74656088411'
   ```

Figure 21: 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 disabled.

4. Enter the root user password and click **Next** to proceed with device discovery.

5. Click **Autoconfigure**.

6. Navigate to Infrastructure > Fabrics > *Fabric Name* > **Fabric Devices** page and verify that the new devices are listed.

7. 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.

Release History Table

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

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 18 on page 90.

Table 18: 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

Table 18: List of Predefined Device Functional Groups *(continued)*

Device Functional Group	Description	OS Version	Routing-Bridging Roles
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 18: 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 .

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:

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

   To create a new device functional group.

   1. Click **Create** on the **Device Functional Groups** tab of the **Infrastructure > Fabrics** page. The **Create Device Functional Group** page appears.

   2. 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.

   3. Click **Create**. The newly created device functional group is listed in the **Device Functional Groups** tab.

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

   To a device functional group.

   1. Click **Create** on the **Fabrics** tab of the **Infrastructure > Fabrics** page. The **Select Provisioning Option** page appears.

   2. Select the **New Fabric** option since device functional groups are supported only on greenfield deployments. The **Create Fabric** page appears.

   3. 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 .

4. Enter the required information as per the steps provided in the "Create a Fabric" on page 36 topic and click **Next.** The **Device discovery** page is displayed.

5. After you have completed the steps provided in the "Discover a Device" on page 46 topic, click **Next**. The **Assign the Roles** page is displayed.

6. 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.

Release History Table

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 | **36**

Assign a Role to a Device | **48**

Return Material Authorization | **134**

Creating Layer 3 PNF Service Chains for Inter-LR Traffic

IN THIS SECTION

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 22: Example Topology
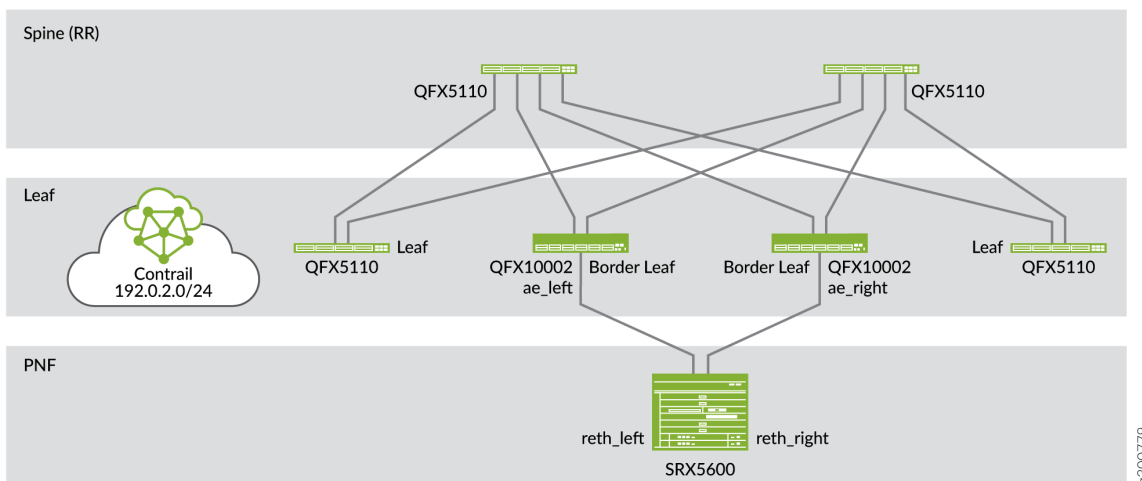
[Figure 22 on page 94](#) shows an example topology of how a PNF device (SRX5600) is used to allow inter-LR traffic to pass through a service chain. You can use the SRX device as a layer 3 PNF device after you have configured the device during device onboarding. The PNF device is connected to border leaf or spine devices.

Getting Started

The general workflow to create a PNF service chain is as follows:
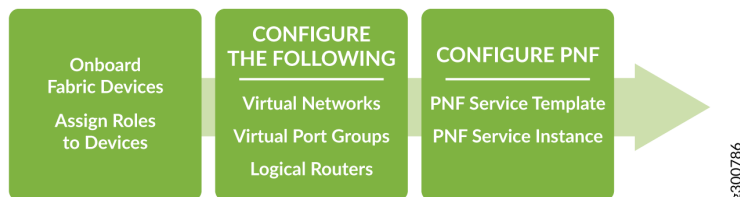
These topics provide instructions to create a PNF service chain.

Onboard Fabric Devices

Follow the steps provided in the "Onboard Devices" on page 66 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 19 on page 95 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 19: 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 70 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 80.

Configure Virtual Port Groups

Follow the steps provided in the "Configuring Virtual Port Groups" on page 113 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 76 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 20: 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 21: Enter PNF Service Instance Information

Field	Action
Name	Enter a name for the PNF service instance.
Service Template	The PNF service template is selected by default.
PNF eBGP ASN	Enter the PNF eBGP AS number.
(Optional) **Configure Static RP**	Select **Configure Static RP** check box to configure static rendezvous point (RP). The RP IP Address field is enabled. The RP is the router that receives multicast traffic. This field is required only when sending multicast traffic through the PNF service chain.
(Optional) **RP IP Address**	Enter the RP IP address iof the router that receives multicast traffic. This field is required only when sending multicast traffic through the PNF service chain.
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	Enter left service VLAN ID. The VLAN ID must be unique.
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 21: Enter PNF Service Instance Information *(continued)*

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

7. Click **Finish** to complete configuration.

The PNF Service Instances page is displayed. For a sample resulting configuration, see .

Figure 23: Resulting Configuration

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:

1. Click **Services** > **Appliances**.

   The Appliances page is displayed.

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

3. 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.

Creating VNF Service Chains for Inter-LR Traffic

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 24: Traffic Flow Between Two Bare Metal Servers
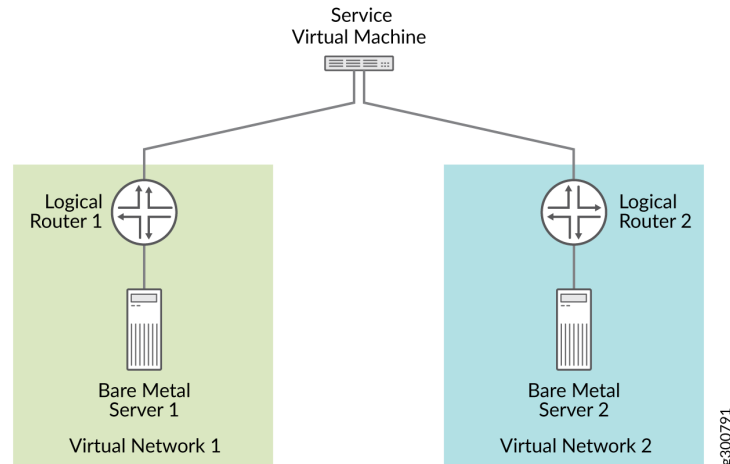
Figure 24 on page 101 shows traffic flowing between two bare metal servers. Each bare metal server is connected to a logical router (virtual routing engine). These logical routers are configured in order to send traffic from the bare metal server in the red virtual network to the bare metal server in the green virtual network, through the service virtual machine.

- A bare metal server and a virtual machine.

Figure 25: Traffic Flow Between a Bare Metal Server and a Virtual Machine
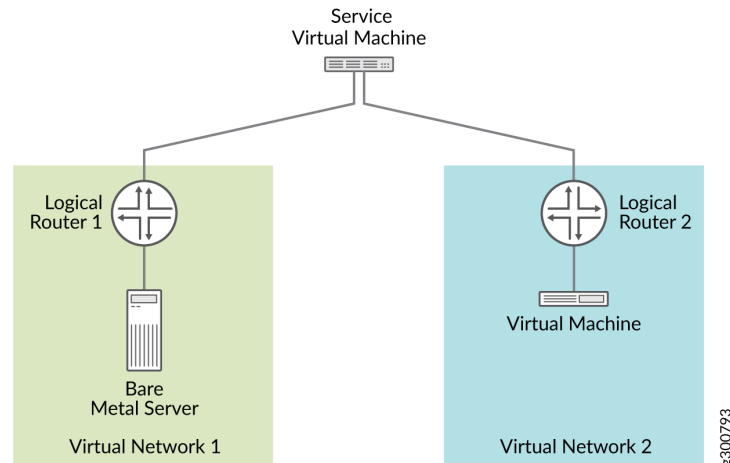
Figure 25 on page 101 shows traffic flowing between a bare metal server and a virtual machine. The bare metal server and the virtual machine are connected to logical routers. These logical routers are configured in order to send traffic from the bare metal server in the red virtual network to the virtual machine in the green virtual network, through the service virtual machine.

- A virtual machine and a bare metal server.

Figure 26: Traffic Flow Between a Virtual Machine and a Bare Metal Servers
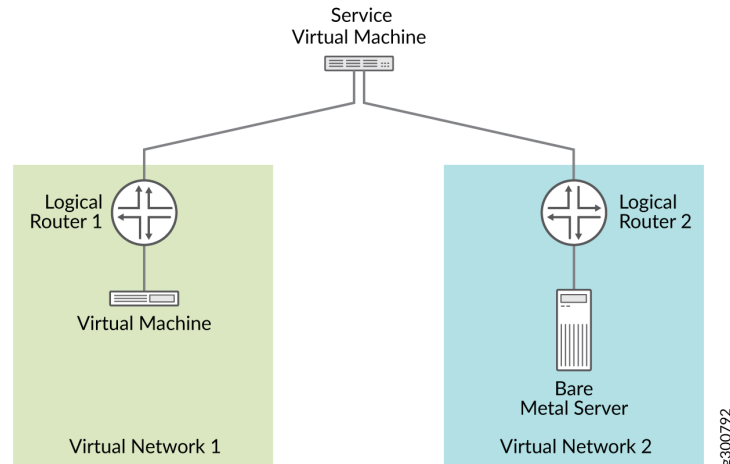
shows traffic flowing between a virtual machine and a bare metal server. The virtual machine and the bare metal server and are connected to logical routers. These logical routers are configured in order to send traffic from the virtual machine in the red virtual network to the bare metal server in the green virtual network, through the service virtual machine.

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

Onboard Devices

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

1. Click **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.

Figure 27: Select Provisioning Option
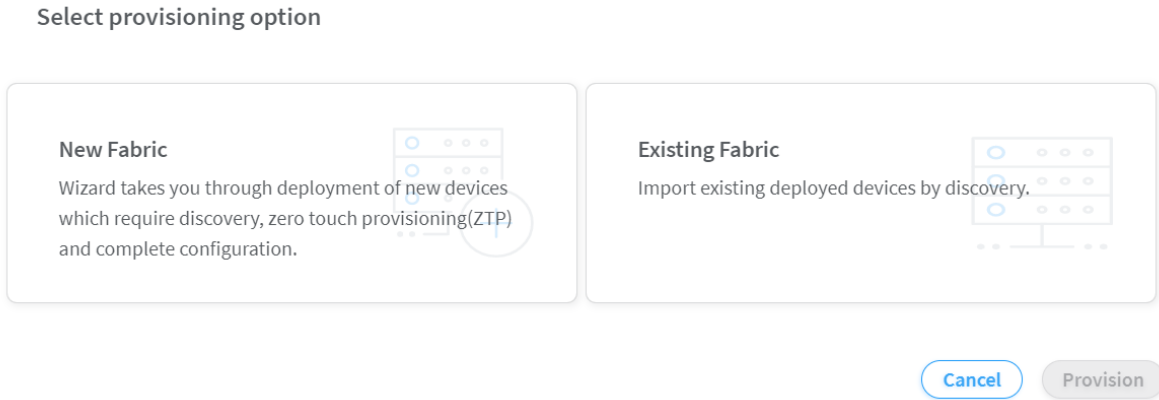
Select provisioning option

New Fabric

Wizard takes you through deployment of new devices which require discovery, zero touch provisioning(ZTP) and complete configuration.

Existing Fabric

Import existing deployed devices by discovery.

Cancel   Provision

4. Click **Provision**.

   The Create Fabric page is displayed.

5. Enter the following information:

   Table 22: Provision Existing Fabric

Field	Action
Name	Enter a name for the fabric.
Username	Enter a username for the device.
Password	Enter a password for the device.
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.
Node profiles	Add node profiles. You can add more than one node profile. All preloaded node profiles are added to the fabric by default. You can remove a node profile by clicking **X** on the node profile.

Table 22: Provision Existing Fabric *(continued)*

Field	Action
Management subnets	Enter the following information:
	CIDR—Enter CIDR network address.
	Gateway—Enter gateway address.
	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.
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.
	• Enter minimum value in **ASN From** field.
	• Enter maximum value in **ASN To** field.
Fabric subnets (CIDR)	Enter fabric CIDR address.
	NOTE: Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.
Loopback subnets (CIDR)	Enter loopback address.
	NOTE: Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.
PNF Servicechain subnets (CIDR)	Enter PNF device CIDR address.
	NOTE: Starting in Contrail Networking Release 5.1, enter the subnet for allocating IP addresses in the PNF Servicechain subnets field to establish EBGP session between PNF device and SPINE switch.

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.

Figure 28: Device Discovery Progress Bar

Device discovery progress

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

7. Select the device(s) 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. Click the **Assign** icon at the end of the row to assign roles.

   The Assign role to devices pop-up is displayed.

10. Assign physical roles and routing bridging roles.

   For Spine Devices:

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

   For Leaf Devices:

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

   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.

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.

11. Click **Assign** to confirm selection and then 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.

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

    The onboarding job is now complete.

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.

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

   The All Networks page is displayed.

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

   The Create Virtual Network page is displayed.

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

4. 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 80.

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

5. 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.

6. Enter subnet information as given in Table 13 on page 71.

Table 23: 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.

7. 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.

8. 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.

9. Enter fat flows information. See Table 14 on page 72.

   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 24: 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.
	Prefix	Enter source subnet prefix length. 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.

Table 24: Configure Fat Flow *(continued)*

Field		Action
Prefix Aggregation Destination	**Destination Subnet**	Enter the destination IP address.
		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.
		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.
	Prefix	Enter the destination subnet prefix length.
		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.

10. 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 .

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

Table 25: 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.

11. Enter advanced configuration information as given in .

Table 26: 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.

Table 26: Advanced Configuration *(continued)*

Field	Action
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.
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 26: Advanced Configuration *(continued)*

Field	Action
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.
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.

12. 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 from Contrail Command UI.

To create virtual port groups:

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

   The Create Virtual Port Group page is displayed.

2. Enter the VLAN ID and network to which the VLAN is associated and select a security group to which the VLAN is to be attached.

   You can select multiple VLANs to include in the virtual port group. Based on the need, you can add or remove VLANs from virtual port group by using the **Edit Virtual Port Group** function.

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 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 , a link aggregation group (LAG) is automatically created on the device.

Figure 29: Select Interfaces on the Same TOR

5. 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 30: Virtual Port Groups

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

|

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. Click **Overlay>Logical Routers**.

   The Logical Routers page is displayed.

2. Click **Create**.

The Create Logical Router page is displayed.

3. Enter the following information.

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.
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.
Logical Router Type	Select **SNAT Routing** or **VXLAN Routing** from the list.
Connected Networks	Select the networks that you want to connect this logical router to.
Public Logical Router	(Optional) Select this check box if you want the logical router to function as a public logical router.
VxLAN Network Identifier	Enter VXLAN network identifier in the range from 1 through 16,777,215. This field is disabled by default.

Field	Action
Route Target(s)	Click **+Add** to add route targets. Enter Autonomous System (AS) number in the **ASN** field. • 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 *L or l (lower-case L)* 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 *L or l (lower-case L)* 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. Enter route target in the **Target** field. • 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.

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

   The Logical Routers page is displayed.

5. Repeat Step 3 and Step 4 to create another logical router.

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. From the Interface section,

   - Select **left** as the interface type from the **Interface Type** list.

   - Click **+ Add**.

     The Interface Type list is added to the table.

     Select **right** as the interface type.

   - Click **+ Add** again.

     Another Interface Type list is added to the table.

     Select **management** as the interface type.

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

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

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

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.

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.

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

6. 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.

Release History Table

Release	Description
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 | 93

Assisted Replication of Broadcast, Unknown Unicast, and Multicast Traffic

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 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 138.

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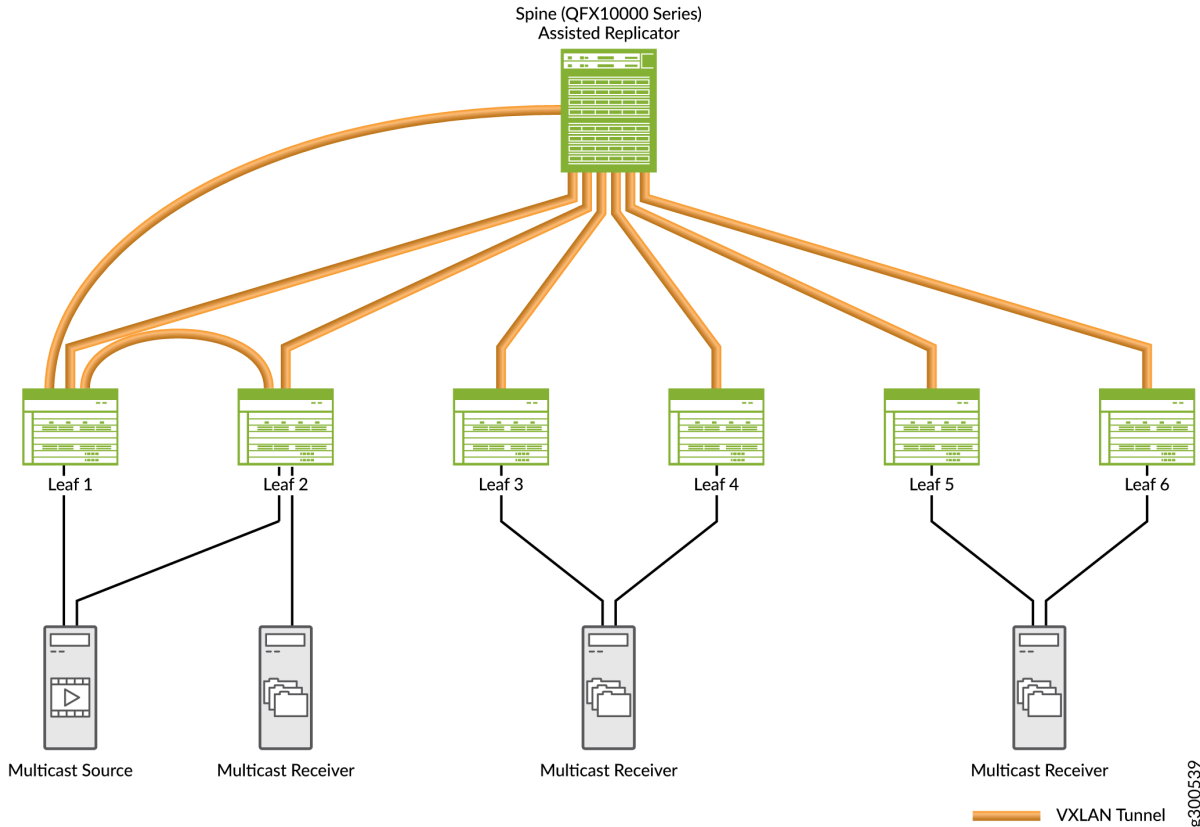
Assisted replication is similar to conventional network segmentation. While the segments in conventional network segmentation could be in different regions, the segments in assisted replication exist in the same region.

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 48 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.

Release History Table

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.

Running Generic Device Operations Commands In Contrail Command

Contrail Networking Release 5.1 and later 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:

- **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 wan to check which among the selected devices has an **xe-0/0/2** interface or an **lo0.0** 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.

- **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.

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 32: Select fabric Devices

3. Click the operation that you want to perform and click **Next**.

Figure 33: Choose an Operation
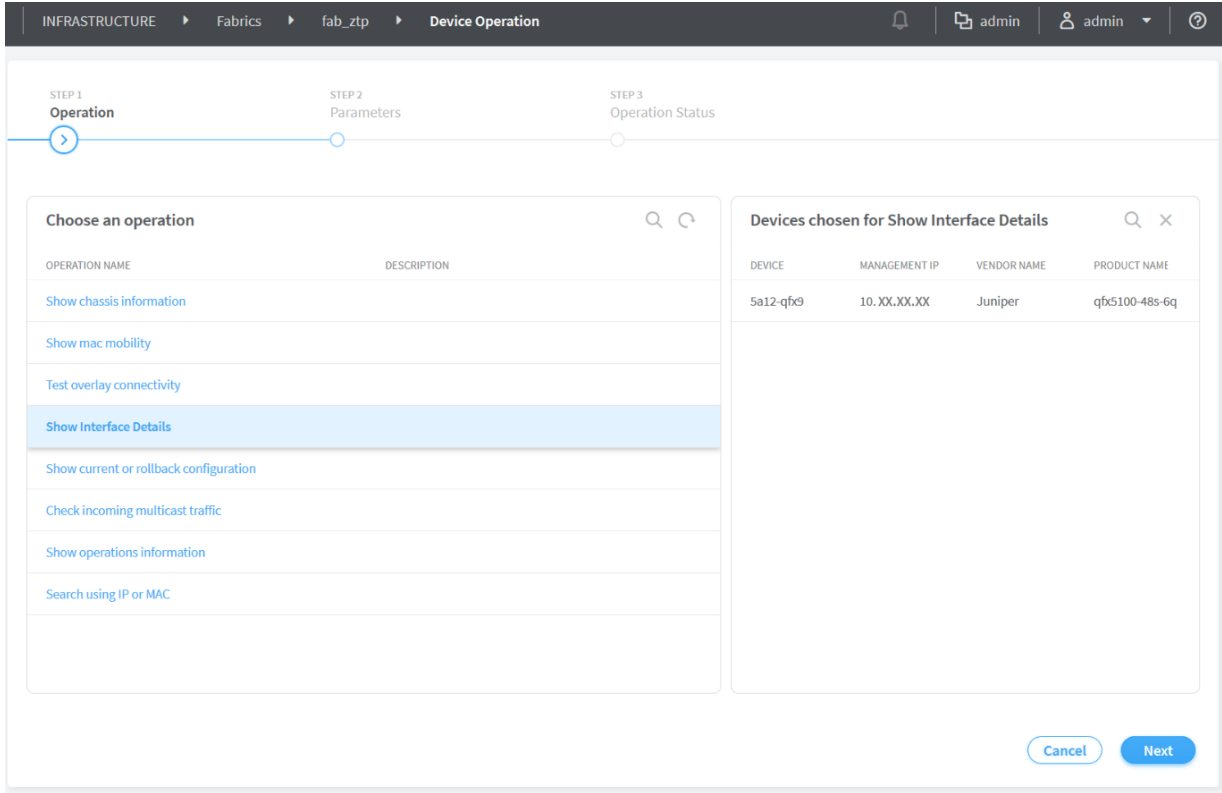
4. Click an operation. Details of the devices selected are displayed as shown in .

Figure 34: Devices Selected for the Operation

5. Click **Next** and select the filters.

Figure 35: Select Filters

6. Click **Execute**.

Information about the selected device is displayed as shown in .

Figure 36: Generic Device Operation Command Results

7. Click **Finish** to complete the operation.

Adding DHCP Server Information for Virtual Networks and Logical Routers

IN THIS SECTION

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 :

- 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 37: Sample Topology
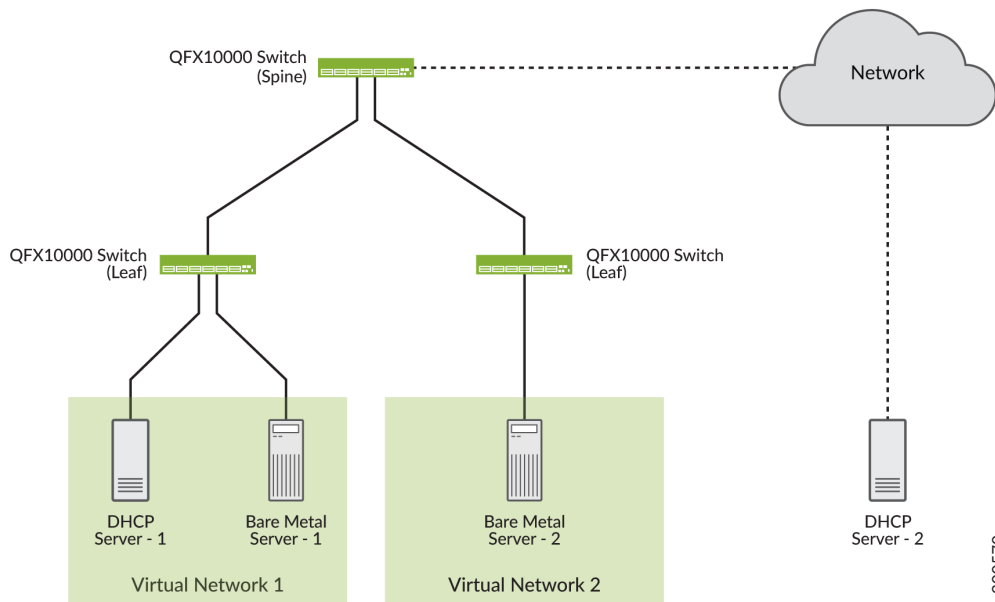
Depending on whether the DHCP server is located in the same virtual network or connected remotely to a routed network through the underlay, the following configuration is pushed to the leaf switch or spine switch.

- 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 | 131
- Adding DHCP Server Information while Creating a Logical Router | 132

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 133.

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 38: DHCP IP Address Field

DHCP Relay Server(s)

DHCP IP Address*

Enter valid IP

+ Add

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 .

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**.
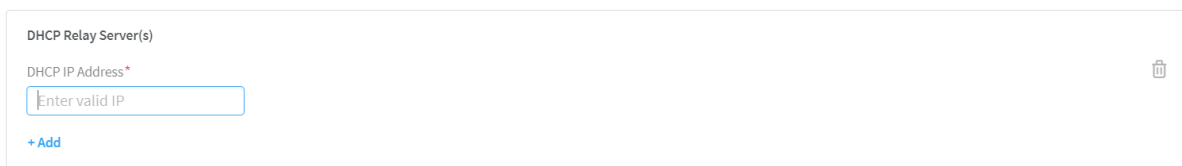
   The Logical Routers page is displayed.

2. Click **Create**.

   The Create Logical Routers page is displayed.

3. Enter the following information as given in .

   Table 27: 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.
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.

4. From the DHCP Relay Server(s) section, click **+Add**.

   The **DHCP IP Address** field is displayed as seen in .

Figure 39: DHCP IP Address Field

5. 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 .

6. (Optional) Click **+Add** to add more DHCP IP addresses.

7. 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.

1. 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.

Release History Table

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

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 to move the selected device to the RMA state.

Figure 40: Return Material Authorization Icon
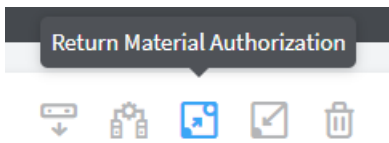
The following confirmation message is displayed:

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 41 on page 136.

Figure 41: Move Out of RMA Icon
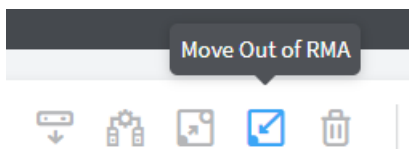
The following confirmation message is displayed:

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 134.

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 36 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.

Release History Table

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.
1907	Contrail Networking Release 1907 supports Return Material Authorization (RMA).

RELATED DOCUMENTATION

Fabric Overview | 17

Create a Fabric | 36

Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles

IN THIS SECTION

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.

Table 28: 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 GW	ERB GW	lean	DC Gateway	DCI Gateway	PNF Service Chaining	Route Reflector	AR Replicator	AR Client
QFX5100 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...	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...	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...	1909 / 19.1R2	1909 / 19.1R2	1912 / 19.1R2			1912 / 19.1R2	1909 / 19.1R2				1909 / 19.1R2		R1909 / 19.1R2
QFX...	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

Table 28: 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-GV	ERB-SV	lean	DC Gateway	DCI Gateway	PNF Service Chaining	Route Reflector	AR_Replicator	AR_Client
QFX5xxC	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
QFX10xxQ	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
QFX02Q QFX003 QFX06	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	R1907 / 18.4R2	R1907 / 18.4R2
QFX0xC	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 29: Supported MX Series Routers

MX Device	Supported from Contrail Networking Release / Supported from Junos OS Releases						
	Physical Roles		Overlay Roles				Gateway Ro...
	Leaf	Spine	ERB-UCAST-Gateway	CRB-Gateway	CRB-MCAST-Gateway	null	DC-Gatewa...
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

Table 29: Supported MX Series Routers *(continued)*

MX Device	**Supported from Contrail Networking Release** **Supported from Junos OS Releases**						
	Physical Roles		**Overlay Roles**				**Gateway R...**
	Leaf	**Spine**	**ERB-UCAST-Gateway**	**CRB-Gateway**	**CRB-MCAST-Gateway**	**null**	**DC-Gatewa...**
MX2008, MX2010, MX2020	5.0.2 17.3R3	5.0.2 17.3R3	R2003 18.4R2-S3	R2003 18.4R2-S3	R2003 18.4R2-S3	5.1 18.1R3	2003 18.4R2-S3 (without SNAT)
MX10003		5.1 18.1R3		2003 18.4R2-S3	2003 18.4R2-S3		
MX204, MX10008, MX10016	5.1 18.1R3	5.1 18.1R3	R2003 18.4R2-S3	R2003 18.4R2-S3	R2003 18.4R2-S3	5.1 18.1R3	2003 18.4R2-S3 (without SNAT)

Table 30: 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 .
- For a list of MX devices, see .
- For a list of SRX devices, see .

Table 31: 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
		Spine	lean, CRB-Access, CRB-Gateway, DC-Gateway, Route-Reflector, DCI-Gateway, CRB-MCAST-Gateway, PNF-Servicechain, AR-Client, AR-Replicator

Table 31: Supported QFX Series Switches *(continued)*

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
QFX10002–72Q, QFX10016, QFX10008	junos-qfx10k	Leaf	CRB-Access, CRB-Gateway, DC-Gateway, Route-Reflector, ERB-UCAST-Gateway, CRB-MCAST-Gateway, PNF-Servicechain, AR-Client, AR-Replicator
		Spine	lean, CRB-Access, CRB-Gateway, DC-Gateway, Route-Reflector, CRB-MCAST-Gateway, PNF-Servicechain, AR-Client, AR-Replicator

Table 31: Supported QFX Series Switches *(continued)*

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
QFX10002-60C	junos-qfx	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. |
| | | Spine | |

Table 31: Supported QFX Series Switches *(continued)*

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
			lean, CRB-Gateway, DC-Gateway, Route-Reflector, ERB-UCAST-Gateway, DCI-Gateway, PNF-Servicechain 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

Table 31: Supported QFX Series Switches *(continued)*

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
			the QFX10002-60C device. For more information, see "Create a Fabric" on page 36.
QFX5100-XX models QFX5200-32C-32Q, QFX5210-64C	junos-qfx	Leaf	CRB-Access, Route-Reflector, AR-Client
		Spine	lean, Route-Reflector, AR-Client
QFX5110-48S-4C, QFX5110-32Q	junos-qfx	Leaf	CRB-Access, Route-Reflector, ERB-UCAST-Gateway, AR-Client
		Spine	lean, CRB-Access, CRB-Gateway, Route-Reflector, AR-Client
QFX5120-48Y-8C	junos-qfx	Leaf	CRB-Access, CRB-Gateway, Route-Reflector, ERB-UCAST-Gateway, AR-Client
		Spine	lean, Route-Reflector, AR-Client

Table 31: Supported QFX Series Switches *(continued)*

QFX Device	Node Profile	Physical Roles	Routing Bridging Roles
QFX5120-32C	juniper-qfx5120	Leaf	CRB-Access, Route-Reflector, AR-Client, ERB-UCAST-Gateway NOTE: Contrail Networking Release 1909 supports QFX5120-32C device running Junos OS Release 19.1R2 and later.
		Spine	lean, Route-Reflector, AR-Client NOTE: Starting in Contrail Networking Release 1910, a QFX5120-32C device can be used in lean-spine deployments.

Table 32: Supported MX Series Routers

MX Device	Node Profile	Physical Roles	Routing Bridging Roles
MX80	**juniper-mx**	Leaf	DC-Gateway, Route-Reflector
		Spine	DC-Gateway, lean, Route-Reflector

Table 32: Supported MX Series Routers *(continued)*

MX Device	Node Profile	Physical Roles	Routing Bridging Roles
MX240, MX480	**juniper-mx**	Leaf	DC-Gateway, Route-Reflector, DCI-Gateway, ERB-UCAST-Gateway
		Spine	DC-Gateway, lean, 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, lean, Route-Reflector, CRB-Gateway, CRB-MCAST-Gateway
MX10003	**juniper-mx**	Spine	Route-Reflector, CRB-Gateway, CRB-MCAST-Gateway

Table 33: 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

Release History Table

Release	Description
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.

5
CHAPTER

Managing Data Center Devices

Data Center Interconnect

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 interconnect of data centers that exist in different fabrics. 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. For more information, see "Creating Data Center Interconnect" on page 152.

> **NOTE:** The gateway devices must support DCI-Gateway routing bridging role.
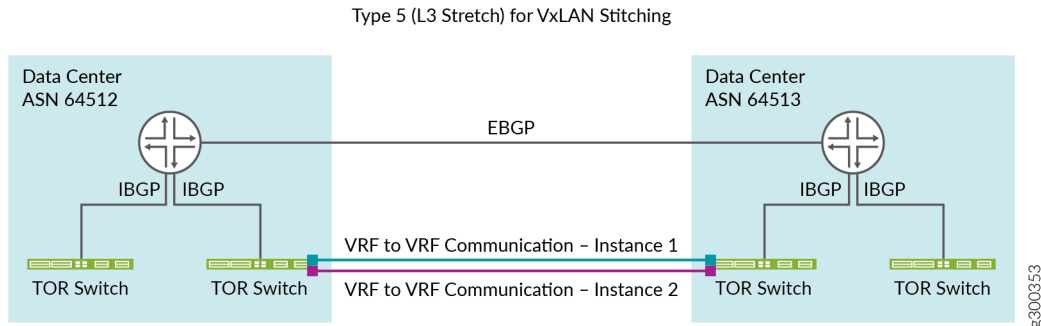
Data Center Interconnect Deployment Topologies

IN THIS SECTION

Contrail Networking supports the following data center interconnect (DCI) deployment topologies.
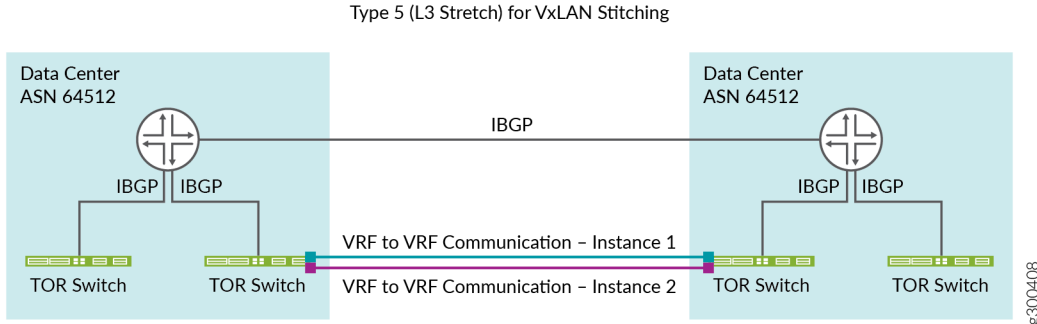
DCI using EBGP

Figure 42: 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 .

DCI using IBGP

Figure 43: DCI using IBGP Connection

Type 5 (L3 Stretch) for VxLAN Stitching

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 .

Creating Data Center Interconnect

IN THIS SECTION

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.

Create a Fabric

Follow these steps to create a fabric with brownfield devices from the Contrail Command user interface (UI):
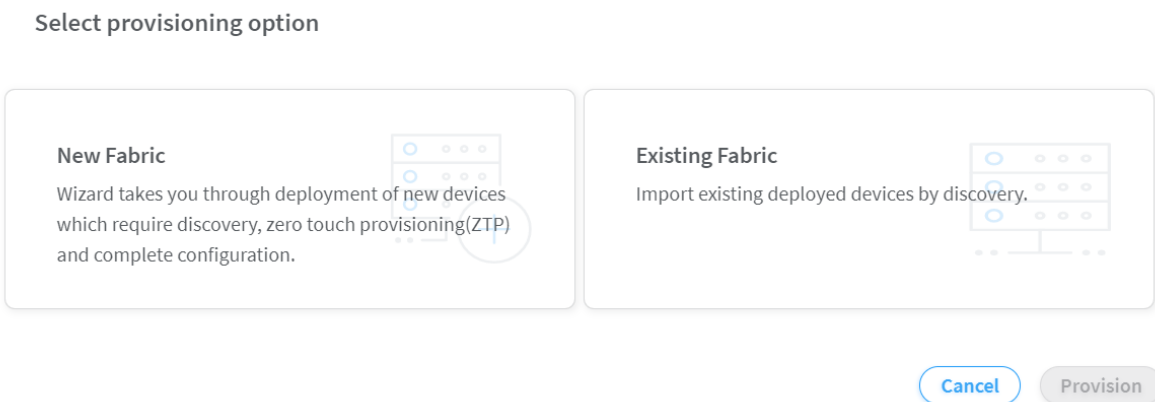
1. Click **Fabrics**.

   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.

Figure 44: Select Provisioning Option

Select provisioning option

New Fabric

Wizard takes you through deployment of new devices which require discovery, zero touch provisioning(ZTP) and complete configuration.

Existing Fabric

Import existing deployed devices by discovery.

Cancel   Provision

4. Click **Provision**.

   The Create Fabric page is displayed.

Figure 45: Create Fabric Page
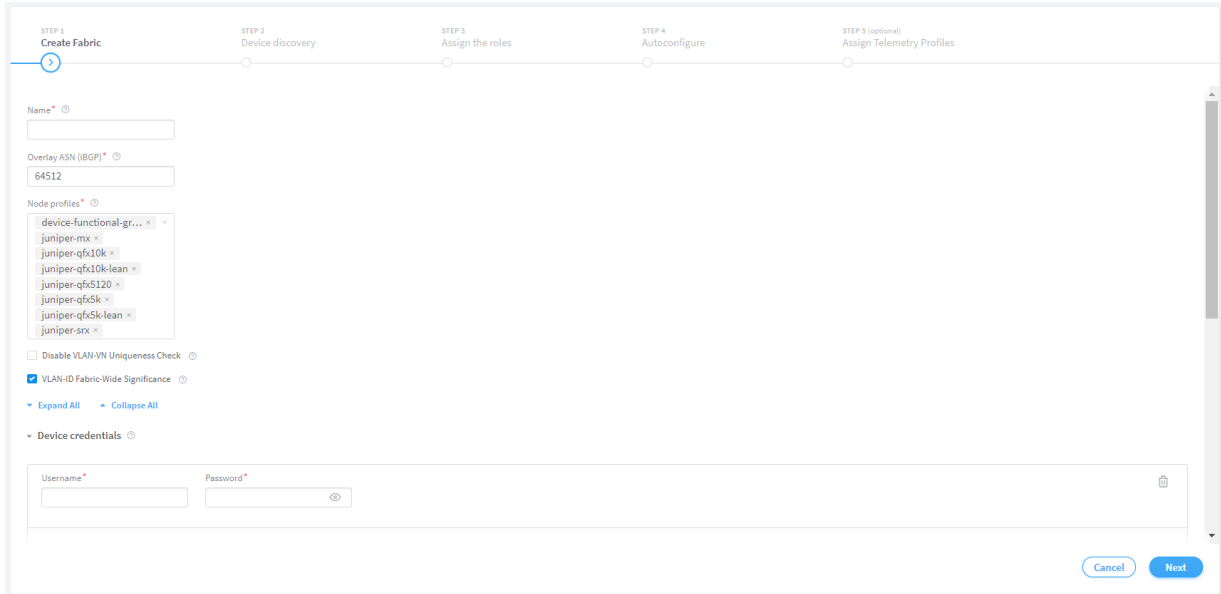
5. Enter the following information:

    Table 34: Provision Existing Fabric

Field	Action
Name	Enter a name for the fabric.
Username	Enter a username for the device.
Password	Enter a password for the device.
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.
Node profiles	Add node profiles. You can add more than one node profile. All preloaded node profiles are added to the fabric by default. You can remove a node profile by clicking **X** on the node profile.

Table 34: Provision Existing Fabric *(continued)*

| Field	Action
Management subnets	Enter the following information: **CIDR**—Enter CIDR network address. **Gateway**—Enter gateway address. 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.
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. • Enter minimum value in **ASN From** field. • Enter maximum value in **ASN To** field.
Fabric subnets (CIDR)	Enter fabric CIDR address. NOTE: Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.
Loopback subnets (CIDR)	Enter loopback address. NOTE: Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.

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.

Figure 46: Device Discovery Progress Bar
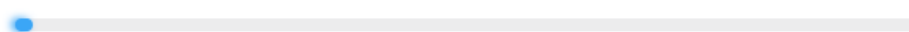
The list of devices discovered are listed in the Discovered devices page.

7. Select the device(s) 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. Click the **Assign** icon at the end of the row to assign roles.

   The Assign role to devices pop-up is displayed.

10. Assign physical roles and routing bridging roles.

    For Spine Devices:

    • Select **spine** from the Physical Role list.

    • Select **DCI-Gateway** from the Routing Bridging Roles list.

Figure 47: Assign Role to Spine Devices

Assign role to 1 devices

Physical Role

spine

Routing Bridging Roles

DCI-Gateway ×

Cancel   Assign

    For Leaf Devices:

    • Select **leaf** from the Physical Role list.

    • Select **DCI-Gateway** from the Routing Bridging Roles list.

Figure 48: Assign Role to Leaf Devices

11. Click **Assign** to confirm selection and then click **Autoconfigure** to initiate the auto-configuration job.

The Autoconfigure page is displayed.

Create Virtual Network

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

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

The All Networks page is displayed.

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

The Create Virtual Network page is displayed.

Figure 49: Create Virtual Network Page

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

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

5. 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.

6. The VXLAN ID is populated by default and is displayed in the **VxLAN Network Identifier** field.

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

8. Enter valid IPv4 address in the **Gateway** field.

9. Click **Create**.

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

Create Logical Routers

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

1. Click **Overlay**>**Logical Routers**.

   The Logical Routers page is displayed.

2. Click **Create**.

   The Create Logical Router page is displayed.

3. Enter the following information.

Field	Action
Name	Enter a name for the Logical Router.
Admin State	Select **Up**.
Extend to Physical Router	Select the routers from the list.
Logical Router Type	Select VXLAN Routing from the list.
Connected Networks	Select the networks from the list.
Public Logical Router	(Optional) Select this check box if you want the logical router to function as a public logical router.
VxLAN Network Identifier	Enter VXLAN network identifier. Range: 1 through 16,777,215

Field	Action
Route Target(s)	Click **+Add** to add route targets. Enter Autonomous System (AS) number in the **ASN** field. • 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 *L* or *l (lower-case L)* 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 *L* or *l (lower-case L)* 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. Enter route target in the **Target** field. • 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.

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

   The Logical Routers page is displayed.

5. Repeat Step 3 and Step 4 to create another logical router.

Create DCI

Creating Data Center Interconnect

Follow these steps to create a DCI of two different data centers from the Contrail Command user interface (UI).

1. Click **Overlay** > **DCI**.

   The DCI page is displayed.

2. Click **Create**.

   The Create DCI page is displayed.

Figure 50: Create DCI Page

3. Enter the following information.

Field	Action
DCI name	Enter a name for the DCI.
BGP Hold Time	Modify BGP hold time. *This field is optional.*
BGP Address Family	Modify the existing BGP address family by selecting BGP address family from the **BGP Address Family** list. You can select more than one option from the list. *This field is optional.*

Field	Action
Connections	Follow these steps to connect two logical routers. a. Select logical router from the **Select logical router** list. b. Select fabric from the **Select fabric** list. c. Select the physical router you want to extend the connection to from the **Extend to Physical Router** list. Repeat the above steps to create the next connection.

4. Click **Create**.

   The connections you create are listed in the DCI page.

RELATED DOCUMENTATION

  VXLAN Data Center Interconnect Using EVPN Overview

Configuring QFX Series Devices as Data Center Gateway

IN THIS SECTION

- Discover a Fabric | **163**
- Add Bare Metal Server | **167**
- Create Tenant Virtual Network | **168**
- Add CSN Nodes | **169**
- Create Logical Router | **170**
- Verification | **171**

You can configure a QFX series device as a Data Center Gateway (DC-GW). DC-GW is an overlay role that is assigned to a QFX series switch 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 devices, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 138.

These sections provide instructions to configure a QFX10000 switch as a DC-GW.

Discover a Fabric

Follow these steps to discover a fabric by using the Contrail Command user interface (UI):

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

2. Click **Create**.

   You are prompted to select a provisioning option.

3. Click **Existing Fabric** to import existing (brownfield) devices by discovery.

Figure 51: Select Provisioning Option
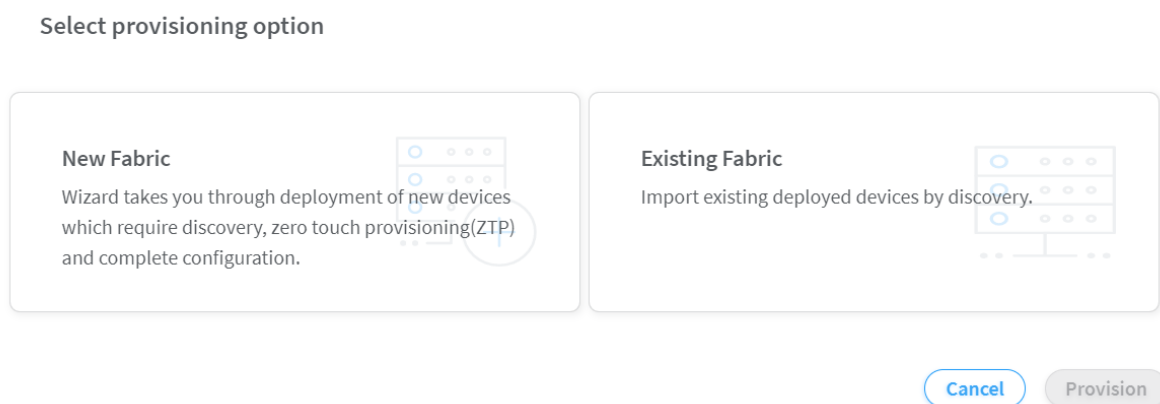
Select provisioning option

New Fabric
Wizard takes you through deployment of new devices which require discovery, zero touch provisioning(ZTP) and complete configuration.

Existing Fabric
Import existing deployed devices by discovery.

Cancel   Provision

4. Click **Provision**.

   The Create Fabric page is displayed.

5. Enter the fabric provisioning information as listed in Table 35 on page 164.

Table 35: Provision Existing Fabric

Field	Action
Name	Enter a name for the fabric.
Username	Enter a username for the device.
Password	Enter a password for the device.
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.
Node profiles	Add node profiles. You can add more than one node profile. All preloaded node profiles are added to the fabric by default. You can remove a node profile by clicking **X** on the node profile.
Management subnets	Enter the following information: **CIDR**—Enter CIDR network address. **Gateway**—Enter gateway address. **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.
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. • Enter minimum value in **ASN From** field. • Enter maximum value in **ASN To** field.
Fabric subnets (CIDR)	Enter fabric CIDR address. **NOTE:** Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.

Table 35: Provision Existing Fabric *(continued)*

Field	Action
Loopback subnets (CIDR)	Enter loopback address. NOTE: Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.
VLAN-ID Fabric Wide Significance	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. You can modify the enterprise style setting to service provider style once configured. However, you cannot modify the service provider style to enterprise style. 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 182. For more information on supported hardware platforms and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 138.

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 page.

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. Click the **Assign** icon at the end of the row to assign roles.

   The Assign role to devices pop-up is displayed.

10. 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 138.

11. Click **Assign** to confirm selection and then click **Autoconfigure** to initiate the auto-configuration job.

   The Autoconfigure page is displayed.

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 Existing Baremetal Server pane:

   Table 36: Add Existing Bare Metal Server Information

Field	Action
Instance Name	Displays the name of the BMS instance.
Baremetal Node	Select a bare metal node.
Interface	Select an interface from the list.
IP Address	Enter IP address of the instance.
VLAN ID	Enter VLAN ID.
Virtual Network	Select a virtual network from the list.
Select Security Groups	Select **default** security group from the list.

Figure 52: Existing Bare Metal Server

5. Click **Create** to confirm.

Create Tenant Virtual 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 by using the Contrail Command UI:

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

    The All networks pane is displayed.

2. Click **Create**.

    The Network page is displayed.

3. Enter the following information:

Table 37: Add Tenant Virtual Network Information

Field	Action
Name	Enter a name for the virtual network.
Network Policies	Select network policy from the list.
Allocation Mode	Select **User defined subnet only** as the allocation mode.
VxLAN Network Identifier	Enter VXLAN Network Identifier in the range from 1 through 16,777,215.
Subnets	Click **+Add** to add subnets.
Network IPAM	Select the default Network IPAM.
CIDR	Enter CIDR address.
Allocation Pools	Enter allocation pool information.
Gateway	Enter gateway IP address.
Auto Gateway	Select Auto Gateway check box.
DHCP	Select DHCP check box.

4. Click **Create** to confirm.

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 38: 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 Router

Follow these steps to create a logical router:

1. Click **Overlay**>**Logical Routers**.

   The Logical Routers page is displayed.

2. Click **Create**.

   The Create Logical Router page is displayed.

3. Enter the following information:

   Table 39: Add Logical Router Information

Field	Action
Name	Enter a name for the logical router.
Admin State	Select **Up** as the state.
Extend to Physical Router	Select the device you want to extend the physical router information to from the list.

Table 39: Add Logical Router Information *(continued)*

Field	Action
Logical Router Type	Select VXLAN Routing from the list.
Connected networks	Select connected network information from the list.
NAT	Enable NAT by selecting the NAT check box.
VxLAN Network Identifier	Edit VXLAN network identifier.

4. Click **Create**.

    The Logical Routers pages is displayed.

Verification

EVPN type 5 configuration is pushed to QFX10000 switch as a DC-GW.

Figure 53: EVPN Type 5 Configuration

Release History Table

Release	Description
1909	Contrail Networking Release 1909 supports QFX10002-60C device running Junos OS Release 19.1R2 and later.

Understanding Underlay Management | 15

Fabric Overview | 17

Edge-Routed Bridging for QFX Series Switches | 183

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 54 on page 173 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 54: Virtual Port Group
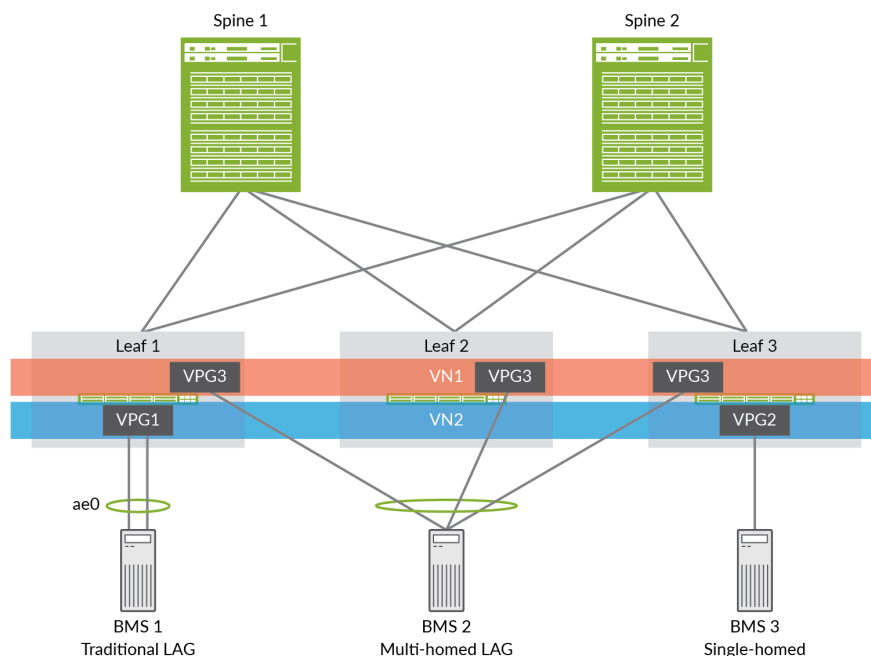
Depending on whether **VLAN-ID Fabric-Wide Significance** field is selected or not, the behavior of virtual port groups is different in enterprise style (**VLAN-ID Fabric-Wide Significance** option enabled) and service provider style (**VLAN-ID Fabric-Wide Significance** option disabled) configurations.

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.

Release History Table

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

This topic describes how to create virtual port groups from Contrail Command UI.

To create virtual port groups:

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

   The Create Virtual Port Group page is displayed.

2. Enter the VLAN ID and network to which the VLAN is associated and select a security group to which the VLAN is to be attached.

   You can select multiple VLANs to include in the virtual port group. Based on the need, you can add or remove VLANs from virtual port group by using the **Edit Virtual Port Group** function.

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 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 , a link aggregation group (LAG) is automatically created on the device.

Figure 55: Select Interfaces on the Same TOR

5. 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 56: Virtual Port Groups

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.

Configuring Storm Control on Interfaces

IN THIS SECTION

●

Starting with Contrail Networking Release 1908, you can configure storm control on the access interfaces of a datacenter fabric managed by Contrail Networking.

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.

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.

Storm control feature is supported in both greenfield and brownfield deployments with enterprise style configuration. 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.

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.

> **NOTE:** Storm control profile feature is supported only on QFX5000 and QFX10000 series devices.

Configuring Storm Control Profiles

To create storm control profiles:

1. Click **Overlay** > **Port Profiles**.

Figure 57: Port Profiles
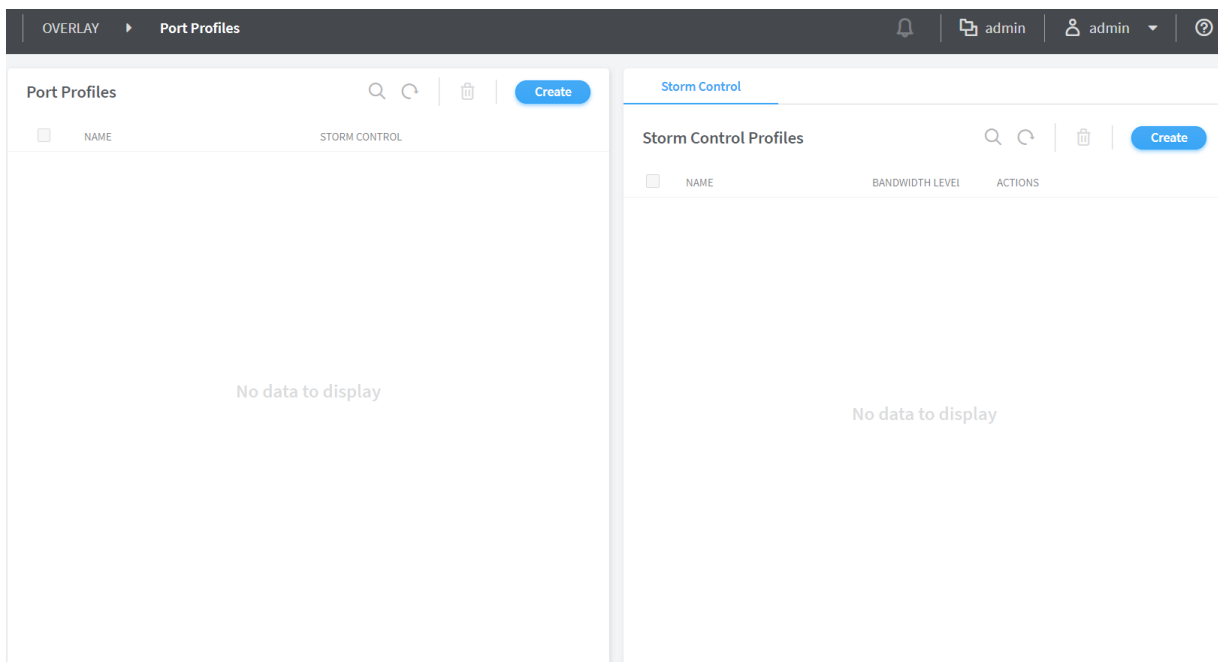
You must first create a storm control profile and then attach it to the port profile. You can attach the storm control profile to existing port profile or attach to a new port profile while creating it.

2. Click **Overlay** > **Port Profiles** > **Create Storm Control Profile**.

Figure 58: 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.
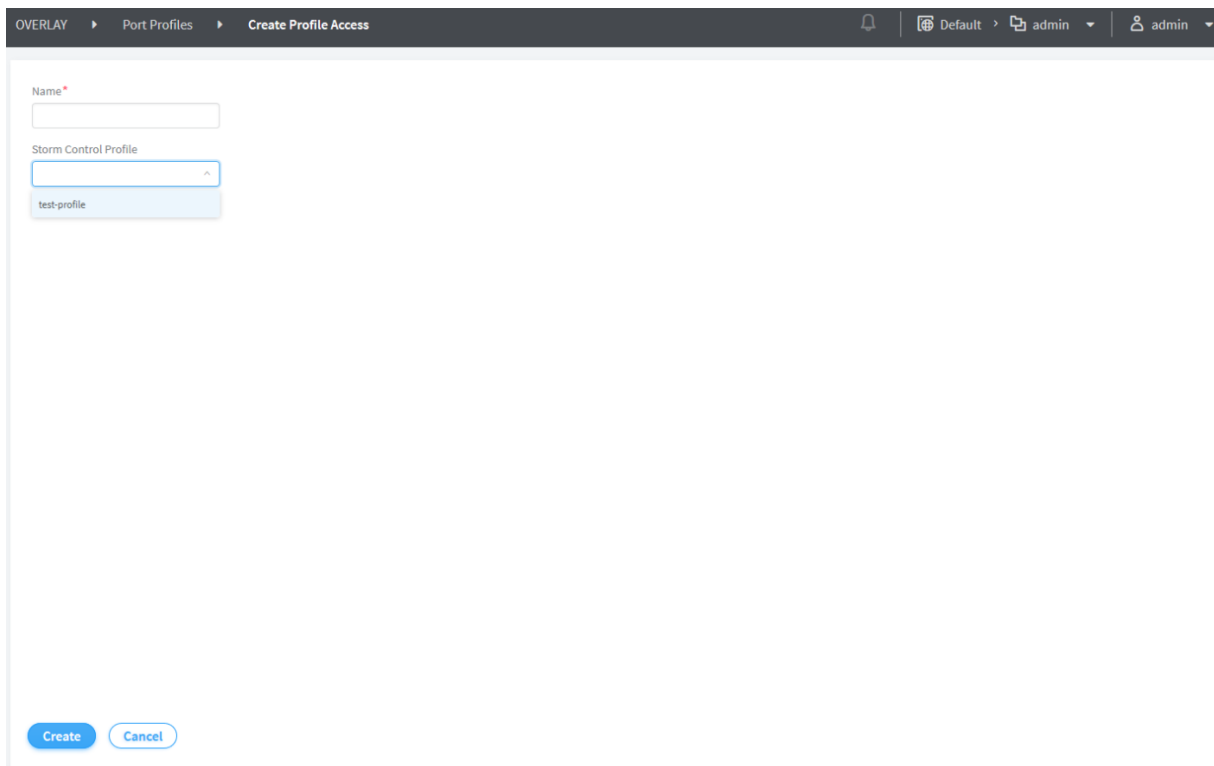
- Bandwidth Level— Enter the maximum value (in percentage) in the range 0–100. If the bandwidth utilized by broadcast, unknown unicast, or multicast (BUM) traffic exceeds this value, the action (default **drop** or configured **Interface shutdown**) specified in the storm control profile is applied on 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 Profile Access**.

Figure 59: Create Port Profile

You must specify a port profile name and select a storm control profile from the profiles created in Step 3. 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.

5. Click **Create**.

After you create a port profile, you can assign it to interfaces or virtual port groups as shown in Figure 60 on page 181.

Figure 60: Attach Port Profile to VPG
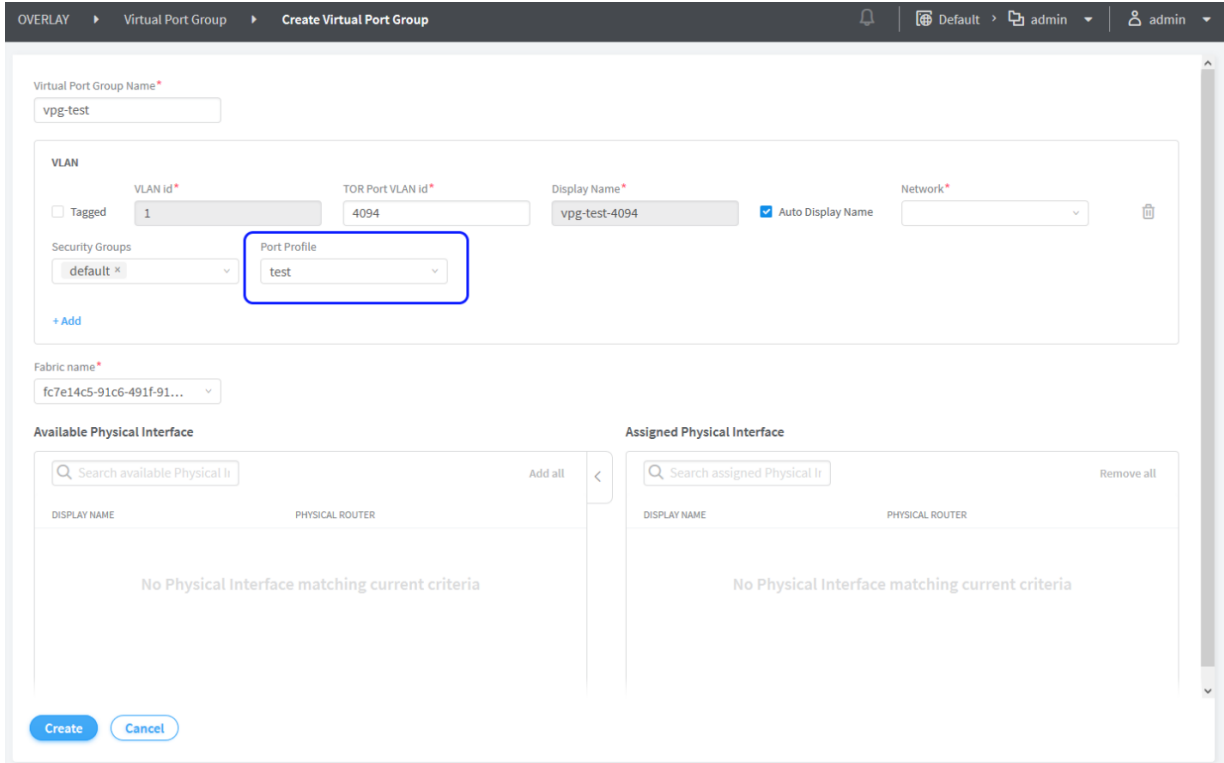
Release History Table

Release	Description
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 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 36.
>
> CRB-Access, CRB-Gateway, AR-Client, DC-Gateway, Route-Reflector, ERB-UCAST-Gateway, DCI-Gateway, and PNF-Servicechain routing bridging roles are supported on QFX10002-60C device. However, CRB-MCAST-Gateway, AR-Client, and AR-Replicator roles are 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 138.

Release History Table

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 | 36

Flexible Ethernet Services Encapsulation

Edge-Routed Bridging for QFX Series Switches

Starting with Contrail Networking Release 5.1, 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 in Contrail Networking Release 5.1. 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 is 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 from Contrail Networking Release 5.1. 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.

RELATED DOCUMENTATION

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**.

Release History Table

Release	Description
1909	Starting with Contrail Networking Release 1909, you can activate maintenance mode on spine and leaf devices in a data center fabric.

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:** You must be running AppFormix version 2.19.11.

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 BMS alone.

Three views are supported:

- Radial
- Vertical
- Horizontal

Select **Infrastructure > Fabrics > *<fabric name>* > Topology View**. Following are two example radial views.
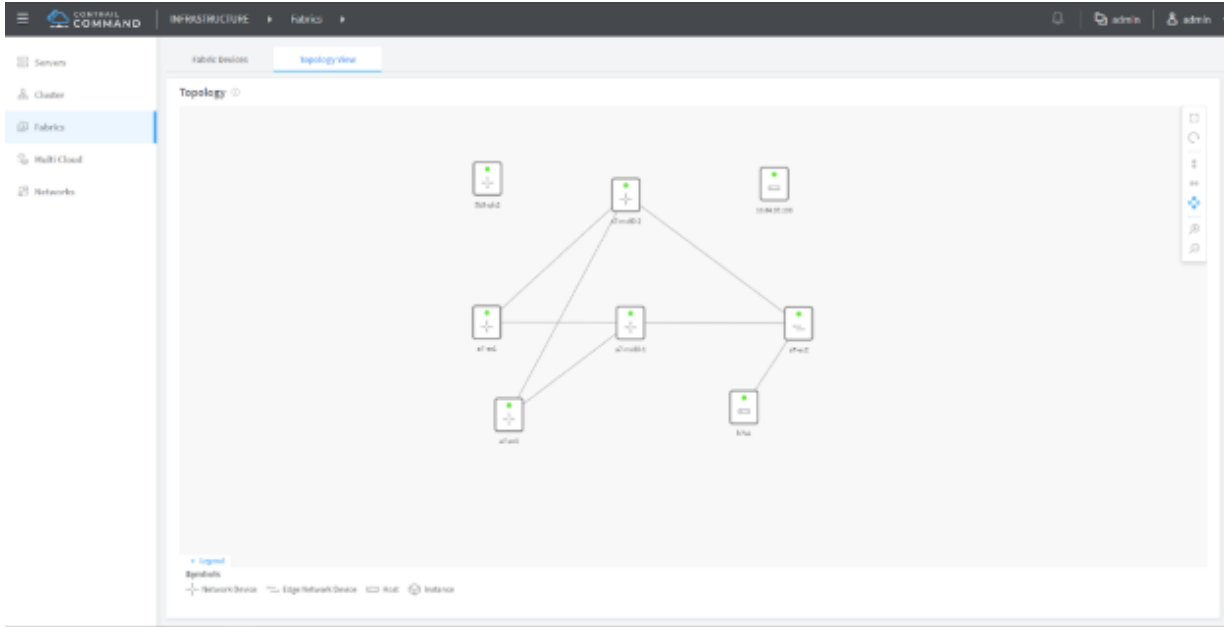
Figure 61: Radial View

Figure 62: Radial View with VMs
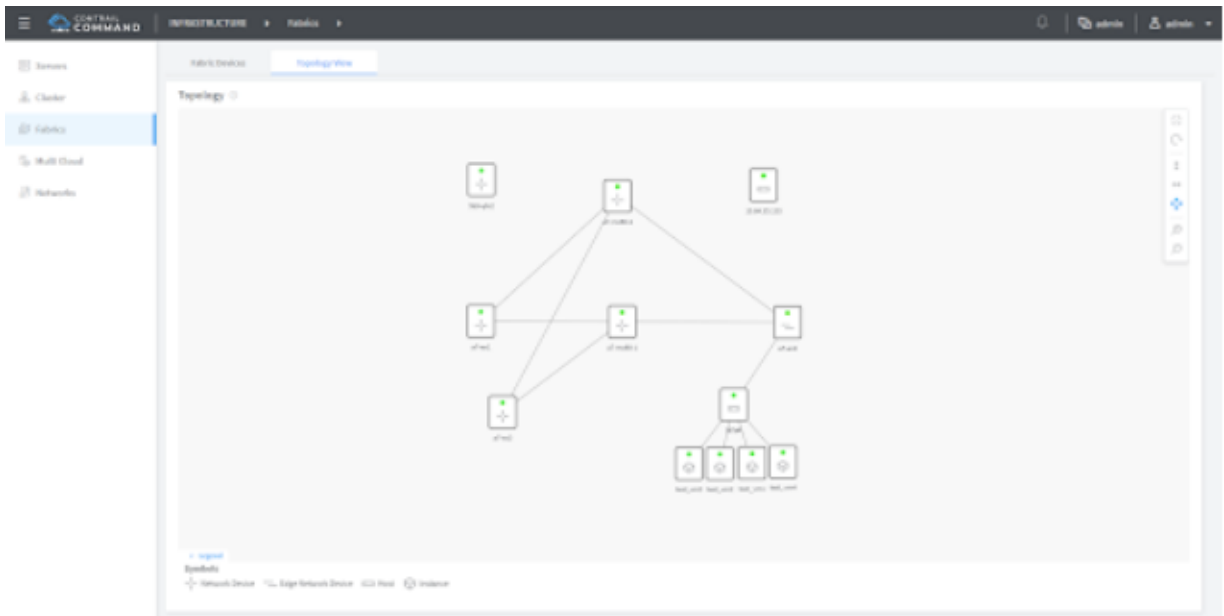
Release History Table

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.

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 63: Hardware Inventory

4. (Optional) Click the **Fetch** button if there is no inventory information available or you want to see an updated inventory information. See .

For more information about the fields displayed on the **Hardware Inventory** page, see show chassis hardware (View).

Release History Table

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).

Certificate Lifecycle Management Using Red Hat Identity Management

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_suffux: '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:KqTDfKQEoKKi7FMzuhBwcO+Y/O9t4rHXQcqPKglJPmI
(ecdsa-sha2-nistp256), SHA256:QSIBCIiRW03eR6+PPyvDWiWEHXC1MewREAt8hMTUOgU
(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
```

6
CHAPTER

Integrating VMware with Contrail Networking Fabric

Understanding VMware-Contrail Networking Fabric Integration

Contrail Networking Release 1910 supports integrating VMware with Contrail Networking fabric. A dedicated Contrail vCenter Fabric Manager (CVFM) plugin is deployed for this integration. This plugin connects various ESXi hosts and helps manage VMware underlay networks. The CVFM plugin is installed when you install the Contrail Command user interface (UI). After the plugin is installed, the plugin runs as a service in a container on the control node. You can enable this plugin when you provision the Contrail Command UI. However, if you do not enable this plugin during provisioning, you can enable the plugin from the **Infrastructure**>**Cluster** 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 ContrailVMs 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 plugin helps synchronize the configuration of VMware Distributed Port Groups (DPG) with the configuration on TOR (leaf) switches. After the CVFM plugin is deployed, Contrail Networking will act 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 Plugin

The following are the benefits of CVFM plugin:

- 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

depicts the CVFM plugin installed on the Contrail Networking control node. The CVFM plugin 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 64: Integrating VMware with Contrail Networking Fabric
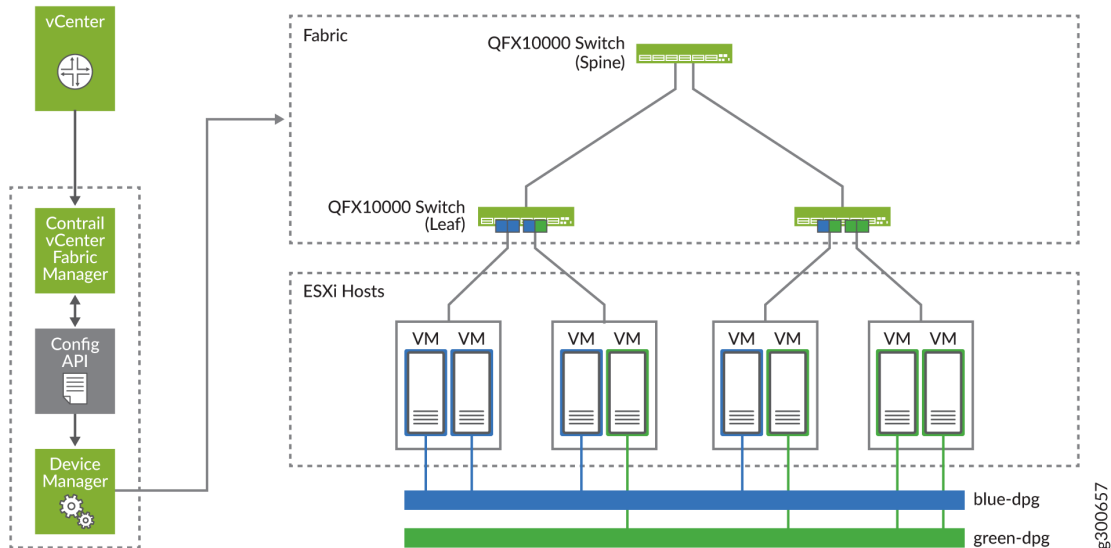
The leaf and spine switches (QFX series) are connected to virtual machines in the ESXi host environment. VLANs are configured on the DPG of these QFX series switches. The CVFM plugin automatically adds

and removes configurations of the VLANs. For more information on deploying the CVFM plugin, see "Deploying Contrail vCenter Fabric Manager Plugin" on page 200.

Getting Started with CVFM Plugin

The CVFM plugin is installed when you install the Contrail Command UI.

1. You can then enable the CVFM plugin while provisioning Contrail Command.

   For more information, see the **Deploying CVFM Plugin while Provisioning Contrail Command** section of the "Deploying Contrail vCenter Fabric Manager Plugin" on page 200 topic.

2. You can also enable the plugin after provisioning Contrail Command.

   For more information, see the **Deploying CVFM Plugin after Provisioning Contrail Command** section of the "Deploying Contrail vCenter Fabric Manager Plugin" on page 200 topic.

3. After you have enabled the plugin, you can update vCenter credentials or override configuration information from Contrail Command.

   For more information, see "Updating vCenter Credentials on Contrail Command" on page 215.

4. After you have enabled the plugin, 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 207.

5. You can also add DPG. For more information, see "Adding Distributed Port Groups" on page 213.

Limitations of the CVFM Plugin

The following are the limitations of the CVFM plugin.

- 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.

Release History Table

Release	Description
1910	Contrail Networking Release 1910 supports integrating VMware with Contrail Networking fabric. A dedicated Contrail vCenter Fabric Manager (CVFM) plugin is deployed for this integration.

RELATED DOCUMENTATION

Deploying Contrail vCenter Fabric Manager Plugin

IN THIS SECTION

Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plugin. With this release, the CVFM plugin is installed when you install the Contrail Command user interface (UI). You can then enable this plugin when you provision Contrail Command. However, if you have not enabled this plugin during provisioning, you can enable the plugin from the **Infrastructure**>**Cluster** page of the Contrail Command UI.

For more information on CVFM plugin, see "Understanding VMware-Contrail Networking Fabric Integration" on page 197.

These topics provide instructions on how to deploy the CVFM plugin.

Prerequisites

Before you deploy the CVFM plugin, 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 Plugin while Provisioning Contrail Command

You enable the CVFM plugin while provisioning Contrail Command.

Follow these steps to enable CVFM plugin:

1. Log in to Contrail Command using the root user credentials.

   When you log in to Contrail Command for the first time, you are directed to the Contrail Command SETUP screen.

2. Click **Credentials**.

   The Available Credentials page is displayed.

3. Click **Add**.

   The Add Credentials page is displayed.

4. Complete the following steps to add credentials.

   a. Enter a name to identify the credentials in the **Name** field.

   b. Enter a username in the **User** field.

   c. Enter a password in the **Password** field.

5. Click **Create**.

   The Available Credentials page is displayed.

6. Click **Servers**.

   The Available Servers page is displayed.

7. Click **Add**.

   The Create Server page is displayed.

8. Complete the following steps to add a server.

   a. Enter the host IP address or host name in the **Hostname** field.

   b. Enter the management IP address in the **Management IP** field.

   c. Enter the management interface name in the **Management Interface** field.

   d. Select credentials from the **Credentials** list.

9. Click **Create**.

   The Available Servers page is displayed.

10. (Optional) Add another server.

   Follow steps 7 through 9 to add another server.

11. Click **Next**.

12. Complete the following steps to create a cluster.

   a. Select the **Contrail Enterprise Multicloud** option button.

   b. Enter the name for the cluster in the **Cluster Name** field.

   c. Enter container registry in the **Container Registry** field.

   d. Enter Contrail version information in the **Contrail Version** field.

   e. Select **Ansible** from the **Provisioner Type** list.

   f. Select **Enable ZTP** check box.

   g. Click **Next**.

13. Complete the following steps to assign control nodes.

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

     The vCenter Credentials section is displayed.

   b. Enter the following information:

     • Enter the vCenter IP address in the **vCenter IP Address** field.

     • In the **Data Center Name** field, enter the name of the data center under vCenter that CVFM will work on.

     • Enter the vCenter username in the **Username** field.

     • Enter the vCenter password in the **Password** field.

   c. Click **>**, next to the name of the server, to assign a server from the Available Servers table as a control node. The server is then added to the Assigned Control Nodes table.

     contrail_vcenter_fabric_manager_node is added to the list of roles.

> **NOTE:** Contrail Command does not provide any restrictions for High Availability mode. You can select more than one control node from the Assign control nodes table.

    d. Click **Next**.

14. Complete the following steps to add orchestrator information.

    a. Select **OpenStack** from the **Orchestrator Type** list.

    b. Click **>**, next to the name of the server, to assign a server from the Available Servers table as an orchestrator node. The server is then added to the Assigned Orchestrator Nodes table.

    c. Click **Next**.

15. Complete the following steps to assign compute nodes.

    a. Click **>**, next to the name of the server, to assign a server from the Available Servers table as a compute node.

       After the server is added to the Assigned Compute Nodes table, the **Default vRouter Gateway** and **Type** fields are enabled.

    b. Enter the gateway IP address in the **Default vRouter Gateway** field.

    c. Select **Kernel** from the **Type** list.

    d. Click **Next**.

16. (Optional) Assign service nodes.

    Click **>**, next to the name of the server, to assign a server from the Available Servers table as an service node. The server is then added to the Assigned Service Nodes table.

17. Click **Next**.

18. (Optional) Assign AppFormix nodes.

    Click **>**, next to the name of the server, to assign a server from the Available Servers table as an AppFormix node. The server is then added to the Assigned AppFormix Nodes table.

19. Click **Next**.

    The Cluster Overview page is displayed.

20. Click **Provision** after you have reviewed cluster overview information.

    The **Provisioning cluster <*Cluster Name*> in progress** bar is displayed. The CVFM plugin is enabled after the provisioning is completed.

    Click **Proceed to login** to log in to Contrail Command UI.

Deploying CVFM Plugin after Provisioning Contrail Command

Follow these steps to deploy the CVFM plugin after provisioning Contrail Command.

1. Navigate to **Infrastructure**>**Cluster** page in Contrail Command.

   The Overview tab is displayed.

2. Click **Manage vCenter** in the Control Nodes widget.

Figure 65: Overview Tab
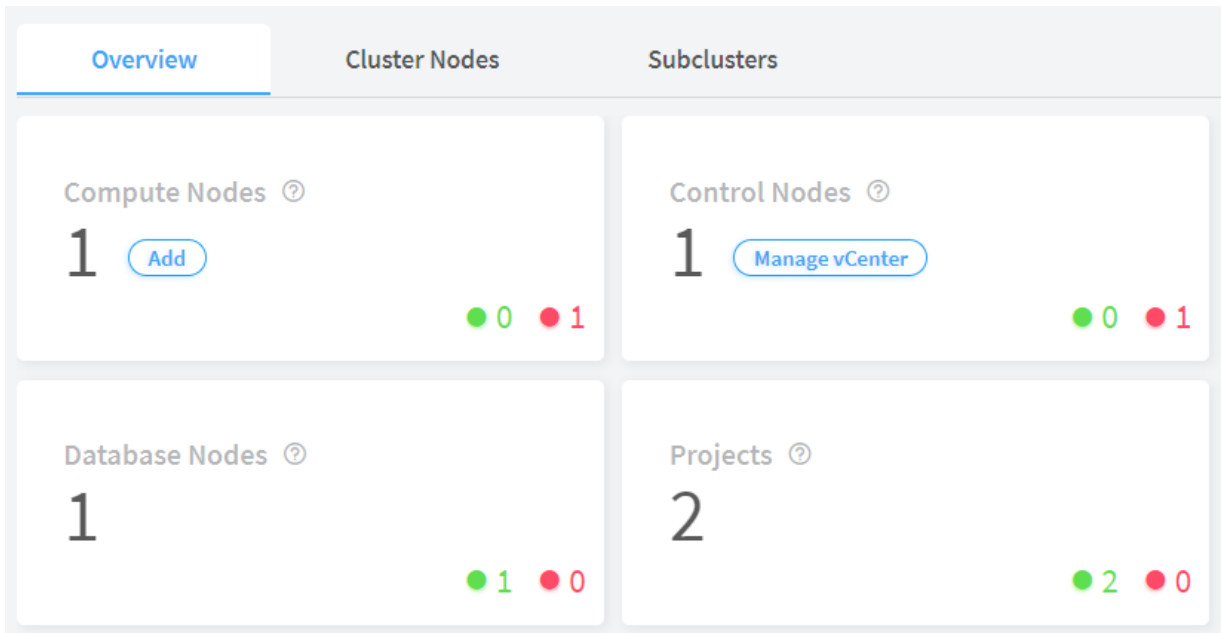
The Deploy vCenter page appears as shown in Figure 66 on page 205.

Figure 66: Deploy vCenter Page

3. Enter the following information:

Field	Description
vCenter IP Address	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 username.
Password	Enter the vCenter password.

4. 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 plugin on.

5. Click **Deploy**.

The CVFM plugin 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.

Release History Table

Release	Description
1910	Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plugin.

Fabric Discovery and ESXi Discovery by Using Contrail Command

IN THIS SECTION

-
-

Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plugin. After you deploy this plugin, 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.
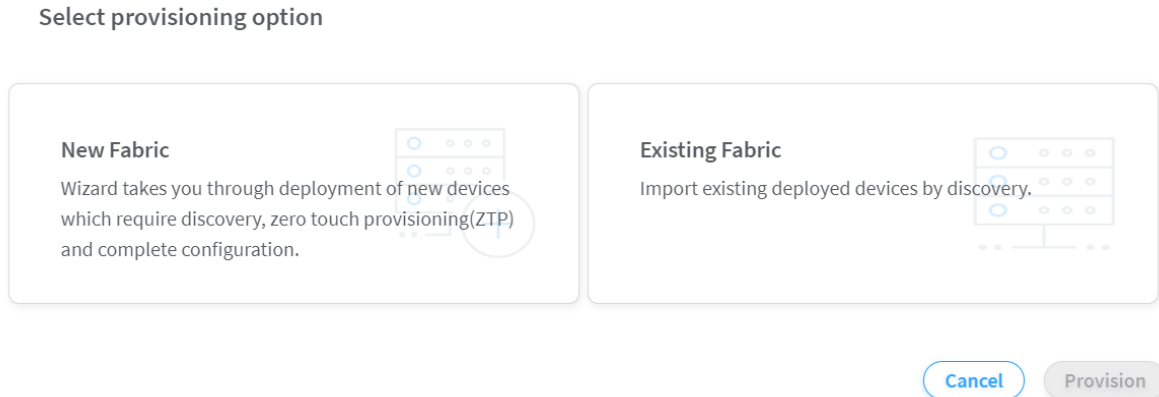
2. Click **Create**.

   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 67: Select Provisioning Option

Select provisioning option

New Fabric

Wizard takes you through deployment of new devices which require discovery, zero touch provisioning(ZTP) and complete configuration.

Existing Fabric

Import existing deployed devices by discovery.

Cancel   Provision

4. Click **Provision**.

   The Create Fabric page is displayed.

5. Enter the fabric provisioning information as listed in .

   Table 40: 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.
Node profiles	Add node profiles. You can add more than one node profile. All preloaded node profiles are added to the fabric by default. You can remove a node profile by clicking **X** on the node profile.

Table 40: Provision Existing Fabric *(continued)*

Field	Action
VLAN-ID Fabric Wide Significance	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. You can modify the enterprise style setting to service provider style once configured. However, you cannot modify the service provider style to enterprise style. 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 182. For more information on supported hardware platforms and roles, see "Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles" on page 138.
Username	Enter a username for the device.
Password	Enter a password for the device.
Management subnets	Enter the following information: **CIDR**—Enter CIDR network address. **Gateway**—Enter gateway address. 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.
Loopback subnets	Click **Loopback subnets** and enter loopback address in the **CIDR** field. NOTE: Loopback subnets are used to auto-assign loopback IP addresses to the fabric devices.

Table 40: Provision Existing Fabric *(continued)*

Field	Action
Underlay ASNs (eBGP)	Click **Additional Configuration** click **+Add** under **Underlay ASNs (eBGP)**. Enter autonomous system (AS) number in the range of 1-65,535 in the **CIDR** field. 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. • Enter minimum value in **ASN From** field. • Enter maximum value in **ASN To** field.
Fabric subnets (CIDR)	Click **+Add** under**Fabric subnets (CIDR)**. Enter fabric CIDR address in the **CIDR** field. NOTE: Fabric subnets are used to assign IP addresses to interfaces that connect to leaf or spine devices.
Advanced interface filters	Click **Advanced interface filters** and select the **Import configured interfaces** check box.

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 138.

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 68: Infrastructure > Servers > Servers Discovery

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 plugin.

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.

Release History Table

Release	Description
1910	Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plugin.
1909	Contrail Networking Release 1909 supports QFX10002-60C device running Junos OS Release 19.1R2 and later.

RELATED DOCUMENTATION

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 207.

Prerequisites

a. 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.

   1. Navigate to **DSwitch**.

   2. Click **Actions**.

3. Select **Settings** > **Edit Settings**.

   The Edit Settings page is displayed.

4. Click **Advanced**.

5. From the Discovery Protocol section, select **Link Layer Discovery Protocol** from the **Type** list.

6. Select an operational mode from the **Operation** list.

7. Click **OK** to confirm.

b. 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

Updating vCenter Credentials on Contrail Command

Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plugin. With this release, the CVFM plugin is installed when you install the Contrail Command user interface (UI).

You can enable this plugin:

- When you provision Contrail Command

- From the **Infrastructure**>**Cluster** page after you provision Contrail Command.

For more information, see .

You can also update vCenter credentials or override the configuration of the CVFM plugin after you have enabled the plugin. These steps provide instructions to update vCenter credentials from Contrail Command.

1. Navigate to the **Infrastructure**>**Cluster** page.

   The Overview tab is displayed.

2. Click **Manage vCenter** in the Control Nodes widget.

   The Update vCenter page appears.

3. You can update the following fields:

   - **vCenter IP Address**—Edit vCenter IP address.

   - **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.

4. Click **Update**.

   The credentials are updated.

Release History Table

Release	Description
1910	Contrail Networking Release 1910 supports the Contrail vCenter Fabric Manager (CVFM) plugin.

7

CHAPTER

Extending Contrail Networking to Bare Metal Servers

Bare Metal Server Management

IN THIS SECTION

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 with in 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 69 on page 219 shows an architectural view of the BMS Management framework.

Figure 69: BMS Management- Detailed Architecture View
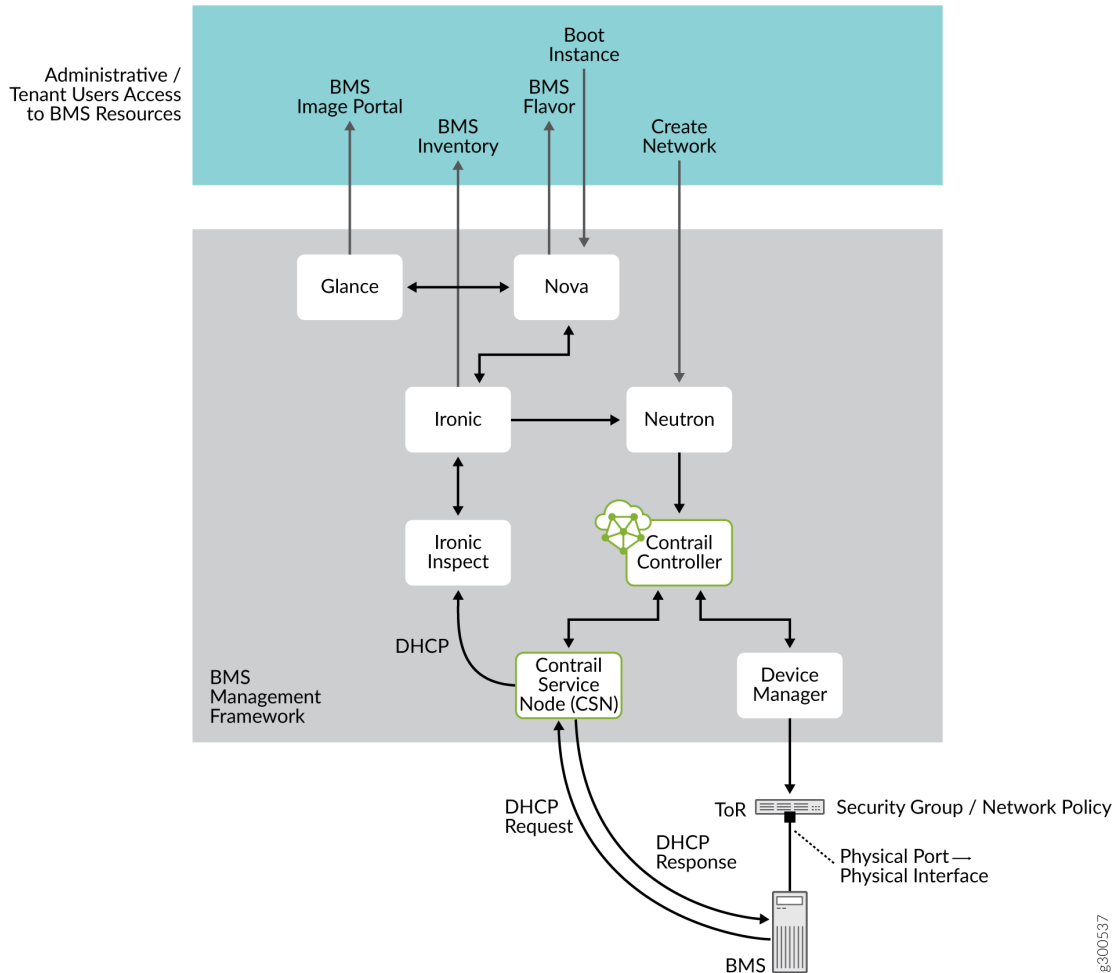
> **NOTE:** In single-tenant environments, administrative and tenant workflows are performed by the same user.

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

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 70.

  > **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 226.

- 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 228.

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 70 on page 225 shows how BMS is connected to a TOR switch.

Figure 70: Connectivity in Multihomed Configuration
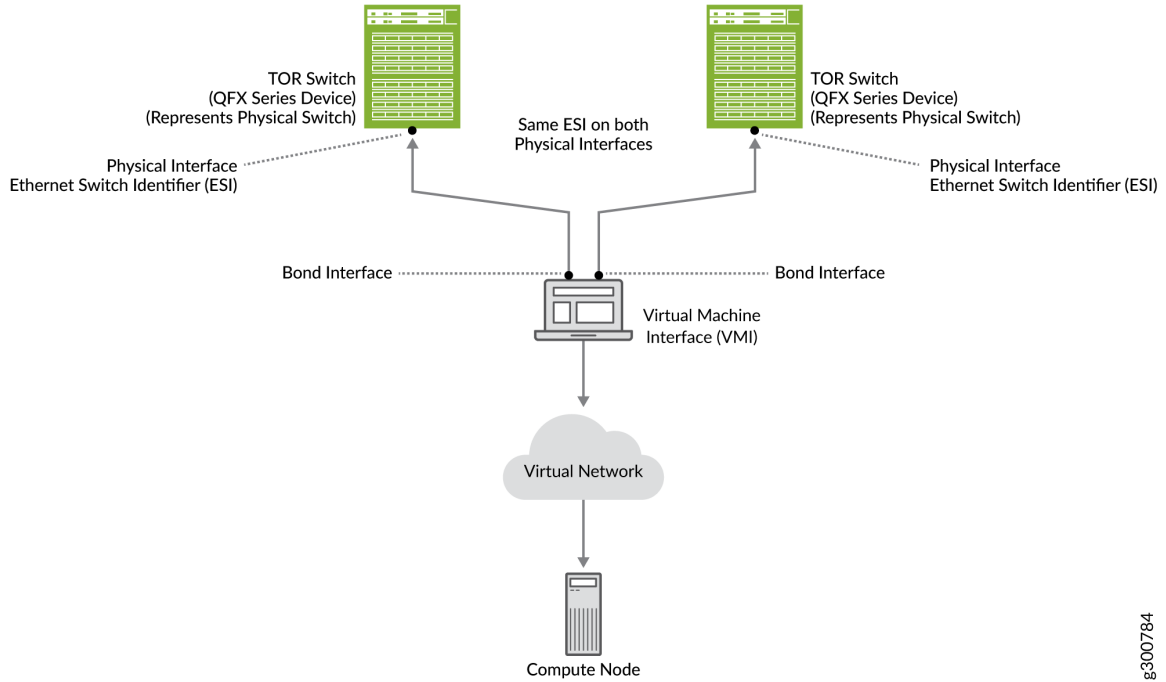
RELATED DOCUMENTATION

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.

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**pxe_ipmitool**. The driver value for Openstack SKU: rocky is **ipmi** . 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 x86_64.
Local gb	Enter Disk Size of BMS Server in gigabytes (Gb).
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.

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 41: 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.

Table 41: Add Existing Bare Metal Server Information *(continued)*

Field	Action
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.

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

5. Click **Create** to launch a new baremetal server.

RELATED DOCUMENTATION

Onboarding and Discovery of Bare Metal Servers

IN THIS SECTION

BMS Manager supports onboarding and discovery of bare metal servers.

Onboarding of Bare Metal Servers

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 tenant users. 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

IN THIS SECTION

- Manual Discovery | **230**
- Auto Discovery | **230**

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 section *Administrative Workflow*.

Auto Discovery

With Contrail Networking Release 5.1, 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

Launching and Deleting a Greenfield Bare Metal Server

This topic describes how to launch a greenfield bare metal server.

1. In the Contrail Command 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.

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