JUNIPER NETWORKS

Junos® OS

Cloud CPE Services on MX Series Routers Feature Guide

Release
14.1

Published: 2014-04-22

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Junos® OS Cloud CPE Services on MX Series Routers Feature Guide
14.1

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

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

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

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

Supported Platforms

For the features described in this document, the following platforms are supported:

- MX Series

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   system {
     scripts {
       commit {
         file ex-script.xsl;
       }
     }
   }
   interfaces {
     fxp0 {
       disable;
       unit 0 {
         family inet {
           address 10.0.0.1/24;
         }
       }
     }
   }
   ```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

   ```
   [edit]
   user@host# load merge /var/tmp/ex-script.conf
   load complete
   ```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   commit {
     file ex-script-snippet.xsl; }
   ```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

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

Table 1: Notice Icons

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

Table 2 on page xv defines the text and syntax conventions used in this guide.

Table 2: 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 2: 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		
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, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at

https://www.juniper.net/cgi-bin/docbugreport/. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Cloud CPE Services

CHAPTER 1

Cloud CPE Services Overview

Understanding How Cloud CPE Virtualizes Customer Premises Equipment (CPE) Services

Juniper Networks cloud CPE (cCPE) services enable service providers to offer their business enterprise customers virtual customer premises equipment (CPE) services. When you transition your customers to cCPE services, you can replace their complex and expensive Layer 3 CPEs with simple, low-cost Layer 2 CPE devices. Layer 3 CPE functions are provided by the MX Series 3D Universal Edge Router in the service provider cloud. cCPE uses existing Junos OS features to provide virtual or cloud CPE services.

Figure 1 on page 4 shows the cCPE services solution versus today's full-featured CPE.

Figure 1: Cloud-Based CPE Versus Full Featured CPE

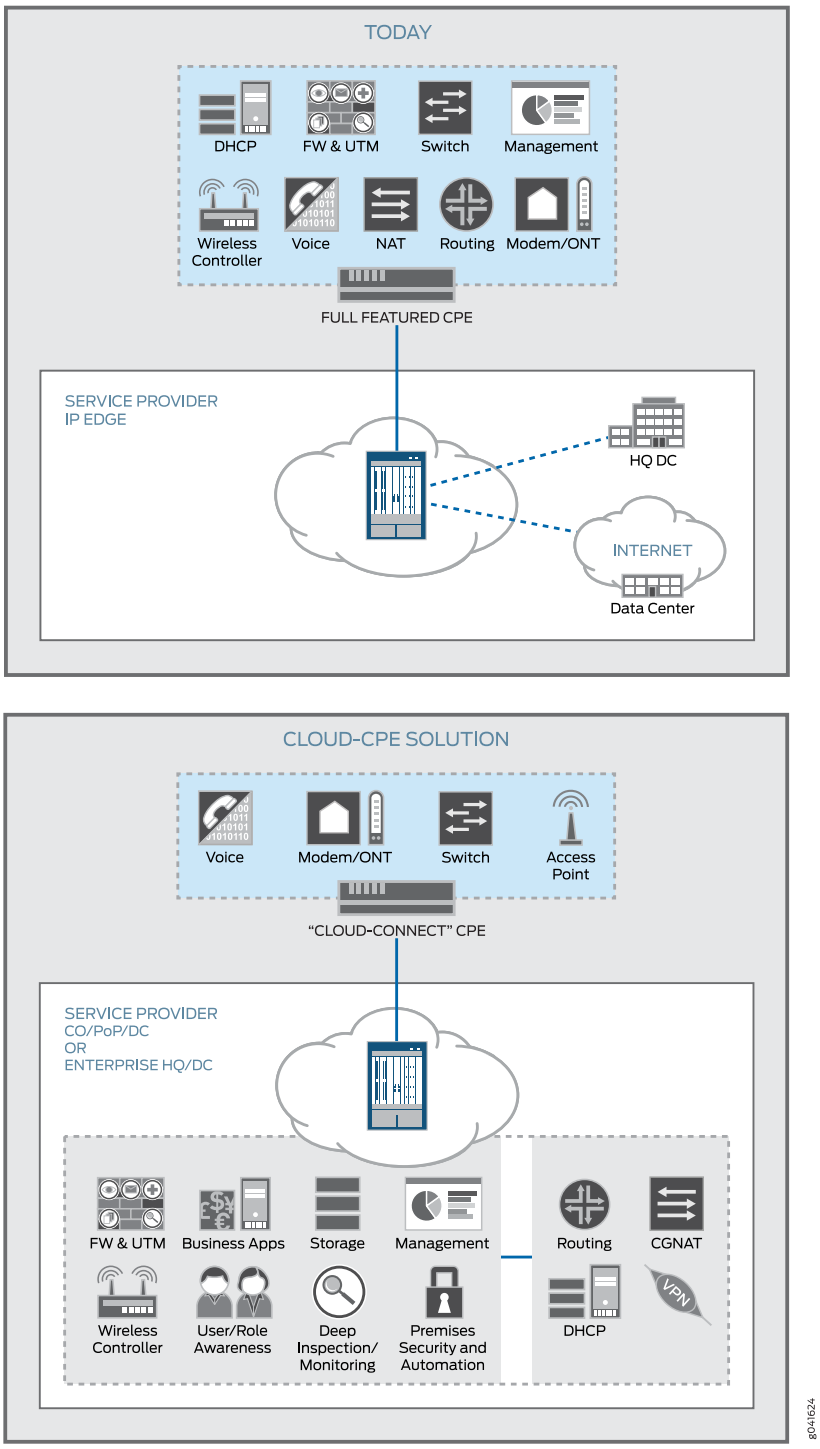

cCPE services redistribute traditional CPE features, reducing onsite CPE requirements like role, value, complexity, and cost, to a minimum and moves these features into the cloud, under service provider control. With cCPE services, the CPE evolves from a physical,

hardware, static, and feature paradigm to a logical, software, dynamic, and service paradigm.

cCPE services are most beneficial at enterprise business sites. The definition of a CPE is broad; however, in general, CPE is the demarcation device that attaches a customer site to the service provider (for example, a branch router). Some functions that are delivered today by other types of CPE devices (for example, security appliances, storage servers, WiFi hotspots, and so forth) may also shift partially or totally to the cloud CPE architecture.

The interfaces on the MX Series routers providing cCPE services must either be MultiServices PICs or MultiServices Dense Port Concentrators (MS-DPCs). These muiltiservice cards provide adaptive services interfaces, which allow you to coordinate multiple services on a single PIC by configuring a set of services and applications. MultiServices PICs and MS-DPCs offer a special range of services you configure in one or more service sets.

The MultiServices DPC is available for MX Series routers; it includes a subset of the functionality supported on the MultiServices PIC. Currently the MultiServices DPC supports the following Layer 3 services: stateful firewall, NAT, IDS, IPsec, active flow monitoring, RPM, and generic routing encapsulation (GRE) tunnels.

NOTE: **Not all Layer 3 services can be configured using the cCPE Selfcare application.**

Service providers can provision and manage cCPE services through the Junos OS CLI on the MX Series router, or the cCPE Selfcare application, which runs on the Junos Space Network Management Platform. In addition, APIs are provided for integrating the cCPE Selfcare application with your operations support systems (OSS). The cCPE Selfcare application also includes web portal that enables subscribers to monitor their cloud services and perform certain, secure configuration changes.

Basic Architecture of Cloud CPE

By using cCPE services, features typically provided by a Layer 2 or Layer 3 CPE are moved out to the network cloud resulting in a simplified and more profitable services architecture. The cCPE Selfcare application is an open platform focused on networking and security services including:

- Routing
- DHCP
- NAT & firewall
- Virtual Router Redundancy Protocol (VRRP)
- Unified Threat Management (UTM)
- WLAN controller (the hotspot is located at the subscriber site)
- Centralized management and provisioning
- Reporting

As an open platform, the cCPE Selfcare application is easily extensible to include new and adjacent services that can be provided by Juniper Networks or through other partnerships, including:

- Network-based storage
- Virtual desktop infrastructure
- Remote maintenance
- Physical site security and automation

Components of cCPE Services

cCPE services use existing Juniper Networks products like the MX Series router, Junos OS, and Junos Space Network Management Platform to move features from the physical CPE out to the network cloud by using the following components:

- An onsite, Layer 2 CPE.
- A cCPE routing context, residing in the provider edge (PE) MX Series router, providing aggregation, isolation, and service binding.
- A service complex, which can include Layer 3 and above services.
- A provisioning and monitoring system. cCPE services can be provisioned and monitored with the Junos OS CLI or the cCPE Selfcare application, which runs on the Junos Space Network Management Platform. APIs are available for the cCPE Selfcare application that enable you to provision and monitor cCPE services using your operations support systems (OSS).

Related Documentation

- cCPE End-To-End Solution Architecture on page 8
- Managing cCPE Services on page 10
- Understanding the cCPE Common Configuration on page 15

Benefits of Using Cloud CPE Services

Using cCPE services has many advantages as shown in Figure 2 on page 7.

Figure 2: Benefits of Cloud CPE

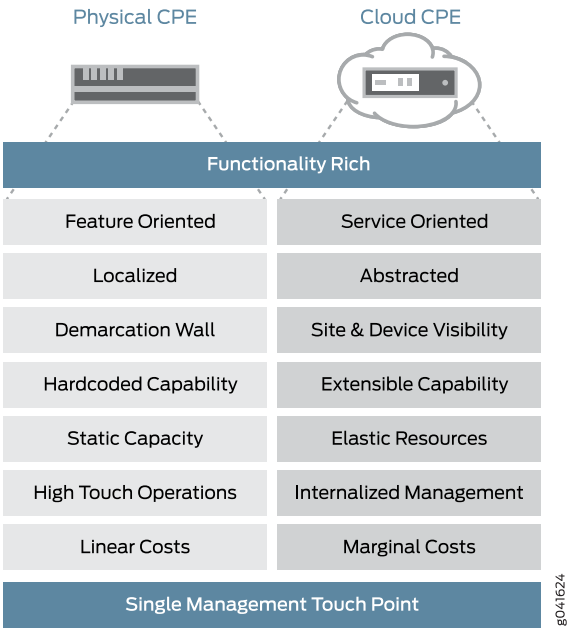

Some of the benefits of using cCPE services include:

- Decreased OPEX — Enabling a minimal physical Layer 2 CPE device is less problematic and requires less support, reducing your overall expenses. Running the majority of the current CPE software in the cloud provides a more stable, controlled, and reliable environment. Software updates are streamlined with a single image upgrading the software for 1,000's of customers in one reliable action.
- Improved efficiency in rolling out services — cCPE services allow you to more efficiently roll out and experiment with new services. Software changes in the cloud can be rolled out rapidly and efficiently, compared to changing the software on physical CPE devices. This results in increased average revenue per user (ARPU), as well as increased market penetration.
- Service oriented — Today's CPEs are very feature rich, even for low end models. However, most CPE deployments do not leverage all these features due to the complexity of configuring them. cCPE services enable you to package features and present them to your customers as services, which can be selected through a web interface. For example, a service could be "remote access." This service could allow access from the Internet to certain on site resources, as compared to a list of features that your subscriber needs to understand and configure (like MAC address, DHCP, NAT, Firewall, DDNS...). This type of service enables you to sell incremental services and also benefits your customers by providing them access to a wide variety of services without requiring network engineering expertise.
- Abstraction — As a cloud solution, CPE features can be highly distributed and are not limited to a single device.

- Device visibility — By moving the demarcation line from the subscriber site to your network, for example using DHCP and NAT cloud services, you have visibility to onsite devices. As a result, you can deploy device specific policies instead of only site granular policies and your subscriber's network can be virtually extended with your resources, for example storage servers. This also eases remote access, for example by mobile devices.
- Extensible capabilities — Physical CPEs are limited to what their firmware supports. Adding features may require a full CPE replacement. By using the cCPE services, you can add features on-demand like changing configurations and adding new virtual machines for supplementary services. This results in greater innovation capabilities and faster time to market.
- Elasticity — Physical CPEs are limited to what their CPU and memory can support. Requirements to process more packets or activate resource-intensive features may require a full CPE replacement. Conversely, the cloud CPE model benefits from horizontal and vertical scaling. More capacity can be added to the MX router by adding more service cards or the service complex by adding more appliances or server capability.
- Simplified management — Instead of having individual management sessions per CPE, there is a discrete number of management touch points.
- Marginal cost model — Traditional telecommunication economical models are based on resource oversubscription, where one resource can be shared across multiple customers. Physical CPEs break this model because they are not shared, implying a linear cost. Cloud CPEs have contexts and additional services that are multiplexed over shared resources, resulting in a marginal cost per CPE.

Related Documentation

- Understanding How Cloud CPE Virtualizes Customer Premises Equipment (CPE) Services on page 3
- cCPE End-To-End Solution Architecture on page 8
- Managing cCPE Services on page 10
- APIs for the cCPE Selfcare Application on page 12
- Understanding the cCPE Common Configuration on page 15

cCPE End-To-End Solution Architecture

Figure 3 on page 9 shows the overall cCPE architecture.

Figure 3: Cloud CPE Architecture

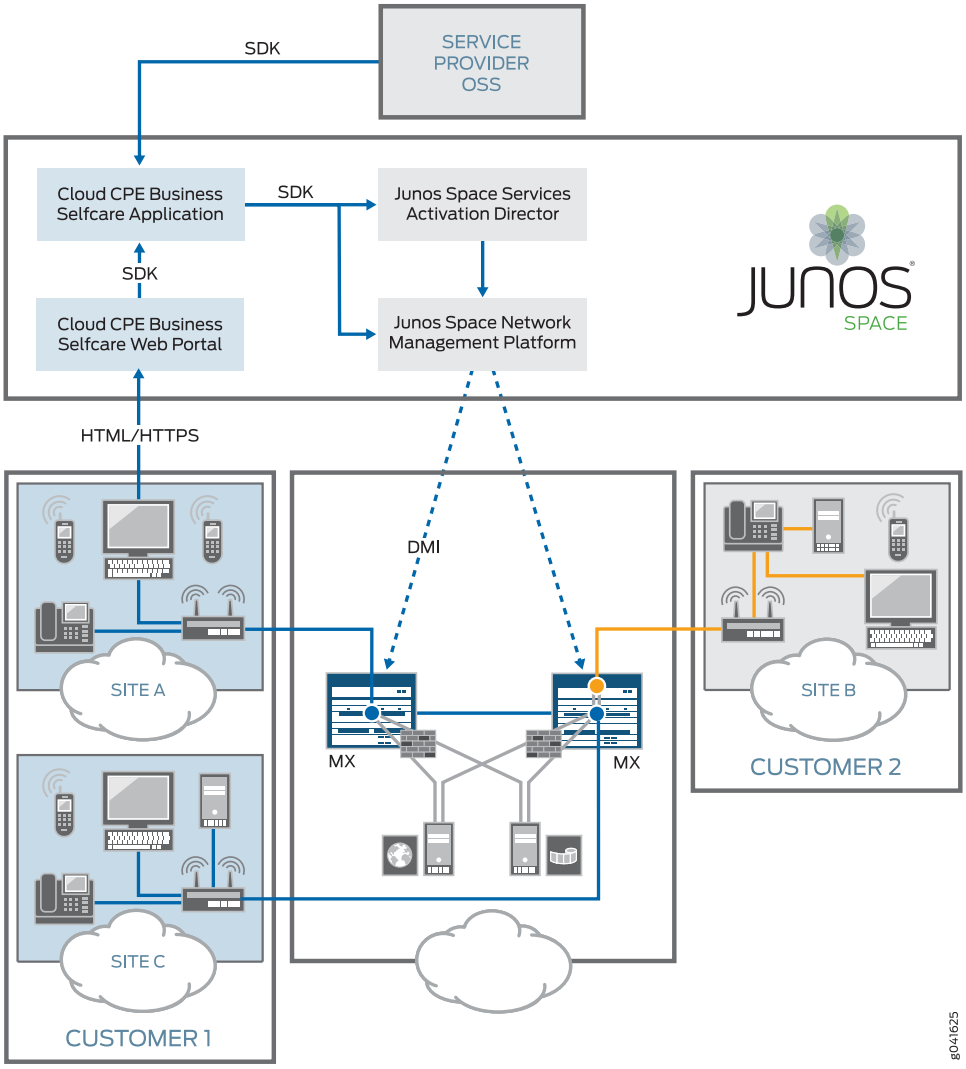

In Figure 3 on page 9, one subscriber has two VPN sites (Site A and Site C) with Internet access. A second subscriber (Site B) has a single VPN site with Internet access. The network consists of MX Series routers providing Layer 2 or Layer 3 connectivity (Internet and VPN access) to the CPEs at the subscriber sites. The cCPE context for each subscriber site is implemented by a routing instance to provide separate routing, DHCP, SNMP, and other services to each subscriber site. Layer 3 connectivity, including NAT for Internet access, and stateful firewalls is implemented by Multiservice-Dense Port Concentrator (MS-DPC) interfaces in the router.

You can use the cCPE Selfcare application to provision and manage cCPE services. The cCPE Selfcare application runs on the Junos Space Network Management Platform. Optionally, you can use Junos Space Services Activation Director to provision the network, create subscriber sites, and associate Layer 3 VPN services with subscribers. If you use Junos Space Services Activation Director, the cCPE Selfcare application can import customers and their associated network inventory (interfaces on the router connecting

to subscriber sites) from Junos Space Services Activation Director. Otherwise, the cCPE Selfcare application provides its own API to enable importing this information from your OSS.

In addition to inventory management functions, the cCPE Selfcare application implements monitoring and service change APIs that enable service providers to integrate their OSS.

The cCPE Selfcare Web portal provides management and monitoring capabilities for enterprise business customers. Using the Web portal, authorized administrators can change certain configuration parameters and view states and statistics for their cCPE resources. For security purposes, the cCPE Selfcare portal exposes only configuration capabilities that are safely changed by the end subscriber; changing the configuration does not affect other cCPE subscribers.

Related Documentation

- Understanding How Cloud CPE Virtualizes Customer Premises Equipment (CPE) Services on page 3
- Managing cCPE Services on page 10
- APIs for the cCPE Selfcare Application on page 12

Managing cCPE Services

You can manage cCPE services in several ways:

- CLI — Service providers can use the Junos OS CLI to provision, manage, and monitor cCPE services.
- cCPE Selfcare application — Service providers can use the cCPE Selfcare application, which runs on the Junos Space Network Management Platform to provision, manage, and monitor cCPE services.
- cCPE Selfcare Web portal — Enterprise business customers can use the cCPE Selfcare Web portal to monitor their cCPE services and perform certain configuration changes.
- Operations support systems (OSS) — Service providers can integrate their OSS with the cCPE Selfcare application APIs to manage and provision cCPE services.

cCPE Selfcare Application

The cCPE Selfcare application:

- Provides configuration, management, and monitoring of cCPE services for service providers.
- Is a Java 2 Platform, Enterprise Edition (J2EE) application based on the Junos Space Network Management Platform software development kit (SDK).
- Acts as the backend server processing configuration, management, and monitoring requests from service providers and subscribers.

- Includes a representational state transfer API that service providers can use to integrate their operations support systems (OSS).
- Uses the Junos Space Network Management Platform authentication and permission control to expose cCPE resources only to authorized administrators at the business subscriber site. For security purposes, only certain configuration capabilities are available to subscribers using the cCPE Selfcare Web portal. Configuration changes made by one subscriber do not affect other subscribers.

The Junos Space Network Management Platform provides network management through the Device Management Interface (DMI) for network devices running Junos OS and is the standard network management platform for Juniper Networks devices. Through the Junos Space Network Management Platform SDK, which contains APIs, applications can change configuration of network devices and issue operational commands to monitor managed devices.

cCPE Selfcare Web Portal

The cCPE Selfcare Web portal:

- Provides configuration, management, and monitoring capabilities for enterprise business subscribers running cCPE services. Authorized administrators at the business subscriber site can make certain configuration changes and view states and statistics of cCPE resources.
- Exposes only configuration capabilities that are safely changed by the subscriber. Configuration changes made by one subscriber do not affect other subscribers.

Downloading and Installing the cCPE Selfcare Application Packages

You install the cCPE Selfcare application like any other application using the Junos Space application management functions. For overall information on installing Junos Space applications, see Application Management Overview. For information about installing Junos Space applications, see Adding a Junos Space Application.

You can download the cCPE Selfcare application from the Juniper Networks software download page, see cCPE Selfcare Application .

NOTE: **This release of the cCPE Selfcare application runs only on Junos Space Network Management Platform release 12.3P2.8. The cCPE Selfcare application is not currently supported on Junos Space Network Management Platform 13.x releases. To run the cCPE Selfcare application on the Junos Space Network Management platform, you first need to download and install the Junos Space 12.3R1.x software image. You then need to upgrade Junos Space by downloading and installing the Junos Space 12.3P2.8 patch image.**

Related Documentation

- Understanding How Cloud CPE Virtualizes Customer Premises Equipment (CPE) Services on page 3
- cCPE End-To-End Solution Architecture on page 8

APIs for the cCPE Selfcare Application

The cCPE Selfcare application provides a representational state transfer API for cCPE monitoring and configuration tasks. It exposes all the service provider and subscriber self-care functionality available through the cCPE Selfcare application and can be used for OSS integration. The representational state transfer API consists of three distinct interfaces:

- Subscriber management API— The subscriber management API provides the following functionality:
  - Enables service providers to create, modify, delete, and retrieve customers and their associated sites and access links.
  - Enables service providers to import customers associated with Layer 3 VPNs provisioned by Network Activate.
  - Enables network administrators at the subscriber site to make configuration changes specific to their cCPE services, as well as provide easy-to-understand names and descriptions to sites and access links. The representational state transfer API hides all network specific information associated with the enterprise administrators access links to the service provider network.
- Configuration management API— The configuration management API provides access to the cCPE configuration on the router. Enterprise administrators can use this API to read, modify, and delete router configuration. Service providers can use this API to provision services on behalf of their subscribers. Using the configuration management API, you can configure:
  - Private IP addresses for an access link
  - Static routes for a VPN site
  - DHCP server and DHCP relay configuration for a VPN site
  - SNMP configuration for a VPN Site
- Monitoring API — Service providers and enterprise administrators can use this API to monitor and audit the health of their cCPE services. It also enables the retrieval of graphs associated with collected SNMP data. This API supports the following functionality:
  - Functional audit — Enables the retrieval of functional audit results. Using the representational state transfer API, you can schedule functional audits to verify the connectivity between different sites of a particular subscriber.
  - Status monitoring — Enables the retrieval of traffic counters for all access links or sites associated with a subscriber.

- Historical data — Enables the retrieval of graphical views for SNMP data using OpenNMS. Because OpenNMS is built into the Junos Space Network Management Platform, no additional software is required. You can specify the type of graph you want to display based on the selected MIB and time interval for the data.

Related Documentation

- Understanding How Cloud CPE Virtualizes Customer Premises Equipment (CPE) Services on page 3
- cCPE End-To-End Solution Architecture on page 8
- Managing cCPE Services on page 10
- Understanding the cCPE Common Configuration on page 15

CHAPTER 2

Configuring the Cloud CPE Common Configuration

Understanding the cCPE Common Configuration

Using cCPE services on MX Series routers, each subscriber VPN site is represented by one VRF routing instance. VPN routes are stored in the VRF routing table and exported to remote sites through internal BGP (IBGP). VPN routes from remote sites received through IBGP are imported into the VRF table. The VRF export and import polices control how routes are passed between sites.

With full-featured Layer 2 and Layer 3 CPEs, the routing instance learns the VPN routes from the local VPN site through routing protocols or a static route. However, when you use cCPE services, the local site appears as a direct route in the Junos OS because hosts connect to the PE router directly through a Layer 2 switch. MX Series routers support direct routes in the VRF and export direct routes through BGP to remote sites.

cCPE services allow heterogeneous setup; sites using the cCPE services can coexist with sites using a full-featured Layer 3 CPE in the same VPN.

To use cCPE services, you must configure the MX Series router with a specific configuration. We refer to this configuration as the *cCPE common configuration* because it is required to run any cCPE services. The common configuration includes defining the physical access link to your subscriber sites; configuring the cCPE context, which is a VRF routing instance for all VPN services; and configuring a generic routing policy for advertising routes. The cCPE common configuration requires a Gigabit Ethernet interface be used as the access link to the subscriber site. The logical interface is configured for Layer 2 mode (**vlan-bridge**) encapsulation and integrated routing and bridging.

NOTE: *Minimum hardware requirements*—cCPE services require a minimum of an MX80 Series 3D Universal Edge Router with MPCs/MICs.

NOTE: Configuration of MPLS and BGP between PE routers is not described here, because their configuration is the same as the configuration of traditional MPLS/BGP VPNs with a Layer 3 CPE.

Related Documentation

- Understanding How Cloud CPE Virtualizes Customer Premises Equipment (CPE) Services on page 3
- cCPE End-To-End Solution Architecture on page 8
- Managing cCPE Services on page 10
- APIs for the cCPE Selfcare Application on page 12
- Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI on page 16

Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI

Configure the cCPE common configuration on each MX Series router that provides cloud services. Complete the following tasks to configure the cCPE common configuration:

1. Configuring the Subscriber Access Link on the PE Router for the cCPE Common Configuration on page 16
2. Configuring the Layer 2 IRB Interface for the cCPE Common Configuration on page 18
3. (Optional) Configuring a Private Subnet on the IRB Interface on page 18
4. Configuring the Bridge Domains for the cCPE Common Configuration on page 19
5. Configuring the Firewall Filters and Policers for the cCPE Common Configuration on page 20
6. Configuring the VPN Routing Instances for the cCPE Common Configuration on page 22
7. Configuring the VRF Import Routing Policies for the cCPE Common Configuration on page 23
8. Configuring the VRF Export Routing Policies for the cCPE Common Configuration on page 24

Configuring the Subscriber Access Link on the PE Router for the cCPE Common Configuration

Complete the following tasks to configure the cCPE access link:

1. Specify the name of the physical interface being used to connect to the subscriber.

   ```
   [edit]
   user@host# edit interfaces type-fpc/pic/port
   ```

   For example:

   ```
   [edit]
   user@host# edit interfaces ge-1/1/1
   ```

2. Configure a description to distinguish the subscriber interface.

```
[edit interfaces ge-1/1/1]
user@host# set description description
```

3. Configure the speed of the interface.

```
[edit interfaces ge-1/1/1]
user@host# set speed (auto | 1Gbps | 100Mbps | 10Mbps)
```

4. Enable CoS hierarchical scheduling on the interface.

```
[edit interfaces ge-1/1/1]
user@host# set hierarchical-scheduler
```

5. Configure the encapsulation type for the interface for flexible Ethernet services.

```
[edit interfaces ge-1/1/1]
user@host# set encapsulation flexible-ethernet-services
```

6. Enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

```
[edit interfaces ge-1/1/1]
user@host# set vlan-tagging
```

7. Configure the link mode for full duplex.

```
[edit interfaces ge-1/1/1]
user@host# set link-mode full-duplex
```

8. Configure the interface for autonegotiation.

```
[edit interfaces ge-1/1/1]
user@host# set gigether-options auto-negotiation
```

9. Create a logical interface and configure it with a unique description that identifies the subscriber.

```
[edit interfaces ge-1/1/1]
user@host# edit unit interface-unit-number
user@host# set description description
```

10. Configure the encapsulation type on the logical interface as Layer 2 Ethernet VLAN bridge encapsulation.

```
[edit interfaces ge-1/1/1 unit 105]
user@host# set encapsulation vlan-bridge
```

11. Configure the subscriber VLAN ID that you want to bind to the logical interface.

```
[edit interfaces ge-1/1/1 unit 105]
user@host# set vlan-id 105
```

12. Verify the configuration.

```
user@host> show interfaces ge-1/2/3
ge-1/1/1 {
    description "VPN-cCPE for ifd-acme-abc";
    hierarchical-scheduler;
    vlan-tagging;
    speed 1g;
    link-mode full-duplex;
    encapsulation flexible-ethernet-services;
    gigether-options {
        auto-negotiation;
    }
```

```
unit 105 {
    description "ifl-acme-site-abc";
    encapsulation vlan-bridge;
    vlan-id 105;
}
```

Configuring the Layer 2 IRB Interface for the cCPE Common Configuration

For the cCPE common configuration, you need to configure the logical interface as an integrated routing and bridging (IRB) interface. For private IP addresses, you can configure multiple addresses on different subnets on the IRB interface. Multiple VPN routes are advertised through the VPN routing protocol, like BGP, to the remote VPN sites. Communication between hosts on different subnets, but the same LAN, goes through the IRB interfaces on the edge router because their gateway addresses are on the IRB interface. To configure the IRB interface:

1. Configure the logical interface as an IRB interface, and provide a description that identifies the subscriber.

   Be sure to specify the unit number from the physical interface used for the subscriber access link.

   [edit]
   user@host# **edit interfaces irb unit 105**

2. Specify a description for the IRB interface that identifies the subscriber.

   [edit interfaces irb unit 105]
   user@host# **set description *description***

3. Specify the IPv4 subnet (subscriber-facing IP address/prefix) for private addresses for the subscriber VPN site.

   [edit interfaces irb unit 105]
   user@host# **set family inet address *address***

4. Specify the bandwidth for the IRB interface.

   [edit interfaces irb unit 105]
   user@host# **set bandwidth *bandwidth***

5. Review the configuration of the IRB interface.

   [edit interfaces irb unit 105]
   user@host# **show interfaces irb**

```
description "IRB interface of Example Customer";
bandwidth 1g;
family inet {
    address 192.168.1.1/24;
}
```

(Optional) Configuring a Private Subnet on the IRB Interface

If the subscriber has set up multiple private subnets in one site, the IRB interface connecting this subscriber site to the PE router needs to be configured with multiple private subnets.

To configure a private subnet:

- [edit]
  user@host# **set interfaces irb unit *vlan-unit-id* family inet address *private customer ip 2/prefix***

Configuring the Bridge Domains for the cCPE Common Configuration

To configure the bridge domains for the cCPE common configuration, you need to associate the IRB, the physical interface, and the VLAN with the bridge domain. Configure one bridge domain for each subscriber site. To configure the bridge domain:

1. Specify the domain name and configure the domain type as **bridge**.

   [edit]
   user@host# **edit bridge-domains acme-bd**

2. Define the bridge domain type as **bridge**.

   [edit bridge-domains acme-bd]
   user@host# **set domain-type bridge**

3. Associate the subscriber's VLAN ID with the bridge domain.

   [edit bridge-domains acme-bd]
   user@host# **set vlan-id 105**

4. Specify the routing interface to include in the bridge domain.

   [edit bridge-domains acme-bd]
   user@host# **set routing-interface irb.105**

5. Specify the logical interfaces to include in the bridge domain.

   [edit bridge-domains acme-bd]
   user@host# **set interface ge-1/1/1.105**

6. Specify the maximum number of MAC addresses allowed to be learned for the bridge domain, and specify that packets for new source MAC addresses be dropped after the MAC address limit is reached.

   [edit bridge-domains acme-bd]
   user@host# **set bridge-options interface-mac-limit 64 packet-action drop**

7. Review the configuration of the bridge domain.

   [edit bridge-domains acme-bd]
   user@host# **show**

```
domain-type bridge;
vlan-id 105;
interface ge-1/1/1.105;
routing-interface irb.105;
bridge-options {
    interface-mac-limit {
        64;
        packet-action drop;
    }
}
```

Configuring the Firewall Filters and Policers for the cCPE Common Configuration

Create a firewall filter and a policer for the bridge domain. Traffic policing is an essential component of network access security designed to minimize the risks of denial of service (DoS) attacks. It enables the control of the maximum rate of traffic sent or received on an interface.

1. Configure an IPv4 firewall filter for Layer 2 traffic. Configure the filter to track Address Resolution Protocol (ARP) packets. In the following procedure, ARP packets are policed by a policer called ARP-Policer, and counted by a counter called ARP-Count. Ethernet packets using ARP are accepted.

   ```
   [edit]
   user@host# edit firewall family bridge filter L2-Traffic
   [edit firewall family bridge filter L2-Traffic]
   user@host# edit term ARP
   [edit firewall family bridge filter L2-Traffic term ARP]
   user@host# set from ether-type arp
   user@host# set then policer ARP-Policer
   user@host# set then count ARP-Count
   user@host# set then accept
   ```

2. Configure filters for broadcast, multicast, and unicast traffic. In this example, broadcast, multicast, and unicast packets are policed by a policer called BMU-Policer and counted by a counter called BMU-Count. Any other types of packets are discarded.

   ```
   [edit firewall family bridge filter L2-Traffic term ARP]
   user@host# up
   [edit firewall family bridge filter L2-Traffic]
   user@host# edit term BMU
   [edit firewall family bridge filter L2-Traffic term BMU]
   user@host# set from traffic-type broadcast
   user@host# set from traffic-type multicast
   user@host# set from traffic-type unknown-unicast
   user@host# set then policer BMU-Policer
   user@host# set then count BMU-Count
   user@host# set then accept
   user@host# up
   user@host# edit term DROP
   [edit firewall family bridge filter L2-Traffic term DROP]
   user@host# set then discard
   ```

3. Configure the ARP policer traffic limits and action to take on nonconforming traffic.

   ```
   [edit firewall family bridge filter L2-Traffic term DROP]
   user@host# top
   user@host# edit firewall policer ARP-Policer
   [edit firewall policer ARP-Policer]
   user@host# set filter-specific
   user@host# set if-exceeding bandwidth-limit 8k
   user@host# set if-exceeding burst-size-limit 1500
   user@host# set then discard
   ```

4. Configure the broadcast, multicast, and unicast policer traffic limits and action to take on nonconforming traffic.

```
[edit firewall policer ARP-Policer]
user@host# up
user@host# edit policer BMU-Policer
[edit firewall policer BMU-Policer]
user@host# set filter-specific
user@host# set if-exceeding bandwidth-limit 8k
user@host# set if-exceeding burst-size-limit 1500
user@host# set then discard
```

5. Apply the firewall filters and policers to the bridge domain.

```
[edit firewall policer BMU-Policer]
user@host# top
user@host# edit bridge-domains acme-bd forwarding-options filter
user@host# set input L2-Traffic
```

6. Review the configuration.

```
user@host>  show firewall
family bridge {
    filter l2-traffic {
        term ARP {
            from {
                ether-type arp;
            }
            then {
                policer ARP-Policer;
                count ARP-Count;
                accept;
            }
        }
        term BMU {
            from {
                traffic-type [ broadcast multicast unknown-unicast ];
            }
            then {
                policer BMU-Policer;
                count BMU-Count;
                accept;
            }
        }
        term DROP {
            then discard;
        }
    }
}
policer ARP-Policer {
    filter-specific;
    if-exceeding {
        bandwidth-limit 8k;
        burst-size-limit 1500;
    }
    then discard;
}
policer BMU-Policer {
    filter-specific;
    if-exceeding {
        bandwidth-limit 8k;
        burst-size-limit 1500;
    }
```

```
        then discard;
    }
```

Configuring the VPN Routing Instances for the cCPE Common Configuration

For the common configuration, you must configure a routing instance that supports Layer 3 VPNs. To configure the routing instance:

1. Configure a name for the routing instance.

   ```
   [edit]
   user@host# edit routing-instances ri-acme-site-abc
   ```

2. Configure a unique description to identify the routing instance.

   ```
   [edit routing-instances ri-acme-site-abc]
   user@host# set description “routing-instance for acme site-abc”
   ```

3. Configure the routing instance as a VRF instance.

   ```
   [edit routing-instances ri-acme-site-abc]
   user@host# set instance-type vrf
   ```

4. (Optional-Required only when using IRB interfaces) Associate the Layer 3 interface with the subscriber.

   ```
   [edit routing-instances ri-acme-site-abc]
   user@host# set interface irb.105
   ```

5. Specify a route distinguisher for the routing instance, enabling you to distinguish which VPN the route belongs to. Each routing instance must have a unique route distinguisher associated with it. The route distinguisher is used to place boundaries around a VPN so that the same IP address prefixes can be used in different VPNs without having them overlap. The format is *as-number:id*.

   ```
   [edit routing-instances ri-acme-site-abc]
   user@host# set route-distinguisher as-number:id
   ```

6. Reference the VRF import and export policies.

   ```
   [edit routing-instances ri-acme-site-abc]
   user@host# set vrf-import acme-import
   user@host# set vrf-export acme-export
   ```

7. Map the inner label of a packet to a specific VRF table. This allows examination of the encapsulated IP header. All routes in the VRF configured with this option are advertised with the label allocated per VRF.

   

   NOTE: **vrf-table-label** is mandatory when the PE router to customer edge (CE) router connection is over a shared medium.

   ```
   [edit routing-instances ri-acme-site-abc]
   user@host# set vrf-table-label
   ```

8. Review the configuration.

   ```
   user@host> show routing-instances
   ri-acme-site-abc {
   ```

```
description routing-instance for acme site-abc;
instance-type vrf;
interface irb.105;
route-distinguisher 65535:1;
}
vrf-import acme-import;
vrf-export acme-export;
vrf-table-label;
```

Configuring the VRF Import Routing Policies for the cCPE Common Configuration

For the common configuration, you need to specify the VRF import routing policies. The import policy is always based on an IBGP session between the PE routers; therefore, BGP is the protocol used for the import policy. To configure the VRF import routing policies, specify the following options:

1. Specify the protocol used between the PE routers.

   ```
   [edit]
   user@host# edit policy-options policy-statement acme-import-policy term a from
     protocol bgp
   user@host# set term a from protocol bgp
   ```

2. (Optional) If the protocol used between the edge routers is BGP, specify the BGP community.

   ```
   [edit policy-options policy-statement acme-import-policy]
   user@host# set term a from community vpn-acme-site-abc
   ```

3. Specify the match condition actions to take for import routing policies.

   ```
   [edit policy-options policy-statement acme-import-policy]
   user@host# set term a then accept
   user@host# set term b then reject
   ```

4. Review the configuration of the import policies.

   ```
   user@host# show
   ```

   ```
   term a {
       from {
           protocol bgp;
           community vpn-acme;
       }
       then accept;
   }
   term b {
       then reject;
   }
   ```

Configuring the VRF Export Routing Policies for the cCPE Common Configuration

For the common configuration, you need to specify the VRF export routing policies to other sites that are in the same VPN. The type of policies you define depends on the type of routing protocol that is configured between the PE routers, the cCPE, and the customer edge (CE) router. PE routers always use the IBPG protocol. To configure the VRF export routing policies, specify the following options:

1. Specify the routing protocol used for routing into the customer's LAN (between the cCPE and the CE router).

   ```
   [edit]
   user@host# edit policy-options policy-statement acme-export-policy
   user@host# set term a from protocol direct
   ```

2. Add the community to IBGP session.

   ```
   [edit policy-options policy-statement acme-export-policy]
   user@host# set term a then community add vpn-acme-site-abc
   ```

3. Specify the export routing policies.

   ```
   [edit policy-options policy-statement acme-export-policy]
   user@host# set term a then accept
   user@host# set term b then reject
   ```

4. (Optional) If the protocol used is BGP, add the subscriber route as a community member.

   ```
   [edit]
   user@host# set policy-options community vpn-acme-site-abc members target:65535:5
   ```

5. ```
   user@host> show

       term a {
           from {
               protocol direct;
               community vpn-acme-site-abc;
           }
           then accept;
       }
       term b {
           then reject;
       }
   }
   community vpn-acme-site-abc members target:65535:5;
   ```

Related Documentation

- Understanding the cCPE Common Configuration on page 15
- Understanding How Cloud CPE Virtualizes Customer Premises Equipment (CPE) Services on page 3
- cCPE End-To-End Solution Architecture on page 8
- Managing cCPE Services on page 10
- APIs for the cCPE Selfcare Application on page 12
   ```

CHAPTER 3

Configuring NAT Services for the cCPE Application

Understanding How to Assign Public and Private Addresses to VPN Subscribers Using NAT cCPE Services

This topic describes how to use the MX Series router Network Address Translation (NAT) function with the cCPE services.

Assigning Primary IP Addresses with NAT cCPE Services

Using the NAT cCPE services, you can assign both public and private IP addresses to VPN subscribers. Public addresses are assigned from the source address pool defined for the NAT service. If the subscriber LAN is connected directly to the edge router, the private address range is defined on the integrated routing and bridging (IRB) interface. When there are routers in the subscriber network, the private network is defined on those routers, which advertise the routes to the VPN routing instance. Alternatively, you can define static routes in the VPN routing instance.

Assigning Multiple Address Ranges (Primary and Secondary Addresses) with NAT cCPE Services

The NAT cCPE service assigns multiple IP address ranges and primary IP addresses in similar ways. To configure multiple address ranges, you define the range in the NAT address pool and on the IRB interface. To support multiple public address ranges, you configure multiple NAT address pools.

Related Documentation

- Configuring NAT cCPE Services to Assign Primary IP Addresses on page 26
- Configuring NAT cCPE Services to Assign Multiple IP Address Ranges (Primary and Secondary Addresses) on page 29

- Verifying NAT cCPE Services on page 30
- Understanding the cCPE Common Configuration on page 15
- Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI on page 16
- cCPE End-To-End Solution Architecture on page 8

Configuring NAT cCPE Services to Assign Primary IP Addresses

This topic describes how to configure NAT cCPE services to assign primary IP addresses to VPN subscribers.

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

Then complete the following tasks to configure NAT cCPE services to assign primary IP addresses:

Then complete the following tasks to configure NAT cCPE services to assign primary IP addresses:

1. Enable NAT on the PE router. See "Configuring NAT on MX Series Routers for cCPE Services" on page 26.
2. Configure the service interfaces. See "Configuring the Service Interfaces for NAT cCPE Services" on page 27.
3. Configure the default route for the cCPE context. See "Configuring the Static Route for the cCPE Context" on page 27.
4. Configure the public address pools. See "Configuring Public Address Pools for NAT cCPE Services" on page 28.
5. Configure the NAT rules. See "Configuring NAT Rules for cCPE Services" on page 28.
6. Configure the logical interfaces. See "Configuring the Inbound and Outbound Logical Interfaces for NAT cCPE Services" on page 29.

Configuring NAT on MX Series Routers for cCPE Services

The Junos OS provides carrier-grade NAT (CGN) for IPv4 and IPv6 networks. To enable the NAT service for the cCPE services:

1. Configure the properties for the MS-DPC and enable the extension provider service package application.

   ```
   [edit]
   edit chassis fpc 4 pic 0 adaptive-services service-package layer-3
   ```

2. Verify the configuration.

   ```
   user@host# show
   ```

```
chassis {
fpc 4 {
    pic 0 {
        adaptive-services {
            service-package layer-3;
        }
    }
}
}
```

Configuring the Service Interfaces for NAT cCPE Services

To configure the service interfaces:

1. Configure the inbound physical and logical interface names, specify the IPv4 protocol (inet), and specify the interface as an inside service domain.

   ```
   [edit]
   user@host# edit interfaces sp-4/0/0 unit 1
   user@host# set service-domain inside
   user@host# set family inet
   ```

2. Configure the outbound physical and logical interface names, specify the IPv4 protocol (inet), and specify the interface as an outside service domain.

   ```
   [edit]
   user@host# edit interfaces sp-4/0/0 unit 2001
   user@host# set service-domain outside
   user@host# set family inet
   ```

3. Review the configuration.

   ```
   user@host> show configuration interfaces sp-4/0/0
   ```

   ```
   unit 1 {
       family inet;
       service-domain inside;
   }
   unit 2001 {
       family inet;
       service-domain outside;
   }
   ```

Configuring the Static Route for the cCPE Context

To configure the static route for the cCPE context:

1. Add a static route for the cCPE context to forward traffic destined for the Internet to the MS-DPC.

   ```
   [edit]
   user@host# set routing-instances vpn-1 routing-options static route 0.0.0.0/0
     sp-4/0/0.1
   ```

2. Review the configuration.

   ```
   user@host> show configuration routing-instances vpn-1 routing-options
   ```

```
static {
    route 0.0.0.0/0 next-hop sp-4/0/0.1;
}
```

Configuring Public Address Pools for NAT cCPE Services

To configure a public address pool for the NAT service:

1. Define the NAT source pool used to provide public addresses to the VPN subscriber.

   ```
   [edit]
   user@host# edit services nat pool pool-name
   ```

2. Specify the public addresses for the VPN subscriber.

   ```
   [edit services nat pool acme-site-abc-pool]
   user@host# set address subscriber-public-subnet
   ```

3. Configure the NAT pool port to be assigned automatically by the router.

   ```
   [edit services nat pool acme-site-abc-pool]
   user@host# set port automatic
   ```

Configuring NAT Rules for cCPE Services

This sample procedure configures a dynamic source rule for the NAT pool. This is just an example. You can also configure other types of rules, such as a dynamic destination NAT rule.

To configure a dynamic source rule for the NAT pool:

1. Specify the direction in which the rule match is applied.

   ```
   [edit]
   user@host# edit services nat rule subscriber-id-nat-rule-name
   ```

2. Specify the direction as input.

   ```
   [edit services nat rule acme-site-abc-nat-rule]
   user@host# set match-direction input
   ```

3. Define the properties for translated traffic.

   ```
   [edit services nat rule acme-site-abc-nat-rule]
   user@host# set term translate then translated source-pool acme-public-pool
   ```

4. Specify the translation type. We recommend setting the translation type to napt-44.

   ```
   [edit services nat rule acme-site-abc-nat-rule]
   user@host# set term translate then translated translation-type napt-44
   ```

5. Specify the NAT rules or rule set included in the subscriber's service-set.

   ```
   [edit services nat rule acme-site-abc-nat-rule]
   user@host# up
   user@host# up
   [edit services]
   user@host# set service-set acme-services nat-rules acme-nat-rule
   ```

Configuring the Inbound and Outbound Logical Interfaces for NAT cCPE Services

To configure the inbound and outbound logical interfaces used by the next-hop service:

1. Specify the name and logical unit number of the service interface associated with the services applied inside and outside the network.

   ```
   [edit services]
   user@host# edit service-set acme-services next-hop-service
   ```

   a. Specify the information for the services applied inside network

      ```
      [edit services service-set acme-services next-hop-service]
      user@host# set inside-service-interface interface-name.inbound-unit-number
      ```

   b. Specify the information for the services applied inside network

      ```
      [edit services service-set acme-services next-hop-service]
      user@host# set outside-service-interface interface-name.outbound-unit-number
      ```

Related Documentation

- Understanding the cCPE Common Configuration on page 15
- Understanding How to Assign Public and Private Addresses to VPN Subscribers Using NAT cCPE Services on page 25
- Verifying NAT cCPE Services on page 30
- Configuring NAT cCPE Services to Assign Multiple IP Address Ranges (Primary and Secondary Addresses) on page 29

Configuring NAT cCPE Services to Assign Multiple IP Address Ranges (Primary and Secondary Addresses)

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

Then complete the following tasks to configure the NAT service to assign multiple IP address ranges for primary and secondary addresses:

1. Enable NAT. See "Configuring NAT on MX Series Routers for cCPE Services" on page 26.
2. Configure the service interfaces for NAT cCPE services. See "Configuring the Service Interfaces for NAT cCPE Services" on page 27.
3. Configure the default route for the cCPE context. See "Configuring the Static Route for the cCPE Context" on page 27.
4. Configure multiple public address pools for NAT cCPE services. See "Configuring Public Address Pools for NAT cCPE Services" on page 28.

5. Configure the NAT rules for cCPE services. See "Configuring NAT Rules for cCPE Services" on page 28.
6. Configure the inbound and outbound logical interfaces for NAT cCPE services. See "Configuring the Inbound and Outbound Logical Interfaces for NAT cCPE Services" on page 29.

Related Documentation

- Configuring NAT cCPE Services to Assign Primary IP Addresses on page 26
- Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI on page 16
- Understanding How to Assign Public and Private Addresses to VPN Subscribers Using NAT cCPE Services on page 25
- *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*
- Verifying NAT cCPE Services on page 30

Configuring Public Address Pools for NAT cCPE Services

To configure a public address pool for the NAT service:

1. Define the NAT source pool used to provide public addresses to the VPN subscriber.

   ```
   [edit]
   user@host# edit services nat pool pool-name
   ```

2. Specify the public addresses for the VPN subscriber.

   ```
   [edit services nat pool acme-site-abc-pool]
   user@host# set address subscriber-public-subnet
   ```

3. Configure the NAT pool port to be assigned automatically by the router.

   ```
   [edit services nat pool acme-site-abc-pool]
   user@host# set port automatic
   ```

Related Documentation

- Configuring NAT cCPE Services to Assign Primary IP Addresses on page 26
- *Configuring Pools of Addresses and Ports for Network Address Translation Overview*
- Understanding How to Assign Public and Private Addresses to VPN Subscribers Using NAT cCPE Services on page 25
- Configuring NAT cCPE Services to Assign Multiple IP Address Ranges (Primary and Secondary Addresses) on page 29
- Verifying NAT cCPE Services on page 30

Verifying NAT cCPE Services

Purpose View the configuration for NAT cCPE services.

Action

- Verify that the primary routing instance has a route for the NAT pool IP address with a next hop of the outside service interface:

user@host> **show route** ***subscriber-public-address/32***

- Verify that the subscriber's routing instance has a default route with a next hop of the inside service interface:

  user@host> **show route table** ***subscriber-id*** **0.0.0.0/0 exact**

- Verify NAT by pinging from the subscriber site into the cloud. Check the NAT pool:

  user@host> **show services nat pool**

Related Documentation

- Configuring NAT cCPE Services to Assign Primary IP Addresses on page 26
- Configuring NAT cCPE Services to Assign Multiple IP Address Ranges (Primary and Secondary Addresses) on page 29
- Understanding How to Assign Public and Private Addresses to VPN Subscribers Using NAT cCPE Services on page 25

CHAPTER 4

Configuring SNMP Monitoring Services for the cCPE Application

Understanding How to Monitor cCPE Services Using SNMP

When you migrate subscribers to cCPE services, you can provide them with read-only access to systems (SNMP MIB-2 System), interfaces (SNMP MIB-2 Interfaces), and some environment-monitoring MIBs. Because the Layer 3 CPE functionality moves from the CPE device to the MX Series router, only the interfaces MIB is relevant to cCPE services. However, you can also provide read-only access to the systems MIB, which enables subscribers to retrieve the location or contact information of the cCPE context. Traps for system-wide events for a specific subscriber are not relevant, because the router hosting the cCPE context is shared by many customers (one cCPE context per subscriber).

The SNMP service supports routing instances in the following ways:

- You can bind an SNMP community to a specific routing instance by using the **edit snmp community *community* routing-instance *routing-instance*** configuration statement.
- If an SNMP request is received on a nondefault routing instance and the routing instance is not bound to the specified community, the request is denied.
- SNMP requests received by a nondefault routing instance through one of its access links must use *community* in the format of ***routing-instance*@*community***, where the ***routing-instance*** must be one of the routing instances bound to the community. In addition, the routing instance used in *community* must be the routing instance where the request is received. Such requests return only MIB objects pertaining to the specified routing instance. For example, an SNMP **walk** command issued on the interfaces table (RFC 2863) with community "acme@public" returns only interfaces defined in the routing instance "acme." To learn more about the MIB objects that are segregated by routing instances, see *Support Classes for MIB Objects*.
- SNMP read requests received on the default routing instance can also use *community* in the format of ***routing-instance*@*community***. The Junos OS SNMP agent returns only MIB objects pertaining to the specified routing instance. If you want to provide MIB information through an NMS system that connects to the default routing instance, you

can use the format ***routing-instance@community*** to provide the necessary MIB object isolation for a subscriber's routing instance.

NOTE: For security purposes, after you configure cCPE services, subscribers cannot retrieve the global router configuration. The only information customers can retrieve pertains to their cCPE context, which includes the interfaces used for implementing the IP VPN services.

You can retrieve the status of VPN services by getting the usage counters of the interfaces.

Related Documentation

- Configuring SNMP CCPE Monitoring Services on MX Series Routers on page 34
- *Support Classes for MIB Objects*
- *Understanding SNMP Support for Routing Instances*
- *Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community*
- *Configuring Access Lists for SNMP Access over Routing Instances*

Configuring SNMP CCPE Monitoring Services on MX Series Routers

To configure SNMP monitoring services for the CCPE application:

1. Define the system MIB view.

   ```
   [edit]
   user@host# set snmp view system-mib-view oid 1.3.6.1.2.1.1
   ```

2. Define the interfaces MIB view. See,

   ```
   [edit snmp view system-mid-view oid 1.3.6.1.2.1.1]
   user@host# set include
   ```

3. Specify the community allowed to view the system MIB.

   ```
   [edit snmp view system-mid-view oid 1.3.6.1.2.1.1]
   user@host# up
   user@host# up
   user@host# edit snmp community acme
   user@host# set view system-mib-view
   ```

4. (Optional) Specify the community allowed to view the interfaces MIB.

   ```
   [edit]
   user@host# set view interfaces-mib-view
   ```

5. Set the access authorization to read-only.

   ```
   [edit]
   user@host# set authorization read-only
   ```

6. Specify the customers routing instance and IP address of the SNMP client hosts that are authorized to use this community. See, *routing-instance* for SNMPv1 and SNMPv2 or *routing-instance* for SNMPv3.

```
[edit]
user@host# set routing-instance subscriber-id clients address
```

Related Documentation

- *Configuring SNMP on a Device Running Junos OS*
- *Understanding the SNMP Implementation in Junos OS*
- *Standard SNMP MIBs Supported by Junos OS*
- *Juniper Networks Enterprise-Specific MIBs*
- *Configuring SNMP Communities*
- *Examples: Configuring the SNMP Community String*
- *Configuring the Trap Target Address*

CHAPTER 5

Configuring CoS Services for the cCPE Application

Understanding How to Prioritize Subscriber Traffic Using Class of Service with cCPE Services

To prioritize subscriber traffic, you can use the CoS features of the MX Series router along with cCPE services. cCPE services require only a Layer 2 CPE device at the subscriber site. If the subscriber is currently using a Layer 3 CPE device, replace it with a Layer 2 CPE that supports CoS, such as the Juniper Networks EX Ethernet switch. To ensure the proper classification and management of upstream traffic in a subscriber's access link, CoS should be supported from CPE to CPE across the provider network. A Layer 2 CPE that supports CoS can classify traffic based on Layer 1 or Layer 2 information, such as physical port, VLAN, or MAC address. It can subsequently mark the traffic with Ethernet 802.1p, then prioritize the traffic through the CPE before passing it to the access link. If there are aggregation switches in the network path located before the provider edge (PE) router, the switches must also support Layer 2 CoS based on 802.1p bits in the traffic.

When the traffic reaches the PE router, the router uses CoS behavior aggregate (BA) classifiers to map 802.1p bits in the traffic to different forwarding classes and loss priorities. The traffic is prioritized to pass through the PE router.

You configure rewrite rules to change CoS values in outgoing packets on the outbound interfaces of the PE router to meet the policies of the core routers. If the remote Layer 2 CPE requires it, the remote PE router can use rewrite rules to change CoS values back to 802.1p on the outbound interface. By default, the MX Series router uses the BA classifier for 802.1p. For more information see, http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/general/default-802-1p-table-cos-config-guide.html.

If the Layer 2 CPE does not support CoS marking, the MX Series router can classify incoming packets by using a multifield classifier. The multifield classifier is a firewall filter

that classifies packets to a forwarding class by examining fields such as source and destination address, port, and IP protocol. CoS rewrite rules are applied to packets on outbound interfaces to set or change CoS markings to meet the policies of the core network.

Related Documentation

- Configuring CoS BA Classifiers and Rewrite Rules to Manage Traffic When Running cCPE Services on page 38
- Configuring CoS to Manage VoIP Traffic When Running cCPE Services on page 39

Configuring CoS BA Classifiers and Rewrite Rules to Manage Traffic When Running cCPE Services

This topic describes how to configure the behavior aggregate (BA) classifier and rewrite rules to manage traffic when you use cCPE services.

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

To configure the BA classifier and rewrite rules:

1. Configure the IEEE 802.1 default classifier to handle incoming packets.

   ```
   [edit]
   user@host# edit class-of-service interfaces interface-name unit vlan-unit-id classifiers
     ieee-802.1
   ```

   For example:

   ```
   [edit]
   user@host# edit class-of-service interfaces ge-1/0/1 unit 3 classifiers ieee-802.1
     default
   user@host# set default
   ```

2. Configure rewrite rules to map traffic to code points when the traffic exits the system, and apply the rewrite rules to a specific interface.

   ```
   [edit class-of-service interfaces ge-1/0/1 unit 105 classifiers ieee-802.1]
   user@host# up
   user@host# up
   user@host# edit rewrite-rules ieee-802.1
   user@host# set default
   ```

3. Review the configuration.

   ```
   user@host> show class-of-service interface ge-1/01

   Physical interface: ge-1/0/1, Index: 192
   Queues supported: 8, Queues in use: 6
     Scheduler map: <default>, Index: 2
     Congestion-notification: Disabled

     Logical interface: ge-1/0/1.0, Index: 65543
       Object                  Name                    Type
   Index
   ```

```
    Classifier              ipprec-compatibility    ip
13
```

```
user@host> show class-of-service classifier name ieee8021p-default

Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
  Code point         Forwarding class                  Loss priority
  000                best-effort                       low
  001                best-effort                       high
  010                expedited-forwarding              low
  011                expedited-forwarding              high
  100                assured-forwarding                low
  101                assured-forwarding                high
  110                network-control                   low
  111                network-control                   high
```

Related Documentation

- Understanding How to Prioritize Subscriber Traffic Using Class of Service with cCPE Services on page 37
- Verifying CoS cCPE Services on page 41
- *CoS Overview*
- *Assigning CoS Components to Interfaces*
- *Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces*
- *Defining CoS Rewrite Rules*
- *Understanding CoS Rewrite Rules*

Configuring CoS to Manage VoIP Traffic When Running cCPE Services

This procedure configures the CoS multifield classifier to classify incoming voice over IP (VoIP) traffic to the corresponding CoS forwarding class. A firewall filter classifies incoming Session Initiation Protocol (SIP) and Real-Time Transport Protocol (RTP) VoIP traffic.

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

Complete the following tasks to configure the multifield classifier:

1. Configuring a Firewall Filter to Classify SIP VoIP Traffic on page 39
2. Configuring a Firewall Filter to Classify RTP Traffic for CoS cCPE Services on page 40
3. Applying the SIP and RTP Firewall Filters for CoS cCPE Services on page 41

Configuring a Firewall Filter to Classify SIP VoIP Traffic

To configure an IPv4 firewall filter to classify session initiation protocol (SIP) VoIP traffic to a named forwarding class:

1. Configure an IPv4 firewall filter that checks for packets from the UDP protocol.

   [edit]

```
user@host# edit firewall family inet filter voip-filter term sip from
user@host# set protocol udp
```

2. Configure the same IPv4 firewall filter to check for packets from the TCP protocol.

```
[edit firewall family inet filter voip-filter term sip from]
user@host# set from protocol tcp
```

3. Specify which port is used for UDP and TCP traffic.

```
[edit firewall family inet filter voip-filter term sip from]
user@host# set port 5060
```

4. Classify the UDP and TCP packets to the named forwarding class.

```
[edit firewall family inet filter voip-filter term sip from]
user@host# up
[edit firewall family inet filter voip-filter term sip]
user@host# set then forwarding-class cos-voice
```

5. Specify that packets matching the filter conditions are accepted.

```
[edit firewall family inet filter voip-filter term sip]
user@host# set then accept
```

6. Verify the configuration.

```
user@host# show
from {
    protocol [ udp tcp ];
    port 5060;
}
then {
    forwarding-class cos-voice;
    accept;
}
```

Configuring a Firewall Filter to Classify RTP Traffic for CoS cCPE Services

This procedure configures an IPv4 firewall filter to classify Real-Time Transport Protocol (RTP) VoIP traffic to a named forwarding class.

1. Configure an IPv4 firewall filter that checks for packets from the UDP protocol.

```
[edit firewall family inet filter voip-filter term sip]
user@host# up
[edit firewall family inet filter voip-filter]
user@host# edit term rtp from
user@host# set protocol udp
```

2. Specify which port is used for UDP and TCP traffic.

```
[edit firewall family inet filter voip-filter term rtp from]
user@host# set port 16384-32767
```

3. Classify the UDP and TCP packets to the named forwarding class. Specify the same forwarding class you configured for SIP traffic.

```
[edit firewall family inet filter voip-filter term rtp from]
user@host# up
[edit firewall family inet filter voip-filter term rtp]
```

```
user@host# set then forwarding-class cos-voice
```

4. Specify that packets matching the filter conditions are accepted.

```
[edit firewall family inet filter voip-filter term rtp]
user@host# set then accept
```

5. Verify the configuration.

```
user@host# show
```

```
from {
    protocol udp;
    port 16384-32767;
}
then {
    forwarding-class cos-voice;
    accept;
}
```

Applying the SIP and RTP Firewall Filters for CoS cCPE Services

This procedure applies the firewall filter for SIP and RTP VoIP traffic to the subscriber's IRB interface.

- To apply the firewall filter, specify the logical interface *unit-number* designated for the subscriber's IRB interface, and specify the filter name you defined for SIP and RTP.

```
[edit]
user@host# edit interfaces irb unit 33 family inet filter input
user@host# set voip-filter
```

Related Documentation

- Understanding How to Prioritize Subscriber Traffic Using Class of Service with cCPE Services on page 37
- *Multifield Classifier Overview*
- *Guidelines for Configuring Firewall Filters*
- *Configuring Multifield Classifiers*
- *Understanding How to Use Firewall Filters*
- *Firewall Filter Nonterminating Actions*

Verifying CoS cCPE Services

Purpose View the information for CoS BA classifiers and rewrite rules for cCPE services.

Action

- To display the logical and physical interface associations for the classifier and rewrite rules for cCPE services:

```
user@host> show class-of-service interface interface-name unit vlan-unit-id
```

- To display the classifier mapping of code point value to forwarding class and loss priority:

```
user@host> show class-of-service classifier name
```

Related Documentation

- Understanding How to Prioritize Subscriber Traffic Using Class of Service with cCPE Services on page 37
- Configuring CoS BA Classifiers and Rewrite Rules to Manage Traffic When Running cCPE Services on page 38
- Configuring CoS to Manage VoIP Traffic When Running cCPE Services on page 39

CHAPTER 6

Configuring Jitter Measurement Services for the cCPE Application

Understanding Jitter Measurement and cCPE Services

MX Series routers support both Layer 2 and Layer 3 jitter measurement. Layer 2 jitter measurement is provided through Ethernet frame delay measurements (referred to as ETH-DM in Ethernet specifications), which is part of Ethernet OAM. The ETH-DM session sends an Ethernet frame delay measurement protocol data unit (PDU) to the remote maintenance association end point (MEP) to measure one-way or two-way frame delays and delay variations. Layer 3 jitter measurement is provided through the real-time performance monitoring (RPM) service, which sends probes (such as ICMP-ping, UDP-ping) to a remote node to measure round-trip time (RTT) and jitter between the two nodes.

NOTE: For cCPE services, jitter is measured between the cCPE contexts.

Layer 2 Jitter Measurement — Ethernet Frame Delay Measurement

Ethernet frame delay measurement is introduced in ITU-T Y.1731 on top of the IEEE 802.1ag Ethernet connectivity fault management (CFM). CFM is an end-to-end per service instance (per VLAN) Ethernet Layer OAM protocol. Two peers of a CFM session do not need to be adjacent. CFM partitions the service network into various administrative domains. For example, operators, providers, and customers may be part of different administrative domains. Each administrative domain is mapped into one maintenance domain, which provides enough information to perform its own management, thus avoiding security breaches and making end-to-end monitoring possible. Each maintenance domain is associated with a maintenance domain level from 0 through 7. Level allocation is based on the network hierarchy, where the outermost domains are assigned a higher level than the innermost domains. Subscriber endpoints have the highest maintenance domain levels (from 5 through 7). In a CFM domain, each service instance is called a *maintenance association*. A maintenance association can be thought of as a full mesh

of MEPs having similar characteristics. MEPs are active CFM entities generating and responding to CFM protocol messages. There is also a maintenance intermediate point (MIP), which is a CFM entity similar to the MEP, but more passive (MIPs respond only to CFM messages).

MEPs in the same maintenance association of the same maintenance domain exchange a continuity check message (CCM), which is a heartbeat message, to maintain the connectivity state with peers. Because CCM is a multicast Ethernet frame, it can pass through switches between two MEPs.

ETH-DM uses CFM as an infrastructure. It measures frame delay and frame delay variations (jitter) between two CFM MEPs, which must be in the same maintenance domain and maintenance association. You can start an ETH-DM in either one-way or two-way (round-trip) mode to gather frame delay statistics. You can also specify frame priority (802.1p) in an ETH-DM session.

Layer 3 Jitter Measurement — Real-Time Performance Monitoring

MX Series routers provide an RPM service that sends probes (such as ICMP-ping, UDP-ping) to a remote node (router) to measure RTT and jitter between the two nodes. The probe packet is timestamped when it exits and when it returns to the source. If the destination is another Juniper Networks router with the RPM service configured, the probe packet is timestamped when it is received at the destination and when the response exits the router. You can configure the probe's DiffServ code point (DSCP) bits so that RPM can measure performance of traffic being managed by a specific CoS configuration.

Because RPM is performed between two PE routers, it is not specific to cCPE services and is not discussed here in detail. For more information about RPM, see *Real-Time Performance Monitoring Services*.

Related Documentation

- Example: Configuring and Running Layer 2 Jitter Measurements with ETH-DM with cCPE Services on page 44
- *Ethernet Frame Delay Measurements*
- *Triggering an Ethernet Frame Delay Measurements Session*
- *Example: Configuring One-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces*
- *Example: Configuring Two-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces*

Example: Configuring and Running Layer 2 Jitter Measurements with ETH-DM with cCPE Services

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

Requirements

In this example, we use the following Juniper Networks hardware and software components:

- MX Series 3D Universal Edge Router for the PE router
- EX Series Ethernet switch for the Layer 2 CPE
- EX Series Ethernet switch for the aggregation switch
- Junos OS Release 13.2 or later

Overview and Topology

This example shows you how to configure and perform Layer 2 jitter measurement with Ethernet frame delay measurement (ETH-DM) when you are using cCPE services. The configuration simulates an access network in which a Layer 2 CPE connects to a PE router through an aggregation switch. Complete the following high-level steps for this test configuration.

On the PE router:

1. Configure the Layer 2 interface, IRB, bridge domain, and VPN routing instance.
2. Configure the maintenance domain and maintenance association under the **[edit protocols oam]** hierarchy.
3. Configure a maintenance end point (MEP) under the maintenance association, and bind the Layer 2 interface to it.

On the Layer 2 CPE:

1. Configure Layer 2 interface and VLAN.
2. Configure the same maintenance domain and maintenance association as those you specified for the PE router.
3. Configure a maintenance end point (MEP) under the same maintenance association you configured for the PE router, and bind the Layer 2 interface to it.

On the aggregation switch:

1. Configure a VLAN with two trunk ports to connect the PE router and Layer 2 CPE.

Figure 4 on page 46 shows the example network setup for the Layer 2 jitter measurement:

Figure 4: Jitter Measurement with Ethernet Frame Delay Measurement

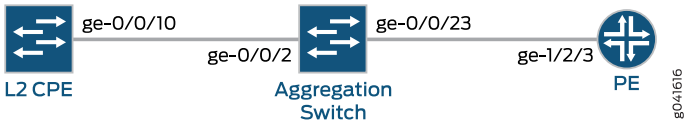

Two maintenance end points (MEPs), one on interface ge-0/0/10 of the Layer 2 CPE and one on interface ge-1/2/3 on the PE router, are defined under the same maintenance association "acme" of maintenance domain "Example-subscriber-md." You can initiate ETH-DM from either end to measure the delay and jitter.

To initiate the ETH-DM session run the command:

user@host> **monitor ethernet delay-measurement** *one-way|two-way* **maintenance-domain** *domain-name* **maintenance-association** *ma-name* **mep** *remote-mep-id*

You can initiate the ETH-DM scan from either end to measure delay and jitter.

Configuring the MX Series Router (PE Router)

The following procedures explain how to configure the MX Series router so that you can perform Layer 2 jitter measurement with ETH-DM.

- Configuring the Layer 2 Interface, IRB, Bridge Domain, and VPN Routing Instance on page 46
- Configuring the Maintenance Domain and Maintenance Association on page 47
- Configuring the MEP Under the Maintenance Association on page 47

Configuring the Layer 2 Interface, IRB, Bridge Domain, and VPN Routing Instance

Step-by-Step Procedure

To configure the Layer 2 interface, IRB interface, bridge domains, and VPN routing instances:

1. Configure the Layer 2 interface and IRB.

```
[edit]
user@host# set interfaces ge-1/2/3 vlan-tagging
user@host# set interfaces ge-1/2/3 encapsulation flexible-ethernet-services
user@host# set interfaces ge-1/2/3 unit 106 encapsulation vlan-bridge
user@host# set interfaces ge-1/2/3 unit 106 vlan-id 106
user@host# set interfaces ge-1/2/3 unit 106 family bridge
user@host# set interfaces irb unit 106 family inet address 10.132.11.1/24
```

2. Configure the bridge domains.

```
[edit]
user@host# set bridge-domains vlan-106 domain-type bridge
user@host# set bridge-domains vlan-106 vlan-id 106
user@host# set bridge-domains vlan-106 interface ge-1/2/3.106
user@host# set bridge-domains vlan-106 routing-interface irb.106
```

3. Configure the VPN routing instances.

```
[edit]
user@host# set routing-instances acme instance-type vrf
user@host# set routing-instances acme interface irb.106
```

Configuring the Maintenance Domain and Maintenance Association

Step-by-Step Procedure

To configure the maintenance domain and maintenance association:

1. Enable hardware timestamping to ensure accurate measurement.

```
[edit]
user@host# set protocols oam ethernet connectivity-fault-management
  performance-monitoring hardware-assisted-timestamping
```

2. Create the maintenance domain and maintenance association.

```
[edit]
user@host# set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md level 3
```

3. Enable the continuity check (CC). The CC interval and hold interval must be set to the same value on all MEPs in the same maintenance association.

```
[edit]
user@host# set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md maintenance-association acme
  continuity-check interval 100ms
user@host# set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md maintenance-association acme
  continuity-check hold-interval 1
```

Configuring the MEP Under the Maintenance Association

Step-by-Step Procedure

To configure the MEP:

- Configure the MEP under the maintenance association, and bind the Layer 2 interface to it.

```
[edit]
user@host# set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md maintenance-association acme
  mep 201 interface ge-1/2/3.106
user@host# set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md maintenance-association acme
  mep 201 direction down
user@host# set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md maintenance-association acme
  mep 201 auto-discovery
```

Configuring the EX Series Ethernet Switch (Layer 2 CPE)

The following procedures explain how to configure the EX Series Ethernet switch as a Layer 2 CPE device and perform Layer 2 jitter measurement with ETH-DM.

- Configuring the Layer 2 Interface and VLANs on page 48
- Creating the Maintenance Domain and Maintenance Association on page 48
- Creating the MEP on page 48

Configuring the Layer 2 Interface and VLANs

Step-by-Step Procedure

To configure the Layer 2 interface and VLANs:

- Configure the Layer 2 interface and VLANs.

```
[edit]
user@host# set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode
  trunk
user@host# set interfaces ge-0/0/9 unit 0 family ethernet-switching port-mode
  access
user@host# set vlans vlan-106 vlan-id 106
user@host# set vlans vlan-106 interface ge-0/0/10.0
user@host# set vlans vlan-106 interface ge-0/0/9.0
```

Creating the Maintenance Domain and Maintenance Association

Step-by-Step Procedure

1. Configure the same maintenance domain and maintenance association as those you specified for the PE router.

```
[edit]
user@host# set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md level 3
```

2. Configure the continuity check.

```
[edit]
user@host# set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md maintenance-association acme
  continuity-check interval 100ms
user@host# set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md maintenance-association acme
  continuity-check hold-interval 1
```

Creating the MEP

Step-by-Step Procedure

Create the MEP and assign the interface to it.

- Configure a MEP under the same maintenance association you configured for the PE router, and bind the Layer 2 interface to it.

```
[edit]
user@host# set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md maintenance-association acme
  mep 101 interface ge-0/0/10.0
user@host# set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md maintenance-association acme
  mep 101 interface vlan-id 106
user@host#  set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md maintenance-association acme
  mep 101 direction down
user@host#  set protocols oam ethernet connectivity-fault-management
  maintenance-domain Example-subscriber-md maintenance-association acme
  mep 101 auto-discovery
```

Configuring the Aggregation Switch

The following procedures describe how to configure the EX Series Ethernet switch as an aggregation switch and perform Layer 2 jitter measurement with ETH-DM.

- Configuring the VLAN to Connect the PE Router and Layer 2 CPE on page 49

Configuring the VLAN to Connect the PE Router and Layer 2 CPE

Step-by-Step Procedure

- Configure a VLAN with two trunk ports to connect the PE router and Layer 2 CPE.

```
[edit]
user@host# set interfaces ge-0/0/23 unit 0 family ethernet-switching port-mode
  trunk
user@host# set interfaces ge-0/0/2 unit 0 family ethernet-switching port-mode
  trunk
user@host# set vlans vlan-106 vlan-id 106
user@host# set vlans vlan-106 interface ge-0/0/2.0
user@host# set vlans vlan-106 interface ge-0/0/23.0
```

NOTE: When the EX Series switch is used as an aggregation switch, you need to disable IGMP snooping or configure a firewall filter to allow CCM multicast packets to pass through the switch.

Verification

Verify the Layer 2 jitter measurement configuration by displaying the connectivity fault management (CFM) sessions at the PE router or CPE by entering the **show oam ethernet connectivity-fault-management mep-database** command.

- Viewing CFM Sessions on page 49

Viewing CFM Sessions

Purpose View CFM sessions.

Action

- To view the CFM sessions, enter the following command at either the PE router or the switch:

```
user@host> show oam ethernet connectivity-fault-management mep-database
```

Related Documentation

- Understanding Jitter Measurement and cCPE Services on page 43

- *Ethernet Frame Delay Measurements*
- *Triggering an Ethernet Frame Delay Measurements Session*
- *Example: Configuring One-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces*

- *Example: Configuring Two-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces*

CHAPTER 7

Configuring DHCP Services for the cCPE Application

Understanding DHCP cCPE Services

Using cCPE services, you can provide DHCP server and DHCP relay services to your customers.

DHCP Server

You can use the MX Series router DHCP server function in a VPN routing instance to provide DHCP services to Layer 2 CPEs. To use this service, enable the DHCP server in at least one routing instance corresponding to the subscriber VPN site; for all other sites you can use the DHCP relay agent pointing to the same DHCP server. In addition, you need to configure the address pools within the VRF routing instance. The address pools you configure for the VPN site need to be part of the subscriber's private subnets configured for the IRB interface.

To protect the Routing Engine from DHCP flooding, you can define interface-specific firewalls.

NOTE: The address assignment pool feature is part of the Junos Subscriber Management Feature Pack license. You must install and properly configure the license to meet the requirements for using the address assignment pool feature. The license includes 1000 address bindings.

DHCP Relay Agent

The MX Series router extended DHCP relay agent forwards DHCP request-reply packets between a DHCP client and a central DHCP server. You configure the relay agent within

the cCPE context; in other words, within the VRF routing instance. Typically, the subscriber owns and operates the central DHCP server, and manages the address pools.

NOTE: The stateless DHCP relay agent, which you can configure under the **[routing-instance *routing-instance* forwarding-options helpers bootp]** edit hierarchy, cannot coexist with the DHCP server in the same MX Series router, even if they are configured in different routing instances.

Related Documentation

- Configuring DHCP cCPE Services on page 52
- Configuring DHCP Relay Agent cCPE Services on page 54

Configuring DHCP cCPE Services

This topic describes how to configure DHCP cCPE services.

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

To configure the DHCP cCPE services, complete the following tasks:

1. Configuring the DHCP Server on the MX Series Router for the cCPE Context on page 52
2. Configuring Subscriber Address Pools for DHCP cCPE Services on page 53

Configuring the DHCP Server on the MX Series Router for the cCPE Context

Configure DHCP local server options on the PE router. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client. To configure the local DHCP server for the cCPE context:

1. Configure the DHCP server in each routing instance that corresponds to a subscriber VPN site.

   ```
   [edit]
   user@host# edit routing-instances acme system services dhcp-local-server group
     default-group interface irb.subscriber-vlan-unit-id
   ```

2. (Optional) When you configure a DHCP server in a VRF, if you want the server to handle DHCP requests sent by DHCP relay agents located at remote sites of the same VPN, you need to add tunnel interfaces to the DHCP server. Specifically, you need to add the following configuration to the DHCP server:

   ```
   user@host set routing-instances acme system services dhcp-local-server group
     default-group interface lsi.0 upto lsi.16385
   ```

Configuring Subscriber Address Pools for DHCP cCPE Services

Configure the address pools within the subscriber's VRF routing instance. The address pools that you configure for the VPN site need to be part of the subscriber's private subnets, configured for the IRB interface.

To configure the address pools:

1. Configure the address pools within the subscriber's routing instance.

   ```
   [edit]
   user@host# edit routing-instances subscriber-id access address-assignment pool
     acme-pool family inet
   ```

2. Configure the subnet information for the IPv4 address-assignment pool.

   ```
   [edit routing-instances acme access address-assignment pool acme-pool family inet]
   user@host# set network prefix/prefix-length
   ```

3. Configure a named range of IPv4 addresses, used within the address-assignment pool.

   a. Define the lower limit of the address range.

      ```
      [edit routing-instances acme access address-assignment pool acme-pool family
        inet]
      user@host# set range r1 low lower-limit
      ```

   b. Define the upper limit of the address range.

      ```
      [edit routing-instances acme access address-assignment pool acme-pool family
        inet]
      user@host# set range r1 high upper-limit
      ```

   c. Repeat the process for each subscriber and address range.

4. Configure address pools that can be used by different client applications. In the following example, we configure the DNS, WINS, and router servers, however, there are many other DHCP attributes that you can configure depending on your network requirements.

   a. Specify a DNS server to which clients can send DNS queries.

      ```
      [edit routing-instances acme access address-assignment pool acme-pool family
        inet]
      user@host# edit dhcp-attributes
      [edit routing-instances acme access address-assignment pool acme-pool family
        inet dhcp-attributes]
      user@host# set name-server dns-server-ip-address
      ```

   b. Specify one or more NetBIOS name servers that the client uses to resolve NetBIOS names.

      ```
      [edit routing-instances acme access address-assignment pool acme-pool family
        inet dhcp-attributes]
      user@host# set wins-server wins-server-ip-address
      ```

   c. Specify one or more routers located on the client's subnet.

[edit routing-instances acme access address-assignment pool acme-pool family inet dhcp-attributes]
user@host# set router *router-ip-address*

5. Review the configuration.

[edit routing-instances acme access address-assignment pool acme-pool family inet]
user@host# show

```
pool acme-pool {
    family inet {
        network 10.10.1.1/32;
        range r1 {
            low 10.10.0.1;
            high 10.10.10.10;
        }
        dhcp-attributes {
            name-server {
                10.10.1.2;
            }
            wins-server {
                10.10.1.3;
            }
            router {
                10.10.1.4;
            }
        }
    }
}
```

Related Documentation

- Understanding DHCP cCPE Services on page 51
- Verifying and Managing DHCP Local Server Configuration on page 56
- *Extended DHCP Local Server Overview*
- *Address-Assignment Pools Overview*
- *Address-Assignment Pools Licensing Requirements*

Configuring DHCP Relay Agent cCPE Services

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

NOTE: The stateless DHCP relay agent, which you configure under the **[routing-instance *routing-instance* forwarding-options helpers bootp]** edit hierarchy, cannot coexist with the DHCP local server in the same MX Series router, even if they are configured in different routing instances.

NOTE: When you configure a DHCP server in a VRF, if you want the server to handle DHCP requests sent by DHCP relay agents located at remote sites of the same VPN, you need to add tunnel interfaces to the DHCP server. Specifically, you need to add the following configuration to the DHCP server:

```
user@host set routing-instances acme system services dhcp-local-server
    group default-group interface lsi.0 upto lsi.16385
```

For more information about configuring the DHCP server, see "Configuring DHCP cCPE Services" on page 52.

Configure one server group and one client group for each subscriber. Define a reference to the server group with the group stanza. To configure the DHCP relay agent:

1. Configure the server group by specifying the name of the group and DHCP server addresses for use by the extended DHCP relay agent.

   ```
   [edit]
   user@host# edit routing-instances acme forwarding-options dhcp-relay
   user@host# set server-group server-group-name server-ip-address
   ```

2. (Optional) If the subscriber is using a DHCP server farm, you can add another server to the group.

   ```
   [edit routing-instances acme forwarding-options dhcp-relay]
   user@host# set server-group server-group-name server-ip-address
   ```

3. (Optional) If the subscriber has more than one server group, configure an active server group to apply a common DHCP relay agent configuration to a named group of DHCP server addresses.

   ```
   [edit routing-instances acme forwarding-options dhcp-relay]
   user@host# set active-server-group server-group-name
   ```

4. (Optional) Specify the interface to be served by the DHCP relay agent. Typically, this is not required, because the IRB interface is the single interface for the cCPE context.

   ```
   [edit routing-instances acme forwarding-options dhcp-relay]
   user@host# set group group-name interface irb.vlan-unit-id
   ```

5. Override the default configuration settings for the extended DHCP relay agent, and enable DHCP snooping.

   ```
   [edit routing-instances acme forwarding-options dhcp-relay]
   user@host# set overrides allow-snooped-clients
   ```

6. Review the configuration.

   ```
   [edit routing-instances acme forwarding-options dhcp-relay]
   user@host# show
   ```

```
overrides {
    allow-snooped-clients;
}
server-group {
    acme-server-group {
        10.1.2.3;
        10.2.3.4;
```

```
    }
}
    active-server-group {
        10.1.2.3;
    }
active-server-group acme-server-group;
group acme-group {
    interface irb.105;
}
```

Related Documentation

- Understanding How to Use VRRP and DHCP with CCPE on page 62
- Configuring VRRP with DHCP cCPE Services on page 70
- Configuring a VRRP IPv4 Group for cCPE Services on page 72
- Understanding DHCP cCPE Services on page 51
- Verifying and Managing DHCP Relay Configuration on page 56

Verifying and Managing DHCP Local Server Configuration

Purpose View or clear information about client address bindings and statistics for the extended DHCP local server.

NOTE: If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the **clear dhcp server binding** command before you delete the DHCP server configuration.

Action

- To display the address bindings in the client table on the extended DHCP local server:

  user@host> **show dhcp server binding routing-instance *customer routing instance***

- To display extended DHCP local server statistics:

  user@host> **show dhcp server statistics routing-instance *customer routing instance***

- To clear the binding state of a DHCP client from the client table on the extended DHCP local server:

  user@host> **clear dhcp server binding routing-instance *customer routing instance***

- To clear all extended DHCP local server statistics:

  user@host> **clear dhcp server statistics routing-instance *customer routing instance***

Related Documentation

- CLI Explorer

Verifying and Managing DHCP Relay Configuration

Purpose View or clear address bindings or statistics for extended DHCP relay agent clients:

Action

- To display the address bindings for extended DHCP relay agent clients:

user@host> **show dhcp relay binding routing-instance** ***customer routing instance***

- To display extended DHCP relay agent statistics:

  user@host> **show dhcp relay statistics routing-instance** ***customer routing instance***

- To clear the binding state of DHCP relay agent clients:

  user@host> **clear dhcp relay binding routing-instance** ***customer routing instance***

- To clear all extended DHCP relay agent statistics:

  user@host> **clear dhcp relay statistics routing-instance** ***customer routing instance***

Related Documentation

- CLI Explorer

CHAPTER 8

Configuring VRRP Services for the cCPE Application

Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links

Traditionally, the Virtual Router Redundancy Protocol (VRRP) is configured between two Layer 3 CPEs to provide redundancy. VRRP monitors the access link state and triggers a redundancy switch when an access link becomes unavailable. When you transition the subscriber to CCPE services, you replace the Layer 3 CPE with a Layer 2 CPE, and VRRP is not applicable to the CPE. Instead, you run VRRP between cloud CPEs, which exist in two MX Series PE routers. VRRP is configured on the IRB interface of cloud CPE context

and monitors the state of the access links to the CCPE subscriber. A virtual IP address, which is the gateway address of the CCPE LAN, is maintained by VRRP. When one link is down, VRRP switches the virtual address to the other CCPE to maintain connectivity.

Figure 5: Cloud CPE VRRP

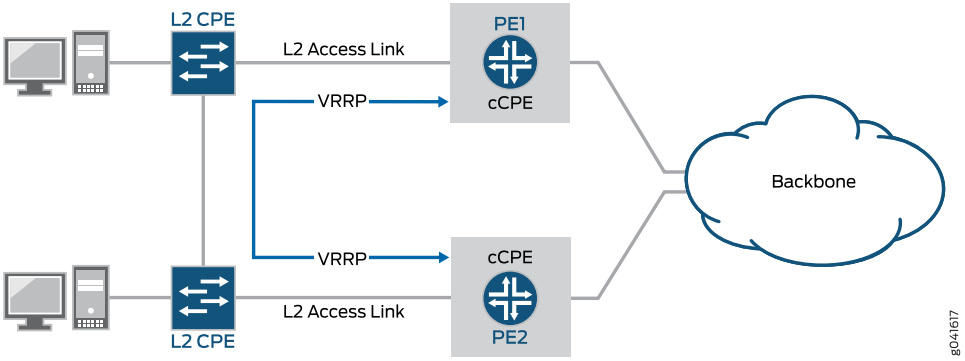

Figure 5 on page 60 shows the scenario of running VRRP between two Layer 2 CPEs running CCPE services in two MX Series PE routers. Each VRRP instance monitors a Layer 2 access link connected to the CCPE.

The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the primary (active) and the other is the backup. If the primary fails, the backup PE router becomes the new active router, providing a virtual default routing platform and allowing traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup router can take over a failed default router within a few seconds. This is done with minimum VRRP traffic and without any interaction with the hosts.

VRRP instances use multicast to communicate with each other. The two cCPEs are interconnected through two access links and two Layer 2 CPEs so VRRP messages can be exchanged using this path. If the link between the two Layer 2 CPEs is disconnected, it creates a split-brain situation so both VRRP instances assume the primary role. In this scenario, virtual IP address exist on both cCPEs.

If there is no aggregation switch between the Layer 2 CPE and PE router, VRRP can monitor the link state of the access link and initiate a redundancy switch if the primary access link goes down.

However, if there is an aggregation switch between the Layer 2 CPE and PE router, the link state is not sufficient to detect connectivity issues. In this case, you can use dynamic routing between the PE router and subscriber's network, which allows VRRP to monitor access link connectivity by monitoring the routing protocol. If there is no dynamic routing but the Layer 2 CPE supports 802.1ag (Ethernet-oam Connectivity Fault Management), you can use CFM in the MX Series PE router to monitor connectivity between the router and Layer 2 CPE. An action profile, for example that shuts down the subscriber facing Layer 2 interface, can be configured under CFM in case of connectivity issues. When the CFM session detects an issue and invokes the action profile to shutdown the subscriber facing Layer 2 interface in the PE router, it triggers VRRP switching redundancy.

Running Multiple VRRP Groups in Multiple Subnets with CCPE (Load Sharing)

NOTE: The MX Series router supports 1024 VRRP instances maximum per chassis.

In Figure 5 on page 60, there is only one virtual address for the CCPE subscriber. At anytime, only one access link is used because traffic from the subscriber only flows through the access link connected to the active/primary VRRP at the time. To support load sharing, you need to split subscriber LAN traffic into two subnets by configuring one VRRP group and one virtual IP address for each subnet. One CCPE has a higher priority than the other CCPE in one VRRP group but lower priority in the other VRRP group. So, when both subscriber access links are available, one virtual IP address exists in one CCPE and the other virtual IP address is in the other CCPE. Traffic from hosts in one subnet flows through the access link connected to the CCPE with the virtual IP address, which is the default gateway for the subnet, so both links are used. When one access link is down, both virtual addresses are moved to the same CCPE so connectivity is maintained.

Running Multiple VRRP Groups in a Single Subnet in the CCPE Application (Load Sharing)

It is possible to configure two VRRP groups in a single subnet to achieve load sharing. Each VRRP group has one virtual address. Hosts on the subscriber network are split into two groups: each group uses one virtual address as the default gateway. Similar to Figure 5 on page 60, you split the primary devices of the VRRP groups into two PE routers to divide outgoing traffic to different access links. For return traffic or incoming traffic, there are two routes through the two PE routers in the core network. By default, when there are multiple equal-cost paths to the same destination for the active route, the Junos OS uses a hash algorithm to choose one of the next-hop addresses to install in the forwarding table. You can also configure the Junos OS to do per-packet or per-flow load balancing across multiple paths.

VRRP with Ethernet-OAM Monitoring the Subscriber's Access Link

If there are aggregation switches between the Layer 2 CPEs and PE routers, link state is not sufficient to detect connectivity issues. If the Layer 2 CPE supports 802.1ag Ethernet-oam Connectivity Fault Management (CFM), you can use Ethernet CFM in the MX Series PE router to monitor connectivity between the router and the Layer 2 CPE. The Ethernet CFM session between the PE router and the Layer 2 CPE exchange heartbeat messages to monitor the connectivity to the remote peer. Junos OS allows you to configure an action profile in CFM, which is invoked when the CFM session detects a connectivity issue. One of the actions you can invoke in an action profile is "interface-down", which brings down the logical interface the CFM session is running on. You can use the "interface-down" action in a CFM session in the PE router to shutdown the subscriber facing VLAN interface when CFM detects a connection loss to the remote Layer 2 CPE. This will further trigger a VRRP redundancy switch. After the connection issue is resolved, CFM automatically brings up the interface and depending on your VRRP configuration, redundancy can be switched back again.

Related Documentation

- Understanding How to Use VRRP and DHCP with CCPE on page 62
- *Configuring the Virtual Router Redundancy Protocol (VRRP) with CCPE Services*
- *Configuring VRRP CCPE Services with Multiple VRRP Groups and Subnets (Load Sharing)*
- *Configuring VRRP CCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing)*
- *Configuring VRRP with Ethernet-OAM CCPE Services*
- *Configuring VRRP with DHCP CCPE Services*
- *Understanding High Availability Features on Juniper Networks Routers*
- *Junos OS Support for VRRPv3*
- *Configuring VRRP Authentication (IPv4 Only)*
- *Configuring Basic VRRP Support*
- *Understanding VRRP*
- *Example: Configuring VRRP*
- *Configuring a Logical Interface to Be Tracked*

Understanding How to Use VRRP and DHCP with CCPE

For customers running VRRP between two CCPE instances for dual access and who also want to use DHCP to allocate addresses for their devices, you must determine the best place to run the DHCP server.

The first option is running two DHCP servers, one in each CCPE instance, and splitting the address pool between the two servers. In this scenario, both DHCP servers are active and devices can receive an IP address from either server. Both DHCP servers configure the VRRP virtual address as the router address (DHCP option 3), so that subscriber traffic travels through the active VRRP interface. When the active access link goes down, the DHCP server in the corresponding MX Series PE router is unreachable to user devices. New DHCP leases are provided by the live DHCP server, however, DHCP renew requests for the leases managed by the unavailable DHCP server will fail. Devices requesting renewals receive new addresses from the live server. Existing TCP connections to or from the device are broken and the application may need to be restarted to recover.

The second option, shown in Figure 6 on page 63, is to run stateless DHCP relay agents in each CCPE instance. Both DHCP relay agents must point to the same DHCP server. When a device sends a DHCP discovery request, the request is forwarded to the DHCP server by both relay agents. The DHCP server may send two offers back to the client through the relay agents, however, the client will select an offer and send a request for it. The DHCP server sends an acknowledgement back to the client and the other offer expires. DHCP renew requests are sent to the sever using unicast. When one of the CCPE instances is unavailable due to the corresponding access link being down, DHCP discover requests are sent to the DHCP server by the other relay agent without interruption. Existing DHCP leases are renewed successfully because renew requests are sent using unicast to the DHCP server without relay agent involvement.

NOTE: The stateless DHCP relay agent, which you can configure under the **[routing-instance *routing-instance* forwarding-options helpers bootp]** edit hierarchy, cannot coexist with the DHCP server in the same MX Series router, even if they are configured in different routing instances.

You must configure the DHCP relay agent by using the **[routing-instance routing-instance-name forwarding-options dhcp-relay]** configuration statement. However, when you use this relay agent, it is possible that renew requests can be dropped if the router has not seen the original discovery request for the address.

Figure 6: VRRP with DHCP Relay

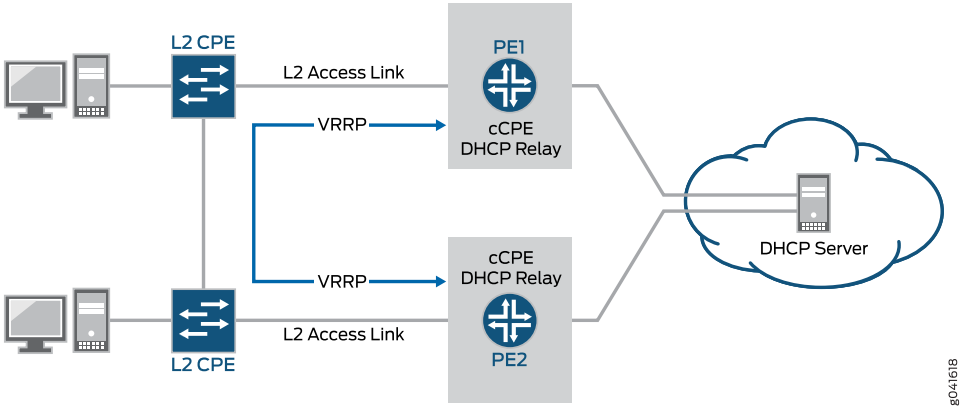

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- *Configuring the Virtual Router Redundancy Protocol (VRRP) with CCPE Services*
- *Configuring VRRP CCPE Services with Multiple VRRP Groups and Subnets (Load Sharing)*
- *Configuring VRRP CCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing)*
- *Configuring VRRP with Ethernet-OAM CCPE Services*
- *Configuring VRRP with DHCP CCPE Services*
- *Understanding High Availability Features on Juniper Networks Routers*
- *Junos OS Support for VRRPv3*
- *Configuring VRRP Authentication (IPv4 Only)*
- *Configuring Basic VRRP Support*
- *Understanding VRRP*
- *Example: Configuring VRRP*
- *Configuring a Logical Interface to Be Tracked*
- *Configuring Server Groups*

- *Extended DHCP Relay Agent Overview*
- *Configuring Active Server Groups*
- *Group-Specific DHCP Relay Options*
- *Overriding the Default DHCP Relay Configuration Settings*
- *Configuring DHCP Snooping for DHCP Relay Agent*
- *DHCP Snooping Support*
- *Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent*
- Verifying and Managing DHCP Relay Configuration on page 56
- *Example: Configuring DHCP Snooping Support for DHCP Relay Agent*
- *Example: Minimum DHCP Relay Agent Configuration*
- *Example: DHCP Relay Agent Configuration with Multiple Clients and Servers*
- *Tracing Extended DHCP Operations*
- *Tracing Extended DHCP Operations for Specific Interfaces*

Configuring VRRP with cCPE Services

This topic describes how to configure Virtual Router Redundancy Protocol (VRRP) with cCPE services.

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

Then complete the following tasks to configure the primary PE router for VRRP cCPE services:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.
3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the VRRP group. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72.
5. (Optional) Enable tracking of the CPE-facing Layer 2 interface. This step is optional if there is only one Layer 2 interface in the bridge domain. If the only Layer 2 interface is down, the bridge domain is down, including the IRB interface. This also brings down the VRRP instance. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.

Complete the following tasks to configure the secondary PE router for VRRP cCPE services:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.
3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72.

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Understanding How to Use VRRP and DHCP with CCPE on page 62
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65
- Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing) on page 67
- Configuring VRRP with Ethernet OAM cCPE Services on page 69
- Configuring VRRP with DHCP cCPE Services on page 70

Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing)

This procedure configures two VRRP groups in two subnets with load balancing. To support load sharing, we need to split the customer LAN into two subnets. There is one VRRP group and one virtual IP address for each subnet. One cCPE has a higher priority than the other cCPE in one VRRP group but a lower priority in the other VRRP group. So, when both access links are available, one virtual IP address exists in one cCPE instance and the other virtual IP address exists in the other cCPE instance. Traffic from hosts in one subnet flows through the access link connected to the cCPE instance with the virtual IP virtual address, which is the default gateway for the subnet. So both links are used. When one access link is down, both virtual addresses are moved to the same cCPE instance so that connectivity is maintained.

NOTE: A maximum of 1024 VRRP instances are supported per chassis.

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

Then complete the following tasks to configure the primary PE router:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.
3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the first VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72.
5. Configure tracking of the CPE-facing Layer 2 interface for the first VRRP group. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.
6. Configure a second VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. Be sure to set either a higher or lower priority as compared to the first VRRP group. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72.
7. Configure tracking of the second CPE-facing Layer 2 interface for the second VRRP group. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.

Complete the following tasks to configure the secondary PE router:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.
3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the first VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72
5. Configure tracking of the CPE-facing Layer 2 interface for the first VRRP. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.
6. Configure a second VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. Be sure to set either a higher or lower priority as compared to the first VRRP group. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72.
7. Configure tracking of the second CPE-facing Layer 2 interface for the second VRRP group. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.

Related Documentation

- Configuring VRRP with cCPE Services on page 64
- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Understanding How to Use VRRP and DHCP with CCPE on page 62
- Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing) on page 67

Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing)

You can configure two VRRP groups in a single subnet to achieve load balancing. Each VRRP group has one virtual address. Hosts on the subscriber network are split into two groups: Each group uses one virtual address as the default gateway. Similar to the configuration described in "Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing)" on page 65, we split masters of the VRRP groups into two PE routers to divide outgoing traffic to different access links. For return traffic or incoming traffic, there are two routes through two PE routers in the core network. By default, when there are multiple equal-cost paths to the same destination for the active route, the Junos OS uses a hash algorithm to choose one of the next-hop addresses to install in the forwarding table. You can also configure per-packet or per-flow load balancing across multiple paths.

This procedure configures two VRRP groups in a single subnet with load balancing.

NOTE: The MX Series router supports up to 1024 VRRP instances per chassis.

NOTE: One cCPE has a higher priority than the other cCPE in one VRRP group but lower priority in the other VRRP group. So when both access links are available, one virtual IP exists in one cCPE and the other virtual IP address is in the other cCPE.

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

Then complete the following tasks to configure the primary PE router:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.
3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the first VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72
5. Configure tracking of the CPE-facing Layer 2 interface for the first VRRP. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.

6. Configure a second VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. Be sure to set either a higher or lower priority as compared to the first VRRP group. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72.
7. Configure tracking of the second CPE-facing Layer 2 interface for the second VRRP group. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.

Complete the following tasks to configure the secondary PE router:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.
3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the first VRRP group. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72
5. Configure tracking of the CPE-facing Layer 2 interface for the first VRRP. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.
6. Configure a second VRRP group. Be sure to set either a higher or lower priority as compared to the first VRRP group. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72.
7. Configure tracking of the second CPE-facing Layer 2 interface for the second VRRP group. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Understanding How to Use VRRP and DHCP with CCPE on page 62
- Configuring VRRP with cCPE Services on page 64
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65
- Configuring VRRP with Ethernet OAM cCPE Services on page 69
- Configuring VRRP with DHCP cCPE Services on page 70

Configuring VRRP with Ethernet OAM cCPE Services

This procedure sets up CFM monitoring for the access link between the PE router and the CPE. You can have multiple aggregation switches between the router and the CPE device. To perform the access link monitoring, you need to configure CFM MEPs in the PE router and CPE so that they can exchange heartbeat messages to monitor connectivity. Intermediate switches do not need a CFM-specific configuration, because the heartbeat messages are just regular Ethernet multicast packets.

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

Then complete the following tasks to configure the primary PE router:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.
3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the first VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72.
5. Configure Ethernet OAM. See "Configuring Ethernet OAM for VRRP cCPE Services" on page 76.

Complete the following tasks to configure the Layer 2 CPE.

NOTE: **This procedure uses a Juniper EX Series Ethernet Switch for the Layer 2 CPE.**

1. Configure the Layer 2 CPE. See "Configuring Ethernet OAM for VRRP cCPE Services on the Layer 2 CPE" on page 77.

Complete the following tasks to configure the secondary PE router:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.
3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the first VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72.
5. Configure Ethernet OAM. See "Configuring Ethernet OAM for VRRP cCPE Services" on page 76.

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Configuring VRRP with cCPE Services on page 64
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65
- Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing) on page 67
- Configuring a VRRP IPv4 Group for cCPE Services on page 72
- *Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch*

Configuring VRRP with DHCP cCPE Services

For customers who run VRRP between two cCPE instances to achieve dual access and who also want to use DHCP to allocate IP addresses for their devices, you can configure DHCP relay agents in the two cCPE instances and point to the same DHCP server.

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

Then complete the following tasks to configure the primary PE router:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.
3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the first VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72
5. Configure tracking of the CPE facing Layer 2 interface for the first VRRP. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.
6. Configure the DHCP relay agent. See "Configuring DHCP Relay Agent cCPE Services" on page 54.

Complete the following tasks to configure the secondary PE router:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.
3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the first VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72

5. Configure tracking of the CPE-facing layer 2 interface for the first VRRP. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.
6. Configure the DHCP relay agent. See "Configuring DHCP Relay Agent cCPE Services" on page 54.

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Understanding How to Use VRRP and DHCP with CCPE on page 62
- Configuring VRRP with cCPE Services on page 64
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65
- Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing) on page 67

Configuring VRRP with DHCP cCPE Services

For customers who run VRRP between two cCPE instances to achieve dual access and who also want to use DHCP to allocate IP addresses for their devices, you can configure DHCP relay agents in the two cCPE instances and point to the same DHCP server.

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

Then complete the following tasks to configure the primary PE router:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.
3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the first VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72
5. Configure tracking of the CPE facing Layer 2 interface for the first VRRP. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.
6. Configure the DHCP relay agent. See "Configuring DHCP Relay Agent cCPE Services" on page 54.

Complete the following tasks to configure the secondary PE router:

1. Configure the subscriber access link. See *Configuring the Subscriber Access Link*.
2. Configure the bridge domain. See "Configuring the Bridge Domain" on page 74.

3. Configure the routing instance. See "Configuring the VPN Routing Instance for cCPE Services" on page 75.
4. Configure the first VRRP group. The virtual IP address for the VRRP group is the gateway address of the corresponding subnet in the cloud CPE LAN. See "Configuring a VRRP IPv4 Group for cCPE Services" on page 72
5. Configure tracking of the CPE-facing layer 2 interface for the first VRRP. See "Enabling Tracking of the CPE-Facing Layer 2 Interface" on page 73.
6. Configure the DHCP relay agent. See "Configuring DHCP Relay Agent cCPE Services" on page 54.

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Understanding How to Use VRRP and DHCP with CCPE on page 62
- Configuring VRRP with cCPE Services on page 64
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65
- Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing) on page 67

Configuring a VRRP IPv4 Group for cCPE Services

To configure a VRRP group for cCPE services, complete the following tasks on the cCPE IRB interface defined for the subscriber:

1. Specify the VRRP group.

   ```
   [edit]
   user@host# edit interfaces irb unit 106 family inet address 10.132.11.1/24 vrrp-group
     106
   ```

2. Specify the address of the virtual router in the IPv4 VRRP group.

   ```
   [edit interfaces irb unit 106 family inet address 10.132.11.1/24 vrrp-group 106]
   user@host# set virtual-address 10.132.11.3
   ```

3. Configure the VRRP router's priority for becoming the master default router.

   ```
   [edit interfaces irb unit 106 family inet address 10.132.11.1/24 vrrp-group 106]
   user@host# set priority 254
   ```

4. Enable the master VRRP virtual router to accept all packets destined for the virtual IP address.

   ```
   [edit interfaces irb unit 106 family inet address 10.132.11.1/24 vrrp-group 106]
   user@host# set accept-data
   ```

5. Configure the authentication scheme for the VRRP group to use simple password authentication.

   ```
   [edit interfaces irb unit 106 family inet address 10.132.11.1/24 vrrp-group 106]
   user@host# set authentication-type simple
   ```

6. Configure the VRRP authentication key.

   ```
   [edit interfaces irb unit 106 family inet address 10.132.11.1/24 vrrp-group 106]
   user@host# set authentication-key key
   ```

7. Verify the configuration.

   ```
   [edit interfaces irb unit 106 family inet address 10.132.11.1/24 vrrp-group 106]
   user@host# up
   user@host# up
   user@host# up
   [edit interfaces irb unit 106]
   user@host# show
   ```

   ```
   address 10.132.11.1/24 {
       vrrp-group 106 {
           virtual-address 10.132.11.3;
           priority 254;
           accept-data;
           authentication-type simple;
          authentication-key "$9$WqbLX-VwY4oZVbQnCtOBcylMxN"; ## SECRET-DATA

       }
   }
   ```

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Configuring VRRP with cCPE Services on page 64
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65
- Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing) on page 67

Enabling Tracking of the CPE-Facing Layer 2 Interface

This topic describes how to configure the VRRP group to track the status of the CPE-facing Layer 2 interface (subnet).

NOTE: If there is only one Layer 2 interface in the bridge domain, this step is optional. If the only Layer 2 interface is down, the bridge domain is also down, including the IRB interface. This also brings down the VRRP instance.

To enable interface tracking for the VRRP group:

1. Enable interface tracking for the VRRP group.

   ```
   [edit interfaces irb unit 106 family inet address 10.132.11.1/24 vrrp-group 106]
   user@host# edit track interface ge-1/2/7.106
   user@host# set priority-cost 100
   ```

2. Review the configuration.

```
[edit interfaces irb unit 106 family inet address 10.132.11.1/24 vrrp-group 106 track]
user@host# up
user@host# show
interface ge-1/2/7.106 {
    priority-cost 100;
}
```

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Understanding How to Use VRRP and DHCP with CCPE on page 62
- Configuring VRRP with cCPE Services on page 64
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65
- Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing) on page 67

Configuring the Bridge Domain

To configure the bridge domain:

1. Configure the bridge domain.

   ```
   [edit]
   user@host# edit bridge-domains vlan-106
   ```

2. Specify the domain type.

   ```
   [edit bridge-domains vlan-106]
   user@host# set domain-type bridge
   ```

3. Configure the VLAN ID.

   ```
   [edit bridge-domains vlan-106]
   user@host# set vlan-id 106
   ```

4. Specify the logical interface to include in the bridge domain.

   ```
   [edit bridge-domains vlan-106]
   user@host# set interface ge-1/2/7.106
   ```

5. Specify a routing interface to include in the bridge domain.

   ```
   [edit bridge-domains vlan-106]
   user@host# set routing-interface irb.106
   ```

6. Review the configuration.

   ```
   [edit bridge-domains vlan-106]
   user@host# show
   domain-type bridge;
   vlan-id 106;
   interface ge-1/2/7.106;
   routing-interface irb.106;
   ```

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Understanding How to Use VRRP and DHCP with CCPE on page 62
- Configuring VRRP with cCPE Services on page 64
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65
- Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing) on page 67

Configuring the VPN Routing Instance for cCPE Services

To configure the VPN routing instance for VRRP cCPE services:

1. Specify a name for the VPN routing instance.

   ```
   [edit]
   user@host# edit routing-instances acme
   ```

2. Configure the routing instance as a VRF instance.

   ```
   [edit routing-instances acme]
   user@host# set instance-type vrf
   ```

3. Specify the interface over which the VPN traffic travels to the PE router.

   Specify the interface name in the format: ***irb.logical-interface***

   ```
   [edit routing-instances acme]
   user@host# set interface irb.106
   ```

4. Configure the route distinguisher.

   ```
   [edit routing-instances acme]
   user@host# set route-distinguisher 65535:1
   ```

5. Specify the VRF target community name.

   ```
   [edit routing-instances acme]
   user@host# set vrf-target target:65535:5
   ```

6. Map the inner label of a packet to a specific VRF table. This enables examination of the encapsulated IP header.

   ```
   [edit routing-instances acme]
   user@host# set vrf-table-label
   ```

7. Review the configuration.

   ```
   [edit routing-instances acme]
   user@host# show
   ```

   ```
   interface irb.106;
   route-distinguisher 65535:1;
   vrf-target target:65535:5;
   vrf-table-label;
   ```

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Understanding How to Use VRRP and DHCP with CCPE on page 62
- Configuring VRRP with cCPE Services on page 64
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65

Configuring Ethernet OAM for VRRP cCPE Services

To configure Ethernet OAM for VRRP cCPE services on the PE routers:

1. Define an action profile with the interface down action.

   ```
   [edit]
   user@host# edit protocols oam ethernet connectivity-fault-management
   user@host# set action-profile acme default-actions interface-down
   ```

2. Configure the maintenance domain and maintenance association.

   ```
   [edit protocols oam ethernet connectivity-fault-management]
   user@host# set maintenance-domain acme-md
   [edit protocols oam ethernet connectivity-fault-management maintenance-domain
     acme-md]
   user@host# set level 6
   user@host# set maintenance-association acme
   [edit protocols oam ethernet connectivity-fault-management maintenance-domain
     acme-md maintenance-association acme]
   user@host# set continuity-check
   [edit protocols oam ethernet connectivity-fault-management maintenance-domain
     acme-md maintenance-association acme continuity-check]
   user@host# set interval 100ms
   user@host# set hold-interval 1
   ```

3. Create a local MEP on top of the Layer 2 interface you specified for the bridge domain.

   ```
   [edit protocols oam ethernet connectivity-fault-management]
   user@host# up
   user@host# set mep 201
   user@host# set interface ge-1/2/7.106
   user@host# set direction down
   user@host# set auto-discovery
   ```

4. Specify the remote MEP and the associated action profile.

   ```
   [edit protocols oam ethernet connectivity-fault-management]
   user@host# set remote-mep 101 action-profile ccpe-default
   ```

5. Repeat this procedure to configure the secondary PE router.

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Understanding How to Use VRRP and DHCP with CCPE on page 62

Configuring Ethernet OAM for VRRP cCPE Services on the Layer 2 CPE

NOTE: This procedure uses a Juniper Networks EX Series Ethernet Switch for the Layer 2 CPE.

To configure the Layer 2 CPE requirements:

1. Configure the Ethernet interfaces and VLAN.

   a. Configure the Ethernet interfaces.

      ```
      [edit]
      user@host# edit interfaces ge-0/0/10 unit 0 family ethernet-switching
      user@host set port-mode trunk
      user@host set port-mode access
      ```

   b. Configure the VLAN.

      ```
      [edit interfaces ge-1/0/0 unit 0 family ethernet-switching]
      user@host# edit vlans vlan-id 106
      user@host# set interface ge-0/0/10.0
      user@host# set interface ge-0/0/9.0
      ```

2. Configure the maintenance domain, maintenance association, and MEP.

   a. Create the maintenance domain.

      ```
      [edit]
      user@host# edit protocols oam ethernet connectivity-fault-management
      user@host# maintenance-domain acme-md
      user@host# set level 3
      ```

   b. Create the maintenance association and configure the continuity check.

      ```
      [edit protocols oam ethernet connectivity-fault-management maintenance-domain
        acme-md]
      user@host# set maintenance-association acme
      user@host# set continuity-check 100ms
      user@host# set interval 100ms
      user@host# set hold-interval 1
      ```

   c. Create local the MEP and assign the interface to it.

      ```
      [edit protocols oam ethernet connectivity-fault-management maintenance-domain
        acme-md]
      user@host# set mep 101
      ```

```
user@host# set interface ge-0/0/10.0 vlan vlan-id 106
user@host# set direction down
user@host# set auto-discovery
```

Related Documentation

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Configuring VRRP with cCPE Services on page 64
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65
- Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing) on page 67
- Configuring VRRP with Ethernet OAM cCPE Services on page 69

Configuring DHCP Relay Agent cCPE Services

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

NOTE: The stateless DHCP relay agent, which you configure under the **[routing-instance *routing-instance* forwarding-options helpers bootp]** edit hierarchy, cannot coexist with the DHCP local server in the same MX Series router, even if they are configured in different routing instances.

NOTE: When you configure a DHCP server in a VRF, if you want the server to handle DHCP requests sent by DHCP relay agents located at remote sites of the same VPN, you need to add tunnel interfaces to the DHCP server. Specifically, you need to add the following configuration to the DHCP server:

```
user@host set routing-instances acme system services dhcp-local-server
  group default-group interface lsi.0 upto lsi.16385
```

For more information about configuring the DHCP server, see "Configuring DHCP cCPE Services" on page 52.

Configure one server group and one client group for each subscriber. Define a reference to the server group with the group stanza. To configure the DHCP relay agent:

1. Configure the server group by specifying the name of the group and DHCP server addresses for use by the extended DHCP relay agent.

```
[edit]
user@host# edit routing-instances acme forwarding-options dhcp-relay
user@host# set server-group server-group-name server-ip-address
```

2. (Optional) If the subscriber is using a DHCP server farm, you can add another server to the group.

   [edit routing-instances acme forwarding-options dhcp-relay]
   user@host# set server-group *server-group-name server-ip-address*

3. (Optional) If the subscriber has more than one server group, configure an active server group to apply a common DHCP relay agent configuration to a named group of DHCP server addresses.

   [edit routing-instances acme forwarding-options dhcp-relay]
   user@host# set active-server-group *server-group-name*

4. (Optional) Specify the interface to be served by the DHCP relay agent. Typically, this is not required, because the IRB interface is the single interface for the cCPE context.

   [edit routing-instances acme forwarding-options dhcp-relay]
   user@host# set group *group-name* interface irb.*vlan-unit-id*

5. Override the default configuration settings for the extended DHCP relay agent, and enable DHCP snooping.

   [edit routing-instances acme forwarding-options dhcp-relay]
   user@host# set overrides allow-snooped-clients

6. Review the configuration.

   [edit routing-instances acme forwarding-options dhcp-relay]
   user@host# show

```
overrides {
    allow-snooped-clients;
}
server-group {
    acme-server-group {
        10.1.2.3;
        10.2.3.4;
    }
}
    active-server-group {
        10.1.2.3;
    }
active-server-group acme-server-group;
group acme-group {
    interface irb.105;
}
```

Related Documentation

- Understanding How to Use VRRP and DHCP with CCPE on page 62
- Configuring VRRP with DHCP cCPE Services on page 70
- Configuring a VRRP IPv4 Group for cCPE Services on page 72
- Understanding DHCP cCPE Services on page 51
- Verifying and Managing DHCP Relay Configuration on page 56

Verifying VRRP cCPE Services

Purpose Verify the configuration for VRRP cCPE services.

Action

- To display VRRP status:

  user@host> **show vrrp**

- To display bridge domain configuration:

  user@host> **show bridge domain *bridge-domain-name***

- To display the routing table of VPN routing instance:

  user@host> **show route table *name*.inet.0**

Related Documentation

- Configuring VRRP with cCPE Services on page 64
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65
- Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing) on page 67
- Configuring VRRP with Ethernet OAM cCPE Services on page 69
- Configuring VRRP with DHCP cCPE Services on page 70
- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59
- Understanding How to Use VRRP and DHCP with CCPE on page 62

Verifying VRRP with Ethernet OAM cCPE Services

Purpose View VRRP with Ethernet OAM cCPE services.

Action

- To display VRRP status:

  user@host> **show vrrp**

- To display bridge domain configuration:

  user@host> **show bridge domain *bridge-domain-name***

- To display the routing table of VPN routing instance:

  user@host> **show route table *name*.inet.0**

- To display CFM sessions at the PE router or the CE router:

  user@host> **show oam ethernet connectivity-fault-management mep-database**

Related Documentation

- Configuring VRRP with cCPE Services on page 64
- Configuring VRRP cCPE Services with Multiple VRRP Groups and Subnets (Load Sharing) on page 65
- Configuring VRRP cCPE Services with Multiple VRRP Groups on a Single Subnet (Load Sharing) on page 67
- Configuring VRRP with Ethernet OAM cCPE Services on page 69

- Understanding How to Use the Virtual Router Redundancy Protocol (VRRP) on cCPE Access Links on page 59

CHAPTER 9

Configuring Multiple Ethernet Interfaces for the cCPE Application

Using Multiple Ethernet Interfaces with CCPE Services

You can use multiple Ethernet interfaces to provide your customers with segregation of their LANs. When you use multiple Ethernet interfaces, the subscriber's LAN is split into multiple VLANs, for example, one for the sales office and one for manufacturing. Each VLAN uses a separate IP subnet. Inter-VLAN communication is done through routing in the Layer 3 CPE. There are a couple of solutions you can implement to achieve this same functionality when you replace the Layer 3 CPE with a Layer 2 CPE.

- The first solution is to extend subscriber VLAN all the way to PE router, either directly or through Q-in-Q tunnel. — The PE router terminates the VLANs. Inter-VLAN communication is done through routing in the VPN routing instance in the PE router. The main drawback to this solution is the inter-VLAN traffic uses up WAN bandwidth.
- The second solution is to configure inter-VLAN communication in the Layer 2 CPE to create a shortcut for inter-VLAN traffic. — Juniper Network's routed VLAN interface (RVI), which is equivalent to the Switch Virtual Interface (SVI) feature on Cisco Integrated Services Routers, is designed to facilitate inter-VLAN communication. RVI is a virtual Layer 3 interface with an IP address. RVIs allow switches to recognize which packets are being sent to another VLAN's MAC addresses—then, packets are bridged (switched) whenever the destination is within the same VLAN and are only routed through the RVI when necessary. If there are two VLANs that need to communicate with each other, you need two RVIs with addresses in each subnet associated with the

VLANS. The switch automatically creates direct routes to these subnets and uses these routes to forward traffic between VLANs. The RVI is actually Layer 3 technology so it only exists in multi-layer switches.

If subscriber VLANs use the same IP subnet, you can use a bridge domain in the Layer 2 CPE to bridge traffic from different VLANS on different interfaces into the same VLAN. When a bridge domain is defined, a VLAN ID is assigned to it. If incoming packets have a different VLAN ID than the bridge domain VLAN ID, the packets are converted to the bridge domain VLAN ID. For outgoing packets, if the egress interface has a different VLAN ID than the bridge domain VLAN ID, the packets are converted to the egress interface VLAN ID. The bridge domain effectively joins multiple VLANS together. You can configure a Layer 2 forwarding filter in the bridge domain if you need to stop broadcasts from reaching into other VLANs. However, if the subscriber's VLANs use different IP subnets, inter-VLAN traffic still travels through the WAN link because the default gateway is the PE router.

Related Documentation

- *Configuring MEI with cCPE Services—Extending Subscriber VLANs to the PE Router (One Bridge Domain per VLAN)*
- *Configuring MEI with cCPE Services—Extending Subscriber VLANs to the PE Router (All VLANs in One Bridge Domain)*
- *Configuring MEI with cCPE Services—Extending Subscriber VLANs to the PE Router Using a Q-in-Q Tunnel*
- *Configuring MEI and RVIs with cCPE Services*

Configuring MEI with CCPE Services—Extending Subscriber VLANs to the PE Router (One Bridge Domain Per VLAN)

This procedure configures one bridge domain and one IRB per subscriber VLAN. Each bridge domain has one Layer 2 interface. The IRB interface is configured with the IP address matching the subnet of the VLAN. Using multiple bridge domains effectively creates a separate broadcast domain for each VLAN. You also need to configure a VRF type routing instance for the VPN site and reference the IRB (one for each subscriber VLAN) as the interface. To use this configuration, verify that multiple subscriber VLANs can be provided.

NOTE: **It is possible that subscriber VLAN ID may be overlapping between multiple customers. VLAN ID translation can be configured on the Layer 2 CPE or the aggregation router to avoid VLAN ID conflict.**

- Configuring the PE Router for Multiple Ethernet Interfaces with cCPE Services (One Bridge Domain per VLAN) on page 85
- Configuring the Layer 2 CPE for Multiple Ethernet Interfaces cCPE Services on page 88
- Verifying Multiple Ethernet Interfaces for cCPE Services on page 89

Configuring the PE Router for Multiple Ethernet Interfaces with cCPE Services (One Bridge Domain per VLAN)

In this procedure you configure two VLANs for the subscriber and one bridge domain and one IRB interface per VLAN. You also need to configure a VRF routing instance for the subscriber VPN site and reference each of the IRB (one for each subscriber VLAN) as the interface.

Complete the following tasks to configure multiple Ethernet interfaces with cCPE services with one bridge domain per VLAN:

1. Configuring the Subscriber Access Link and VLANs for Multiple Ethernet Interfaces with cCPE Services on page 85
2. Configuring the IRB Interfaces for Multiple Ethernet Interfaces with cCPE Services on page 86
3. Configuring the Bridge Domains for Multiple Ethernet Interfaces with cCPE Services on page 86
4. Configuring the Routing Instance for Multiple Ethernet Interfaces with cCPE Services on page 87

Configuring the Subscriber Access Link and VLANs for Multiple Ethernet Interfaces with cCPE Services

To configure the access link and VLANs:

1. Configure the physical interface.

```
[edit]
user@host# edit interfaces ge-1/2/3
user@host# set vlan-tagging
user@host# set encapsulation flexible-ethernet-services
```

2. Configure the first logical interface and VLAN.

```
[edit interfaces ge-1/2/3]
user@host# edit unit 105
user@host# set encapsulation vlan-bridge
user@host# set vlan-id 105
user@host# set family bridge
```

3. Configure the second logical interface and VLAN.

```
[edit interfaces ge-1/2/3 unit 105]
user@host# up
user@host# edit unit 106
[edit interfaces ge-1/2/3 unit 106]
user@host# set encapsulation vlan-bridge
user@host# set vlan-id 106
user@host# set family bridge
```

4. Review your configuration.

```
[edit interfaces ge-1/2/3 unit 106]
user@host# up
user@host# show
```

```
unit 105 {
    encapsulation vlan-bridge;
    vlan-id 105;
    family bridge;
}
unit 106 {
    encapsulation vlan-bridge;
    vlan-id 106;
    }
```

Configuring the IRB Interfaces for Multiple Ethernet Interfaces with cCPE Services

Configure the IRB interfaces, with the IP addresses matching the subnets of the VLAN.

To configure the IRB interfaces:

1. Configure the first IRB interface.

   ```
   [edit]
   user@host# edit interfaces irb unit 105 family inet address 192.168.1.1/24
   ```

2. Configure the second IRB interface.

   ```
   [edit interfaces irb unit 105]
   user@host# up
   [edit interfaces irb]
   user@host# edit unit 106 family inet address 192.168.2.1/24
   ```

3. Review the configuration.

   ```
   [edit interfaces irb unit 106 family inet address 192.168.2.1/24]
   user@host# up
   user@host# up
   user@host# up
   [edit interfaces irb]
   user@host# show
   ```

```
unit 105 {
    family inet {
        address 192.168.1.1/24;
    }
}
unit 106 {
    family inet {
        address 192.168.2.1/24;
    }
}
```

Configuring the Bridge Domains for Multiple Ethernet Interfaces with cCPE Services

To configure the bridge domains:

1. Configure the first bridge domain.

   ```
   [edit]
   user@host# edit bridge-domains bd-105
   user@host# set vlan-id 105
   user@host# set interface ge-1/2/3.105
   user@host# set routing-interface irb.105
   ```

2. Configure the second bridge domain.

```
[edit bridge-domains bd-105]
user@host# up
user@host# edit bd-106
user@host# set vlan-id 106
user@host# set interface ge-1/2/3.106
user@host# set routing-interface irb.106
```

3. Review the configuration.

```
[edit bridge-domains]
user@host# show
```

```
bd-105 {
    vlan-id 105;
    interface ge-1/2/3.105;
    routing-interface irb.105;
}
bd-106 {
    vlan-id 106;
    interface ge-1/2/3.106;
    routing-interface irb.106;
}
```

Configuring the Routing Instance for Multiple Ethernet Interfaces with cCPE Services

Configure a VRF routing instance for the subscriber VPN site, and reference each of the IRB interfaces.

To configure the routing instance:

1. Configure the name of the routing instance.

```
[edit]
user@host# edit routing-instances ccpe1
```

2. Configure the VRF routing instance.

```
[edit routing-instances ccpe1]
user@host# set instance-type vrf
```

3. Reference the first IRB interface.

```
[edit routing-instances ccpe1]
user@host# set interface irb.105
```

4. Reference the second IRB interface.

```
[edit routing-instances ccpe1]
user@host# set interface irb.106
```

5. Review the configuration.

```
[edit routing-instances ccpe1]
user@host# show
```

```
instance-type vrf;
interface irb.105;
interface irb.106;
```

Configuring the Layer 2 CPE for Multiple Ethernet Interfaces cCPE Services

NOTE: This procedure uses a Juniper Networks EX Series Ethernet Switch for the Layer 2 CPE.

To configure the Layer 2 CPE for multiple Ethernet interfaces support:

1. Configure the first Ethernet interface on the CPE.

   ```
   [edit]
   user@host# edit interfaces ge-0/0/8 unit 0
   user@host# set family ethernet-switching
   user@host# set port-mode access
   ```

2. Configure the second Ethernet interface.

   ```
   [edit interfaces ge-0/0/8 unit 0]
   user@host# up
   user@host# up
   user@host# edit interfaces ge-0/0/9 unit 0
   user@host# set family ethernet-switching
   user@host# set port-mode access
   ```

3. Configure the WAN link for trunk mode.

   ```
   [edit interfaces ge-0/0/9 unit 0]
   user@host# up
   user@host# up
   user@host# edit interfaces ge-0/0/10 unit 0
   user@host# set family ethernet-switching
   user@host# set port-mode trunk
   ```

4. Configure the VLAN for the first department; for example, Sales.

   a. Specify the VLAN name.

      ```
      [edit interfaces ge-0/0/10]
      user@host# top
      user@host# edit vlans sales-vlan
      ```

   b. Specify the VLAN ID.

      ```
      [edit vlans sales-vlan]
      user@host# set vlan-id 105
      ```

   c. Specify the interface associated with the VLAN.

      ```
      [edit vlans sales-vlan]
      user@host# set interface ge-0/0/8.0
      ```

   d. Specify the trunk port interface associated with the VLAN.

      ```
      [edit vlans sales-vlan]
      user@host# set interface ge-0/0/10.0
      ```

5. Configure the VLAN for the second department; for example, Engineering.

   a. Specify the VLAN name.

```
[edit vlans sales-vlan]
user@host# top
user@host# edit vlans eng-vlan
```

b. Specify the VLAN ID.

```
[edit vlans eng-vlan]
user@host# set vlan-id 106
```

c. Specify the interface associated with the VLAN.

```
[edit vlans eng-vlan]
user@host# set interface ge-0/0/9.0
```

d. Specify the trunk port interface associated with the VLAN.

```
[edit vlans eng-vlan]
user@host# set interface ge-0/0/10.0
```

Verifying Multiple Ethernet Interfaces for cCPE Services

Purpose Verify multiple Ethernet interfaces for cCPE services.

Action

- To display the ARP table associated with the Ethernet interface:

  user@host> **show arp interface *interface-name***

- To display the bridge domain configuration:

  user@host> **show bridge domain *domain-name***

- To display the route table of the VPN routing instance:

  user@host> **show route table *name.inet.0***

Related Documentation

- *Layer 2 Bridging Interfaces Overview*
- *Configuring Layer 2 Bridging Interfaces*
- *Configuring Routing Instances on PE Routers in VPNs*
- *Understanding Bridging and VLANs on EX Series Switches*
- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- *Understanding Interface Naming Conventions on EX Series Switches*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *EX Series Switches Interfaces Overview*
- *Configuring VLANs for EX Series Switches (CLI Procedure)*

Configuring MEI with CCPE Services—Extending Subscriber VLANs to the PE Router (All VLANs in One Bridge Domain)

This procedure configures all subscriber VLANs in the same bridge domain so only one IRB interface is required per subscriber. You need to configure multiple IP addresses (one per subnet) on the IRB interface. The bridge domain joins multiple VLANs together. To

control broadcasts from one VLAN reaching into other VLANs, you do not enable local-switching in the bridge domain. Alternatively, you can control selected broadcast traffic with a Layer 2 forwarding filter.

NOTE: It is possible that subscriber VLAN ID may be overlapping between multiple customers. VLAN ID translation can be configured on the Layer 2 CPE or the aggregation router to avoid VLAN ID conflict.

Configuring the PE Router for Multiple Ethernet Interfaces with cCPE Services (All VLANs in One Bridge Domain)

Complete the following tasks to configure the PE router for multiple Ethernet interfaces with cCPE services:

Configuring the Subscriber Access Link and VLANs for Multiple Ethernet Interfaces with cCPE Services

To configure the access link and VLANs:

1. Configure the physical interface.

   ```
   [edit]
   user@host# edit interfaces ge-1/2/3
   user@host# set vlan-tagging
   user@host# set encapsulation flexible-ethernet-services
   ```

2. Configure the first logical interface and VLAN.

   ```
   [edit interfaces ge-1/2/3]
   user@host# edit unit 105
   user@host# set encapsulation vlan-bridge
   user@host# set vlan-id 105
   user@host# set family bridge
   ```

3. Configure the second logical interface and VLAN.

   ```
   [edit interfaces ge-1/2/3 unit 105]
   user@host# up
   ```

```
user@host# edit unit 106
[edit interfaces ge-1/2/3 unit 106]
user@host# set encapsulation vlan-bridge
user@host# set vlan-id 106
user@host# set family bridge
```

4. Review your configuration.

```
[edit interfaces ge-1/2/3 unit 106]
user@host# up
user@host# show
```

```
unit 105 {
    encapsulation vlan-bridge;
    vlan-id 105;
    family bridge;
}
unit 106 {
    encapsulation vlan-bridge;
    vlan-id 106;
    }
```

Configuring the Bridge Domain for Multiple Ethernet Interfaces with cCPE Services

To configure the bridge domain:

1. Configure the bridge domain for the first subscriber VLAN by specifying the name of the bridge domain.

```
[edit]
user@host# edit bridge-domains bd-105
```

2. Associate the subscriber's VLAN ID with the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set vlan-id 105
```

3. Specify the interface names for the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set interface ge-1/2/3.105
user@host# set interface ge-1/2/3.106
```

4. Specify the routing interface to include in the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set routing-interface irb.105
```

5. To control broadcast from one VLAN reaching into other VLANs, you can either enable **no-local-switching** in the bridge domain or configure a Layer 2 filter.

   - (Optional) To enable **no-local-switching** in the bridge domain:

```
[edit bridge-domains bd-105]
user@host# set routing-interface irb.105 no-local-switching
```

   - (Optional) To configure a Layer 2 filter:

```
[edit bridge-domains bd-105]
user@host# set forwarding-options filter input no-dhcp
user@host# top
```

```
user@host# edit firewall family bridge filter no-dhcp
[edit firewall family bridge filter no-dhcp]
user@host# set term dhcp from port dhcp
user@host# set term dhcp then discard
user@host# set term accept-all then accept
```

6. Review the configuration.

```
[edit bridge-domains bd-105]
user@host# show
```

```
vlan-id 105;
no-local-switching;
interface ge-1/2/3.105;
interface ge-1/2/3.106;
routing-interface irb.105;
```

Configuring the Routing Instances for Multiple Ethernet Interfaces with cCPE Services

Configure a VRF routing instance for the subscriber VPN site, and reference the IRB as the interface. To configure the routing instance:

1. Configure a name for the routing instance.

```
[edit]
user@host# edit routing-instances ccpe1
```

2. Configure the VRF routing instance.

```
[edit routing-instances ccpe1]
user@host# set instance-type vrf
```

3. Reference the IRB interface by specifying irb.*vlan-id*.

```
[edit routing-instances ccpe1]
user@host# set interface irb.105
```

4. Review the configuration.

```
[edit routing-instances ccpe1]
user@host# show
```

```
instance-type vrf;
interface irb.105;
```

Configuring the IRB Interface for Multiple Ethernet Interfaces with cCPE Services

Configure the IRB interface, with the IP address matching the subnet of the VLAN.

To configure the IRB interface:

1. Configure the first subnet on the IRB interface.

```
[edit]
user@host# edit interfaces irb unit 105 family inet
user@host# set address 192.168.1.1/24
```

2. Configure the second subnet on the IRB interface.

```
[edit interfaces irb unit 105 family inet]
```

```
user@host# set address 192.168.2.1/24
```

3. Review the configuration.

```
[edit interfaces irb unit 106 family inet]
user@host# show
```

```
unit 105 {
    family inet {
        address 192.168.1.1/24;
        address 192.168.2.1/24;
    }
}
```

Configuring the Layer 2 CPE for Multiple Ethernet Interfaces cCPE Services

NOTE: This procedure uses a Juniper Networks EX Series Ethernet Switch for the Layer 2 CPE.

To configure the Layer 2 CPE for multiple Ethernet interfaces support:

1. Configure the first Ethernet interface on the CPE.

```
[edit]
user@host# edit interfaces ge-0/0/8 unit 0
user@host# set family ethernet-switching
user@host# set port-mode access
```

2. Configure the second Ethernet interface.

```
[edit interfaces ge-0/0/8 unit 0]
user@host# up
user@host# up
user@host# edit interfaces ge-0/0/9 unit 0
user@host# set family ethernet-switching
user@host# set port-mode access
```

3. Configure the WAN link for trunk mode.

```
[edit interfaces ge-0/0/9 unit 0]
user@host# up
user@host# up
user@host# edit interfaces ge-0/0/10 unit 0
user@host# set family ethernet-switching
user@host# set port-mode trunk
```

4. Configure the VLAN for the first department; for example, Sales.

   a. Specify the VLAN name.

```
[edit interfaces ge-0/0/10]
user@host# top
user@host# edit vlans sales-vlan
```

   b. Specify the VLAN ID.

```
[edit vlans sales-vlan]
user@host# set vlan-id 105
```

c. Specify the interface associated with the VLAN.

```
[edit vlans sales-vlan]
user@host# set interface ge-0/0/8.0
```

d. Specify the trunk port interface associated with the VLAN.

```
[edit vlans sales-vlan]
user@host# set interface ge-0/0/10.0
```

5. Configure the VLAN for the second department; for example, Engineering.

a. Specify the VLAN name.

```
[edit vlans sales-vlan]
user@host# top
user@host# edit vlans eng-vlan
```

b. Specify the VLAN ID.

```
[edit vlans eng-vlan]
user@host# set vlan-id 106
```

c. Specify the interface associated with the VLAN.

```
[edit vlans eng-vlan]
user@host# set interface ge-0/0/9.0
```

d. Specify the trunk port interface associated with the VLAN.

```
[edit vlans eng-vlan]
user@host# set interface ge-0/0/10.0
```

Verifying Multiple Ethernet Interfaces for cCPE Services

Purpose Verify multiple Ethernet interfaces for cCPE services.

Action

- To display the ARP table associated with the Ethernet interface:

  user@host> **show arp interface *interface-name***

- To display the bridge domain configuration:

  user@host> **show bridge domain *domain-name***

- To display the route table of the VPN routing instance:

  user@host> **show route table *name.inet.0***

Related Documentation

- *Layer 2 Bridging Interfaces Overview*
- *Configuring Layer 2 Bridging Interfaces*
- *Configuring Routing Instances on PE Routers in VPNs*
- *Understanding Bridging and VLANs on EX Series Switches*
- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- *Understanding Interface Naming Conventions on EX Series Switches*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

- *EX Series Switches Interfaces Overview*
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- *Guidelines for Configuring Firewall Filters*

Configuring MEI with CCPE Services—Extending Subscriber VLANs to the PE Router Using a Q-in-Q Tunnel

With the subscriber VLAN in a Q-in-Q tunnel, you do not need to be concerned about conflicting subscriber VLAN IDs between multiple customers. You can configure the Q-in-Q tunnel in the Layer 2 CPE or an aggregation switch.

NOTE: **This configuration uses one bridge domain and two subscriber VLANs from the Layer 2 CPE.**

At a high-level, you need to configure the following for this configuration:

1. On the PE router, configure the logical interface with a stacked VLAN.
2. On the PE router, configure one bridge-domain and one IRB for each subscriber VLAN.
3. On the PE router, configure a VRF type routing-instance for the VPN site and reference the IRB as the interface.
4. On the Layer 2 CPE (EX Series switch), configure a Q-in-Q tunnel.

- Configuring the PE Router for Multiple Ethernet Interfaces cCPE Services (Q-in-Q Tunnel) on page 95
- Configuring the Layer 2 CPE for Multiple Ethernet Interfaces cCPE Services on page 98
- Configuring the Q-in-Q Tunnel on the Aggregation Switch (EX Series Ethernet Switch) on page 99
- Verifying Multiple Ethernet Interfaces for cCPE Services on page 101

Configuring the PE Router for Multiple Ethernet Interfaces cCPE Services (Q-in-Q Tunnel)

Complete the following tasks to configure the PE router for multiple Ethernet interfaces cCPE services (Q-in-Q tunnel):

- Configuring the Subscriber Access Link and VLANs for Multiple Ethernet Interfaces cCPE Services (Q-in-Q Tunnel) on page 96
- Configuring the IRB Interface for Multiple Ethernet Interfaces with cCPE Services on page 96
- Configuring the Bridge Domain for Multiple Ethernet Interfaces cCPE Services (Q-in-Q Tunnel) on page 97
- Configuring the Routing Instances for Multiple Ethernet Interfaces with cCPE Services on page 98

Configuring the Subscriber Access Link and VLANs for Multiple Ethernet Interfaces cCPE Services (Q-in-Q Tunnel)

To configure the subscriber access link and VLANs:

1. Configure the physical interface.

   ```
   [edit]
   user@host# edit interfaces ge-1/2/3
   user@host# set vlan-tagging
   user@host# set encapsulation flexible-ethernet-services
   ```

2. Configure the first logical interface for Ethernet VLAN bridge encapsulation.

   ```
   [edit interfaces ge-1/2/3]
   user@host# set stacked-vlan-tagging
   user@host# edit unit 105
   user@host# setencapsulation vlan-bridge
   user@host# set vlan-tags outer 1000
   user@host# set vlan-tags inner 105
   user@host# set family bridge
   ```

3. Configure the second logical interface for Ethernet VLAN bridge encapsulation.

   ```
   [edit interfaces ge-1/2/3 unit 105]
   user@host# up
   user@host# edit unit 106
   user@host# set encapsulation vlan-bridge
   user@host# set vlan-tags outer 1000
   user@host# set vlan-tags inner 106
   user@host# set family bridge
   ```

4. Review the configuration.

   ```
   [edit interfaces ge-1/2/3 unit 106]
   user@host# up
   user@host# show
   ```

```
stacked-vlan-tagging;
encapsulation flexible-ethernet-services;
unit 105 {
    encapsulation vlan-bridge;
    vlan-tags outer 1000 inner 105;
    family bridge;
}
unit 106 {
    encapsulation vlan-bridge;
    vlan-tags outer 1000 inner 106;
    family bridge;
}
```

Configuring the IRB Interface for Multiple Ethernet Interfaces with cCPE Services

Configure the IRB interface, with the IP address matching the subnet of the VLAN.

To configure the IRB interface:

1. Configure the first subnet on the IRB interface.

```
[edit]
user@host# edit interfaces irb unit 105 family inet
user@host# set address 192.168.1.1/24
```

2. Configure the second subnet on the IRB interface.

```
[edit interfaces irb unit 105 family inet]
user@host# set address 192.168.2.1/24
```

3. Review the configuration.

```
[edit interfaces irb unit 106 family inet]
user@host# show
```

```
unit 105 {
    family inet {
        address 192.168.1.1/24;
        address 192.168.2.1/24;
    }
}
```

Configuring the Bridge Domain for Multiple Ethernet Interfaces cCPE Services (Q-in-Q Tunnel)

To configure the bridge domain:

1. Configure the name of the bridge domain.

```
[edit]
user@host# edit bridge-domains bd-105
```

2. Associate the subscriber's VLAN ID with the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set vlan-id 105
```

3. Specify the interface names for the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set interface ge-1/2/3.105
user@host# set interface ge-1/2/3.106
```

4. Specify the routing interface to include in the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set routing-interface irb.105
```

5. Review the configuration.

```
[edit bridge-domains bd-105]
user@host# show
```

```
vlan-id 105;
interface ge-1/2/3.105;
interface ge-1/2/3.106;
routing-interface irb.105;
```

Configuring the Routing Instances for Multiple Ethernet Interfaces with cCPE Services

Configure a VRF routing instance for the subscriber VPN site, and reference the IRB as the interface. To configure the routing instance:

1. Configure a name for the routing instance.

   ```
   [edit]
   user@host# edit routing-instances ccpe1
   ```

2. Configure the VRF routing instance.

   ```
   [edit routing-instances ccpe1]
   user@host# set instance-type vrf
   ```

3. Reference the IRB interface by specifying irb.*vlan-id*.

   ```
   [edit routing-instances ccpe1]
   user@host# set interface irb.105
   ```

4. Review the configuration.

   ```
   [edit routing-instances ccpe1]
   user@host# show
   ```

   ```
   instance-type vrf;
   interface irb.105;
   ```

Configuring the Layer 2 CPE for Multiple Ethernet Interfaces cCPE Services

NOTE: This procedure uses a Juniper Networks EX Series Ethernet Switch for the Layer 2 CPE.

To configure the Layer 2 CPE for multiple Ethernet interfaces support:

1. Configure the first Ethernet interface on the CPE.

   ```
   [edit]
   user@host# edit interfaces ge-0/0/8 unit 0
   user@host# set family ethernet-switching
   user@host# set port-mode access
   ```

2. Configure the second Ethernet interface.

   ```
   [edit interfaces ge-0/0/8 unit 0]
   user@host# up
   user@host# up
   user@host# edit interfaces ge-0/0/9 unit 0
   user@host# set family ethernet-switching
   user@host# set port-mode access
   ```

3. Configure the WAN link for trunk mode.

   ```
   [edit interfaces ge-0/0/9 unit 0]
   user@host# up
   user@host# up
   user@host# edit interfaces ge-0/0/10 unit 0
   ```

```
user@host# set family ethernet-switching
user@host# set port-mode trunk
```

4. Configure the VLAN for the first department; for example, Sales.

   a. Specify the VLAN name.

   ```
   [edit interfaces ge-0/0/10]
   user@host# top
   user@host# edit vlans sales-vlan
   ```

   b. Specify the VLAN ID.

   ```
   [edit vlans sales-vlan]
   user@host# set vlan-id 105
   ```

   c. Specify the interface associated with the VLAN.

   ```
   [edit vlans sales-vlan]
   user@host# set interface ge-0/0/8.0
   ```

   d. Specify the trunk port interface associated with the VLAN.

   ```
   [edit vlans sales-vlan]
   user@host# set interface ge-0/0/10.0
   ```

5. Configure the VLAN for the second department; for example, Engineering.

   a. Specify the VLAN name.

   ```
   [edit vlans sales-vlan]
   user@host# top
   user@host# edit vlans eng-vlan
   ```

   b. Specify the VLAN ID.

   ```
   [edit vlans eng-vlan]
   user@host# set vlan-id 106
   ```

   c. Specify the interface associated with the VLAN.

   ```
   [edit vlans eng-vlan]
   user@host# set interface ge-0/0/9.0
   ```

   d. Specify the trunk port interface associated with the VLAN.

   ```
   [edit vlans eng-vlan]
   user@host# set interface ge-0/0/10.0
   ```

Configuring the Q-in-Q Tunnel on the Aggregation Switch (EX Series Ethernet Switch)

NOTE: This procedure assumes you are using a Juniper Networks EX Series Ethernet Switch for the aggregation switch.

To configure the Q-in-Q tunnel:

1. Configure the physical and logical interfaces for the subscriber's Ethernet LAN.

   a. Configure the physical and logical interface.

      ```
      [edit]
      user@host# edit interfaces ge-0/0/2 unit 0
      ```

   b. Configure the protocol family for the logical interface on the switch.

      ```
      [edit interfaces ge-0/0/2 unit 0]
      user@host# edit family ethernet-switching
      ```

   c. Configure the interface to operate in access mode.

      ```
      [edit interfaces ge-0/0/2 unit 0 family ethernet-switching]
      user@host# set port-mode access
      ```

   d. Bind the subscriber VLAN to the logical interface.

      ```
      [edit interfaces ge-0/0/2 unit 0 family ethernet-switching]
      user@host# set vlan members qinq-1000
      ```

2. Configure the VLAN trunk port for the WAN link.

   a. Configure the physical and logical interface.

      ```
      user@host# top
      user@host# edit interfaces ge-0/0/23 unit 0
      ```

   b. Configure the protocol family.

      ```
      [edit interfaces ge-0/0/23 unit 0]
      user@host# edit family ethernet-switching
      ```

   c. Configure the WAN link as a trunk interface.

      ```
      [edit interfaces ge-0/0/23 unit 0 family ethernet-switching]
      user@host# set port-mode trunk
      ```

   d. Configure the subscriber VLAN on the trunk interface.

      ```
      [edit interfaces ge-0/0/23 unit 0 family ethernet-switching]
      user@host# set vlan members qinq-1000
      ```

3. Configure the switching options for the tunnel.

   ```
   user@host# top
   [edit]
   user@host# set ethernet-switching-options dot1q-tunneling ether-type 0x8100
   ```

4. Configure the subscriber VLANs.

   a. Configure the VLAN name.

      ```
      [edit]
      user@host# edit vlans qinq-1000
      ```

   b. Configure the VLAN ID.

      ```
      [edit vlans qinq-1000]
      user@host# set vlan-id 1000
      ```

   c. Configure the interfaces for the VLAN.


```
[edit vlans qinq-1000]
user@host# set interface ge-0/0/2.0
user@host# set interface ge-0/0/23.0
```

d. Enable Q-in-Q tunneling on the subscriber VLANs.

```
[edit vlans qinq-1000]
user@host# set dot1q-tunneling customer-vlans 105
user@host# set dot1q-tunneling customer-vlans 106
```

Verifying Multiple Ethernet Interfaces for cCPE Services

Purpose Verify multiple Ethernet interfaces for cCPE services.

Action

- To display the ARP table associated with the Ethernet interface:

  user@host> **show arp interface *interface-name***

- To display the bridge domain configuration:

  user@host> **show bridge domain *domain-name***

- To display the route table of the VPN routing instance:

  user@host> **show route table *name.inet.0***

Related Documentation

- *Layer 2 Bridging Interfaces Overview*
- *Configuring Layer 2 Bridging Interfaces*
- *Configuring Routing Instances on PE Routers in VPNs*
- *Understanding Bridging and VLANs on EX Series Switches*
- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- *Understanding Interface Naming Conventions on EX Series Switches*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *EX Series Switches Interfaces Overview*
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- *Understanding Q-in-Q Tunneling on EX Series Switches*
- *Example: Setting Up Q-in-Q Tunneling on EX Series Switches*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *EX Series Switches Interfaces Overview*
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- *Configuring Q-in-Q Tunneling (CLI Procedure)*

Configuring MEI with CCPE Services—Routed VLAN Interfaces (RVI)

When you configure routed VLAN interfaces (RVIs) in the Layer 2 CPE, it creates a shortcut for inter-VLAN traffic. There is one RVI per subscriber VLAN. The RVI creates a direct route in the Layer 2 CPE for inter-VLAN traffic.

NOTE: To run this configuration the host connecting to the customer VLAN needs a static route for inter-VLAN traffic with the RVI address as the gateway.

NOTE: This procedure uses a Juniper EX Series Ethernet Switch for the Layer 2 CPE.

- Configuring the PE Router for Multiple Ethernet Interfaces with cCPE Services (All VLANs in One Bridge Domain) on page 102
- Configuring the Layer 2 CPE with Multiple Ethernet Interfaces and RVIs for cCPE Services on page 105
- Verifying Multiple Ethernet Interfaces with RVIs for cCPE Services on page 107

Configuring the PE Router for Multiple Ethernet Interfaces with cCPE Services (All VLANs in One Bridge Domain)

Complete the following tasks to configure the PE router for multiple Ethernet interfaces with cCPE services:

- Configuring the Subscriber Access Link and VLANs for Multiple Ethernet Interfaces with cCPE Services on page 102
- Configuring the Bridge Domain for Multiple Ethernet Interfaces with cCPE Services on page 103
- Configuring the Routing Instances for Multiple Ethernet Interfaces with cCPE Services on page 104
- Configuring the IRB Interface for Multiple Ethernet Interfaces with cCPE Services on page 105

Configuring the Subscriber Access Link and VLANs for Multiple Ethernet Interfaces with cCPE Services

To configure the access link and VLANs:

1. Configure the physical interface.

   ```
   [edit]
   user@host# edit interfaces ge-1/2/3
   user@host# set vlan-tagging
   user@host# set encapsulation flexible-ethernet-services
   ```

2. Configure the first logical interface and VLAN.

   ```
   [edit interfaces ge-1/2/3]
   ```

```
user@host# edit unit 105
user@host# set encapsulation vlan-bridge
user@host# set vlan-id 105
user@host# set family bridge
```

3. Configure the second logical interface and VLAN.

```
[edit interfaces ge-1/2/3 unit 105]
user@host# up
user@host# edit unit 106
[edit interfaces ge-1/2/3 unit 106]
user@host# set encapsulation vlan-bridge
user@host# set vlan-id 106
user@host# set family bridge
```

4. Review your configuration.

```
[edit interfaces ge-1/2/3 unit 106]
user@host# up
user@host# show
```

```
unit 105 {
    encapsulation vlan-bridge;
    vlan-id 105;
    family bridge;
}
unit 106 {
    encapsulation vlan-bridge;
    vlan-id 106;
    }
```

Configuring the Bridge Domain for Multiple Ethernet Interfaces with cCPE Services

To configure the bridge domain:

1. Configure the bridge domain for the first subscriber VLAN by specifying the name of the bridge domain.

```
[edit]
user@host# edit bridge-domains bd-105
```

2. Associate the subscriber's VLAN ID with the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set vlan-id 105
```

3. Specify the interface names for the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set interface ge-1/2/3.105
user@host# set interface ge-1/2/3.106
```

4. Specify the routing interface to include in the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set routing-interface irb.105
```

5. To control broadcast from one VLAN reaching into other VLANs, you can either enable **no-local-switching** in the bridge domain or configure a Layer 2 filter.

- (Optional) To enable **no-local-switching** in the bridge domain:

  ```
  [edit bridge-domains bd-105]
  user@host# set routing-interface irb.105 no-local-switching
  ```

- (Optional) To configure a Layer 2 filter:

  ```
  [edit bridge-domains bd-105]
  user@host# set forwarding-options filter input no-dhcp
  user@host# top
  user@host# edit firewall family bridge filter no-dhcp
  [edit firewall family bridge filter no-dhcp]
  user@host# set term dhcp from port dhcp
  user@host# set term dhcp then discard
  user@host# set term accept-all then accept
  ```

6. Review the configuration.

   ```
   [edit bridge-domains bd-105]
   user@host# show
   ```

   ```
   vlan-id 105;
   no-local-switching;
   interface ge-1/2/3.105;
   interface ge-1/2/3.106;
   routing-interface irb.105;
   ```

Configuring the Routing Instances for Multiple Ethernet Interfaces with cCPE Services

Configure a VRF routing instance for the subscriber VPN site, and reference the IRB as the interface. To configure the routing instance:

1. Configure a name for the routing instance.

   ```
   [edit]
   user@host# edit routing-instances ccpe1
   ```

2. Configure the VRF routing instance.

   ```
   [edit routing-instances ccpe1]
   user@host# set instance-type vrf
   ```

3. Reference the IRB interface by specifying irb.*vlan-id*.

   ```
   [edit routing-instances ccpe1]
   user@host# set interface irb.105
   ```

4. Review the configuration.

   ```
   [edit routing-instances ccpe1]
   user@host# show
   ```

   ```
   instance-type vrf;
   interface irb.105;
   ```

Configuring the IRB Interface for Multiple Ethernet Interfaces with cCPE Services

Configure the IRB interface, with the IP address matching the subnet of the VLAN.

To configure the IRB interface:

1. Configure the first subnet on the IRB interface.

   ```
   [edit]
   user@host# edit interfaces irb unit 105 family inet
   user@host# set address 192.168.1.1/24
   ```

2. Configure the second subnet on the IRB interface.

   ```
   [edit interfaces irb unit 105 family inet]
   user@host# set address 192.168.2.1/24
   ```

3. Review the configuration.

   ```
   [edit interfaces irb unit 106 family inet]
   user@host# show
   ```

   ```
   unit 105 {
       family inet {
           address 192.168.1.1/24;
           address 192.168.2.1/24;
       }
   }
   ```

Configuring the Layer 2 CPE with Multiple Ethernet Interfaces and RVIs for cCPE Services

NOTE: This procedure uses a Juniper Networks EX Series Ethernet Switch for the Layer 2 CPE.

To configure the Layer 2 CPE for Multiple Ethernet Interfaces with RVIs:

1. Configure the first Ethernet interface on the Layer 2 CPE.

   ```
   [edit]
   user@host# set interfaces ge-0/0/8 unit 0 family ethernet-switching port-mode
     access
   ```

2. Configure the second Ethernet interface by specifying a different physical interface. You can use the same logical interface number, which in our example is 0.

   ```
   [edit]
   user@host# set interfaces ge-0/0/9 unit 0 family ethernet-switching port-mode
     access
   ```

3. Configure the VLAN trunk port for a WAN link.

   ```
   [edit]
   user@host# set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode
     trunk
   ```

4. Configure the RVIs.

   ```
   [edit]
   ```

```
user@host# set interfaces vlan unit 0 family inet address 192.168.1.254/24
user@host# set interfaces vlan unit 1 family inet address 192.168.2.254/24
```

5. Configure the VLAN for the first department; for example, Sales.

   a. Specify the VLAN name and VLAN ID.

      ```
      [edit]
      user@host# set vlans sales-vlan vlan-id 105
      ```

   b. Specify the interface associated with the VLAN.

      ```
      [edit vlans sales-vlan]
      user@host# set interface ge-0/0/8.0
      ```

   c. Specify the trunk port interface associated with the VLAN.

      ```
      [edit vlans sales-vlan]
      user@host# set interface ge-0/0/10.0
      ```

   d. Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, whereas across VLANs, traffic is routed.

      ```
      [edit vlans sales-vlan]
      user@host# set l3-interface vlan.0
      ```

6. Configure the VLAN for the second department; for example, Engineering.

   a. Specify the VLAN name and VLAN ID.

      ```
      [edit]
      user@host# set vlans eng-vlan vlan-id 106
      ```

   b. Specify the interface associated with the VLAN.

      ```
      [edit vlans eng-vlan]
      user@host# set interface ge-0/0/9.0
      ```

   c. Specify the trunk port interface associated with the VLAN.

      ```
      [edit vlans eng-vlan]
      user@host# set interface ge-0/0/10.0
      ```

   d. Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, whereas across VLANs, traffic is routed.

      ```
      [edit vlans eng-vlan]
      user@host# set l3-interface vlan.1
      ```

7. Finally, configure a static route with the RVI addresses as the gateway on the host connecting to the subscriber VLAN for inter-VLAN traffic. If the customer host device uses DHCP to get IP addresses, a static route can be pushed by the DHCP server through option 121; otherwise, you must manually configure the static route.

   - Add a static route for the subnet using the RVI as the gateway.

     ```
     route add -net 192.168.2.0/24 gw 192.168.1.254
     ```

   - Add a static route using the IRB as the gateway.

     ```
     route add -net 0.0.0.0/0 gw 192.168.1.1
     ```

Verifying Multiple Ethernet Interfaces with RVIs for cCPE Services

Purpose View multiple Ethernet interfaces with RVIs for cCPE services.

Action

- To display the ARP table associated with the Ethernet interface:

  user@host> **show arp interface *interface-name***

- To display the bridge domain configuration:

  user@host> **show bridge domain *domain-name***

- To display the route table of the VPN routing instance:

  user@host> **show route table *name.inet.0***

- To display the RVIs and their current states in EX Series switch:

  user@host> **show interfaces vlan *vlan-id***

Verifying Multiple Ethernet Interfaces for cCPE Services

Purpose Verify multiple Ethernet interfaces for cCPE services.

Action

- To display the ARP table associated with the Ethernet interface:

  user@host> **show arp interface *interface-name***

- To display the bridge domain configuration:

  user@host> **show bridge domain *domain-name***

- To display the route table of the VPN routing instance:

  user@host> **show route table *name.inet.0***

Related Documentation

- *Configuring MEI with cCPE Services—Extending Subscriber VLANs to the PE Router (One Bridge Domain per VLAN)*
- *Configuring MEI with cCPE Services—Extending Subscriber VLANs to the PE Router (All VLANs in One Bridge Domain)*
- *Configuring MEI with cCPE Services—Extending Subscriber VLANs to the PE Router Using a Q-in-Q Tunnel*
- *Configuring MEI and RVIs with cCPE Services*
- Using Multiple Ethernet Interfaces with CCPE Services on page 83

Verifying Multiple Ethernet Interfaces with RVIs for cCPE Services

Purpose View multiple Ethernet interfaces with RVIs for cCPE services.

Action

- To display the ARP table associated with the Ethernet interface:

  user@host> **show arp interface *interface-name***

- To display the bridge domain configuration:

  user@host> **show bridge domain *domain-name***

- To display the route table of the VPN routing instance:

  user@host> **show route table *name.inet.0***

- To display the RVIs and their current states in EX Series switch:

  user@host> **show interfaces vlan *vlan-id***

Related Documentation

- Using Multiple Ethernet Interfaces with CCPE Services on page 83
- *Configuring MEI with cCPE Services—Extending Subscriber VLANs to the PE Router (One Bridge Domain per VLAN)*
- *Configuring MEI with cCPE Services—Extending Subscriber VLANs to the PE Router (All VLANs in One Bridge Domain)*
- *Configuring MEI with cCPE Services—Extending Subscriber VLANs to the PE Router Using a Q-in-Q Tunnel*
- *Configuring MEI and RVIs with cCPE Services*
- Verifying Multiple Ethernet Interfaces for cCPE Services on page 89

CHAPTER 10

Configuring Routed Internet Traffic Through a Subscriber-Owned NAT Gateway for the cCPE Application

Understanding How to Use cCPE Services to Route Internet Traffic to a Subscriber-Owned NAT Gateway

This topic describes how you can migrate from a scenario where the customer site requires a router to forward Internet traffic to their NAT device, to a scenario where only a Layer 2 CPE is required and the MX Series router routes the traffic to the Internet.

In the scenario depicted in Figure 7 on page 110, the NAT functions are provided by a subscriber-owned device with Layer 3 capabilities. Two logical interfaces are configured between the subscriber site and the PE router. One logical interface is for VPN internal traffic, and the other logical interface is for public Internet traffic. For Internet-bound traffic, the subscriber CE router has a route defined in the VPN routing instance, which forwards the traffic to the subscriber's NAT device through the VPN internal interface. After address translation, the Internet-bound traffic is converted to a public address and sent through the public interface. In the PE router, the VPN internal interface is defined in the VPN routing instance, and the public interface is defined in the default routing instance.

Figure 7: Routing Internet Traffic Through a Subscriber NAT Device

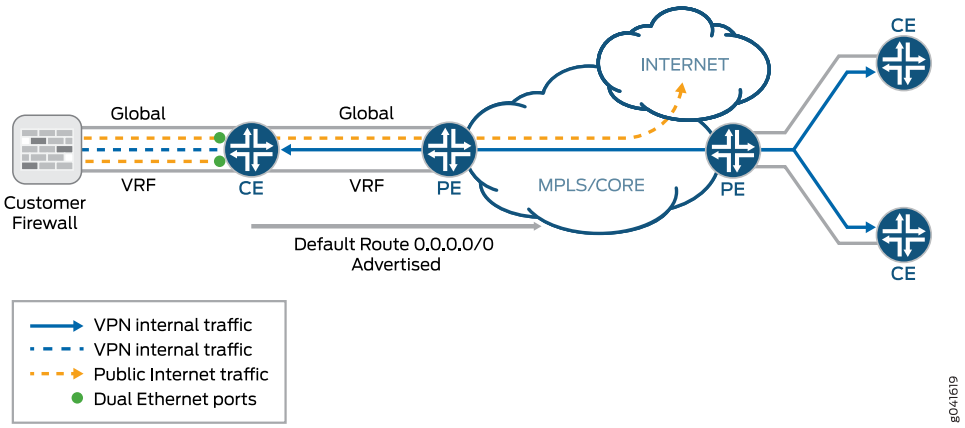

When cCPE services are introduced into this same scenario, the CE router at the subscriber site is replaced with a Layer 2 CPE and the routing is moved out to the PE router. Two VLAN interfaces are connected to the Layer 2 CPE: one as a VPN internal interface and one as a public interface for Internet traffic.

In the VPN site with Internet access, there is a NAT gateway that performs address translation between private and public addresses. The inside interface of the NAT gateway on the LAN side has a private address. In the VPN routing instance in the adjacent PE router, a static route is configured, which sends Internet-bound traffic to the private address of the inside interface of NAT gateway. This static route is further propagated to the VPN in the remote sites through an internal BGP (IBGP) session. The WAN side of the NAT gateway also has a route defined with the adjacent PE router as the next hop for Internet-bound traffic. Packets going to the Internet are sent to the inside interface of the NAT gateway, where the addresses are translated to public address and then routed to the PE router.

In the VPN site without Internet access, you do not need to configure a static route in the VPN routing instance. This site relies on the static route propagated through the BGP routing protocol from the VPN routing instance with Internet access.

Related Documentation

- *Configuring Routed Internet Traffic Through a Subscriber-Owned NAT Gateway for cCPE Services*
- Configuring the PE1 Router for Routed Internet Traffic Through a Subscriber-Owned NAT Gateway on page 113
- Configuring the PE2 Router for Routed Internet Traffic Through a Subscriber-Owned NAT Gateway on page 115
- Configuring the Layer 2 CPE at Site 1 for Routed Internet Traffic Through a Subscriber-Owned NAT Gateway on page 111
- Understanding How to Run Carrier-Grade NAT (CGN) cCPE Services to Route Subscriber Internet Traffic on page 118

Configuring Internet Access with VPNs Using CPE-Based Dual Ethernet (NAT Functions Provided by Subscriber-Owned Gateway)

This topic describes how to configure a VPN subscriber called acme, with two sites. The subnet for site 1 is 192.168.1.0/24 and the subnet for site 2 is 192.168.2.0/24. Site 1, which is connected to PE1, has Internet access. The NAT gateway is located in site 1 with address 192.168.1.2. A static route is added to the VPN routing instance to send Internet traffic to the NAT gateway. Site 1 has two VLAN interfaces into the PE1 router: VLAN 105 is for VPN internal traffic and belongs to the VPN routing instance. VLAN 106 is for public Internet traffic and is terminated in the global routing instance in PE1.

Internet access from site 2 travels through the NAT gateway in site 1. The static route configured in the VPN routing instance of site 1 is propagated to the VPN routing instance at site 2 through BGP.

This procedure requires the following configuration:

1. In PE1, configure a bridge domain, IRB and a VLAN interface for the VPN internal interface.
2. In PE1, configure a VPN routing instance that includes the IRB interface.
3. In PE1, configure the second VLAN interface for the public interface for the subscriber.
4. Optionally, configure a static or dynamic route for internet access from the site.
5. Optionally, configure dynamic routing between the PE and CE for VPN internal routes.
6. In PE2, configure the bridge domain, IRB and VLAN interface
7. In PE2, configure the VPN routing instance that includes the IRB interface you configured on PE2.
8. In the Layer 2 CPE, configure two VLAN interfaces: one as VPN internal interface and one as a Internet public interface.

Configuring the Layer 2 CPE at Site 1 for Routed Internet Traffic Through a Subscriber-Owned NAT Gateway

NOTE: This procedure uses a Juniper Networks EX Series Ethernet Switch for the Layer 2 CPE.

To configure the Layer 2 CPE:

1. Configure the first Ethernet interface on the CPE — ge-0/0/8 is the VPN Internal interface.

   ```
   [edit]
   user@host# set interfaces ge-0/0/8 unit 0 family ethernet-switching port-mode
     access
   ```

2. Configure the second Ethernet interface — ge-0/0/9 is the public Internet interface connecting to the WAN link of the subscriber's NAT gateway.

   ```
   [edit]
   user@host# set interfaces ge-0/0/9 unit 0 family ethernet-switching port-mode
     access
   ```

3. Configure the VLAN trunk port for the WAN link — ge-0/0/10 is a trunk port connecting to the PE router.

   ```
   [edit]
   user@host# set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode
     trunk
   ```

4. Configure the first VLAN for VPN internal traffic.

   a. Specify the VLAN name.

      ```
      [edit]
      user@host# edit vlans vlan-vpn
      ```

   b. Specify the VLAN ID.

      ```
      [edit]
      user@host# set vlan-id 105
      ```

   c. Specify the interface associated with the VLAN.

      ```
      [edit vlans vlan-vpn]
      user@host# set interface ge-0/0/8.0
      ```

   d. Specify the trunk port interface associated with the VLAN.

      ```
      [edit vlans vlan-vpn]
      user@host# set interface ge-0/0/10.0
      ```

5. Configure the second VLAN for public Internet traffic.

   a. Specify the VLAN name.

      ```
      [edit vlans vlan-vpn]
      user@host# top
      user@host# edit vlans vlan-public
      ```

   b. Specify the VLAN ID.

      ```
      [edit vlans vlan-public]
      user@host# set vlan-id 106
      ```

   c. Specify the interface associated with the VLAN.

      ```
      [edit vlans vlan-public ]
      user@host# set interface ge-0/0/9.0
      ```

d. Specify the trunk port interface associated with the VLAN.

```
[edit vlans vlan-public ]
user@host# set interface ge-0/0/10.0
```

Configuring the PE1 Router for Routed Internet Traffic Through a Subscriber-Owned NAT Gateway

This topic describes how to configure the PE1 router for routed Internet traffic through a subscriber-owned NAT device. This configuration uses CPE-based dual-Ethernet. Complete the following tasks to configure PE1:

1. Configuring the Subscriber VLANs — Routed Internet Traffic Through a Subscriber NAT Device on page 113
2. Configuring the IRB Interface, Bridge Domain, and Routing Instance on page 114

Configuring the Subscriber VLANs — Routed Internet Traffic Through a Subscriber NAT Device

To configure the subscriber VLANs:

1. Configure the physical interface for VLAN tagging and flexible Ethernet services encapsulation.

```
[edit]
user@host# edit interfaces ge-1/2/3
user@host# set vlan-tagging
user@host# set encapsulation flexible-ethernet-services
```

2. Set up the first VLAN for internal VPN traffic.

   a. Configure the logical interface.

```
[edit interfaces ge-1/2/3]
user@host# edit unit 105
```

   b. Configure the encapsulation for the logical interface.

```
[edit interfaces ge-1/2/3 unit 105 ]
user@host# set encapsulation vlan-bridge
```

   c. Bind an 802.1Q VLAN tag ID to the logical interface.

```
[edit interfaces ge-1/2/3 unit 105]
user@host# set vlan-id 105
```

   d. Configure the logical interface for Layer 2 bridging.

```
[edit interfaces ge-1/2/3 unit 105]
user@host# set family bridge
```

3. Set up the second VLAN as the Internet public interface. This public interface belongs to the global routing instance. You can configure dynamic or static routing between this interface and the WAN interface at the subscriber site.

   a. Configure the second logical interface.

```
[edit interfaces ge-1/2/3 unit 105]
user@host# up
```



```
[edit interfaces ge-1/2/3]
user@host# edit unit 106
```

b. Configure the encapsulation for the logical interface.

```
[edit interfaces ge-1/2/3 unit 106 ]
user@host# set encapsulation vlan-bridge
```

c. Bind an 802.1Q VLAN tag ID to the logical interface.

```
[edit interfaces ge-1/2/3 unit 106]
user@host# set vlan-id 106
```

d. Configure the logical interface for IPv4 protocol (**inet**) and configure the IP address of the WAN link.

```
[edit interfaces ge-1/2/3 unit 106]
user@host# set family inet address IP prefix of WAN link
```

4. Review the configuration.

```
[edit interfaces ge-1/2/3 unit 106]
user@host# up
[edit interfaces ge-1/2/3]
```

```
unit 105 {
    encapsulation vlan-bridge;
    vlan-id 105;
    family bridge;
}
unit 106 {
    encapsulation vlan-bridge;
    vlan-id 106;
    family inet {
        address 192.1.1.1/32;
    }
}
```

Configuring the IRB Interface, Bridge Domain, and Routing Instance

To configure the IRB interface, bridge domain, and routing instance:

1. Configure the IRB interface.

   a. Configure the logical interface used for internal VPN traffic as the IRB interface.

```
[edit]
user@host# edit interfaces irb unit 105
```

   b. Specify the private subnet of the VPN site on the IRB interface.

```
[edit interfaces irb unit 105]
user@host# set family inet address 192.168.2.1/24
```

2. Configure the bridge domain.

   a. Configure the bridge domain name.

```
[edit]
user@host# edit bridge-domains bd-105
```

   b. Associate the subscriber's VLAN ID with the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set vlan-id 105
```

c. Specify the interface to include in the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set interface ge-1/2/3.105
```

d. Specify the routing interface to include in the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set routing-interface irb.105
```

3. Configure the routing instance.

a. Configure the name of the routing instance.

```
[edit]
user@host# edit routing-instances acme
```

b. Configure the routing instance as VRF.

```
[edit routing-instances acme]
user@host# set instance-type vrf
```

c. Reference the IRB interface by specifying irb.*vlan-id*.

```
[edit routing-instances acme]
user@host# set interface irb.105
```

d. Specify a route distinguisher attached to the route, enabling you to distinguish which VPN the route belongs to. Each routing instance must have a unique route distinguisher associated with it. The route distinguisher is used to place boundaries around a VPN so that the same IP address prefixes can be used in different VPNs without having them overlap. The format for the route distinguisher is *as-number:id*.

```
[edit routing-instances acme]
user@host# set route-distinguisher 65535:1
```

e. Specify the VPN's community. VRF import and export policies are automatically generated.

```
[edit routing-instances acme]
user@host# set vrf-target target:65535:5
```

f. Map the inner label of a packet to a specific VRF table. This enables examination of the encapsulated IP header.

```
[edit routing-instances acme]
user@host# set vrf-table-label
```

Configuring the PE2 Router for Routed Internet Traffic Through a Subscriber-Owned NAT Gateway

Complete the following tasks to configure PE2:

1. Configuring the Subscriber VLAN on page 116
2. Configuring the IRB Interface, Bridge Domain, and Routing Instance on page 116

Configuring the Subscriber VLAN

To configure the subscriber VLANs:

1. Configure the physical interface for VLAN tagging and flexible Ethernet services encapsulation.

   ```
   [edit]
   user@host# edit interfaces ge-1/2/3
   [edit interfaces ge-1/2/3]
   user@host# set vlan-tagging
   user@host# set encapsulation flexible-ethernet-services
   ```

2. Set up the VLAN for internal VPN traffic.

   a. Configure the logical interface.

      ```
      [edit interfaces ge-1/2/3]
      user@host# edit unit 105
      ```

   b. Configure the encapsulation for the logical interface.

      ```
      [edit interfaces ge-1/2/3 unit 105 ]
      user@host# set encapsulation vlan-bridge
      ```

   c. Bind an 802.1Q VLAN tag ID to the logical interface.

      ```
      [edit interfaces ge-1/2/3 unit 105]
      user@host# set vlan-id105
      ```

   d. Configure the logical interface for Layer 2 bridging.

      ```
      [edit interfaces ge-1/2/3 unit 105]
      user@host# set family bridge
      ```

Configuring the IRB Interface, Bridge Domain, and Routing Instance

To configure the IRB interface, bridge domain, and routing instance:

1. Configure the IRB interface.

   a. Configure the logical interface used for internal VPN traffic as the IRB interface.

      ```
      [edit]
      user@host# edit interfaces irb unit 105
      ```

   b. Specify the private subnet of the VPN site on the IRB interface.

      ```
      [edit interfaces irb unit 105]
      user@host# set family inet address 192.168.2.1/24
      ```

2. Configure the bridge domain.

   a. Configure the bridge domain name.

      ```
      [edit]
      user@host# edit bridge-domains bd-105
      ```

   b. Associate the subscriber's VLAN ID with the bridge domain.

      ```
      [edit bridge-domains bd-105]
      user@host# set vlan-id 105
      ```

c. Specify the interface to include in the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set interface ge-1/2/3.105
```

d. Specify the routing interface to include in the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set routing-interface irb.105
```

3. Configure the routing instance.

a. Configure the name of the routing instance.

```
[edit]
user@host# edit routing-instances acme
```

b. Configure the routing instance as VRF.

```
[edit routing-instances acme]
user@host# set instance-type vrf
```

c. Reference the IRB interface by specifying irb.*vlan-id*.

```
[edit routing-instances acme]
user@host# set interface irb.105
```

d. Specify a route distinguisher attached to the route, enabling you to distinguish which VPN the route belongs to. Each routing instance must have a unique route distinguisher associated with it. The route distinguisher is used to place boundaries around a VPN so that the same IP address prefixes can be used in different VPNs without having them overlap. The format for the route distinguisher is *as-number:id*.

```
[edit routing-instances acme]
user@host# set route-distinguisher 65535:1
```

e. Specify the VPN's community. VRF import and export policies are automatically generated.

```
[edit routing-instances acme]
user@host# set vrf-target target:65535:5
```

f. Map the inner label of a packet to a specific VRF table. This enables examination of the encapsulated IP header.

```
[edit routing-instances acme]
user@host# set vrf-table-label
```

Verifying Routed Internet Traffic for NAT for cCPE Services

Purpose Verify that the subscriber's routing instance has a default route with a next hop of the inside service interface.

Action

- Display the subscriber's routing instance and verify that it has a default route with a next hop of the inside service interface:

```
user@host> show route table table.inet.0
```

Related Documentation

- *Configuring a Bridge Domain*
- *Examples: Configuring Static Routes*
- *Filtering Packets in Layer 3 VPNs Based on IP Headers*
- *Routing Instances Overview*

Understanding How to Run Carrier-Grade NAT (CGN) cCPE Services to Route Subscriber Internet Traffic

This topic describes how you can migrate from a scenario where the NAT function is performed by a Layer 3 CPE device to a scenario where the Layer 3 CPE is replaced by a Layer 2 CPE and the MX Series router provides carrier-grade NAT (CGN). The Junos OS provides CGN for IPv4 and IPv6 networks.

In the scenario shown in Figure 8 on page 118, the NAT function is performed by a Layer 3 CPE provided by the service provider. In the Layer 3 CPE, a default route or route for public prefixes is configured to send Internet-bound traffic to the inside interface of NAT. Routes for private address prefixes, configured statically or learned through routing protocols, are used to send VPN internal traffic to the VPN interfaces (the subscriber LAN or the VPN internal interface to the PE router).

If dynamic routing is enabled between the CE router and the adjacent PE router, the default route configured in the CE router is advertised to the VPN routing instance in the PE router, which further propagates it to other VPN sites through IBGP. The default route in the remote VPN sites, sends Internet-bound traffic to the NAT inside interface of the CE router through the VPN. The addresses are translated, and the traffic is sent back to the PE router through the public interface. Internet-bound traffic from the remote site travels through the WAN link of the VPN site with Internet access twice: First from the PE router to the CE router and then from the CE router back to the PE router after address translation by NAT. The process is the same for traffic from the Internet.

Figure 8: Routing Internet Traffic Through Carrier-Grade NAT

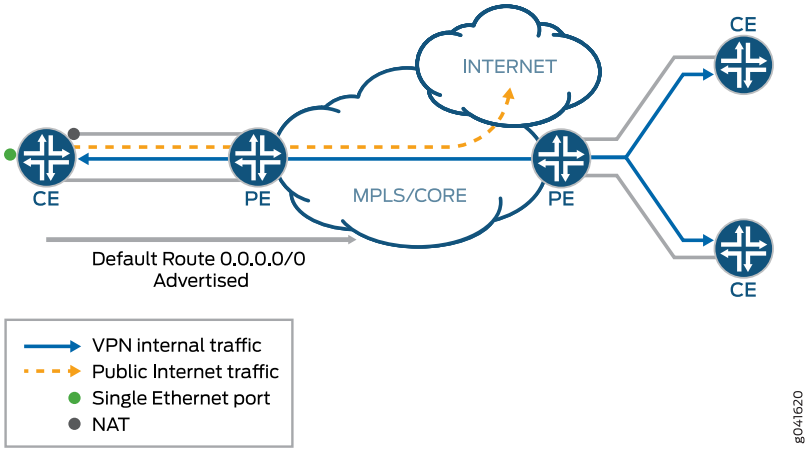

There is one interface from the CE router to the subscriber LAN, and two interfaces between the CE router and the PE router: one for VPN internal traffic and one for Internet

traffic. The VPN internal interface belongs to the VPN routing instance in the PE router, and the public interface is terminated in the global routing instance in the PE router.

When you implement cCPE services into this same scenario, the Layer 3 CPE is replaced by a Layer 2 CPE. Layer 3 functions, including routing, NAT, and firewall, are provided as cloud services by the MX Series router. In the VPN routing instance in the PE router, a static route is needed to send Internet-bound traffic to the inside interface of the NAT gateway and the firewall. After address translation, traffic travels through the global routing instance and is then routed to the Internet. Return traffic travels back through the global routing instance and is then routed to the outside interface of the NAT gateway and firewall. After the traffic is converted to the private destination address, the packets in the routing instance are routed within the VPN.

Implementing cCPE services in this scenario provides two benefits:

- Internet traffic for remote sites (sites without Internet access) does not need to go through the WAN link twice for address translation by NAT.
- Only one interface is required between the PE router and the CE router.

Related Documentation

- Understanding How to Use cCPE Services to Route Internet Traffic to a Subscriber-Owned NAT Gateway on page 109
- *Configuring Routed Internet Traffic Through a Subscriber-Owned NAT Gateway for cCPE Services*
- *Configuring Routed Internet Traffic and MX Series Router NAT Functions for cCPE Services*

Configuring Internet Access for VPN Subscribers Using CGN cCPE Services

This procedure describes how to configure Internet access for VPN subscribers with only a single Ethernet LAN. The MX Series router at the provider edge provides the NAT service. A default route is added to the VPN routing instance to send Internet traffic to the NAT inside interface. Site 1 has only one VLAN interface into the PE1 router with VLAN ID 105.

Internet access from site 2 goes through the NAT gateway in site 1. The default route configured in the site 1 VPN routing instance is propagated to the site 2 VPN routing instance through BGP.

This procedure requires the following configuration in PE1, see "Configuring the PE1 Router for Routed Internet Traffic and MX Series Router NAT Functions" on page 120:

1. Configure a bridge domain, IRB, and a VLAN interface as the VPN internal interface.
2. Configure a VPN routing instance that includes the IRB interface.
3. Configure two service interfaces for NAT: one as the inside interface and one as the outside interface.

4. Configure a next-hop style service set that includes the NAT server. If required, stateful firewall service can be added to the service set.
5. Create a default route to send Internet traffic to the inside interface of above service set.

The PE2 requires the following configuration, see "Configuring the PE2 Router for Routed Internet Traffic and MX Series Router NAT Functions" on page 124:

1. Configure a bridge domain, IRB, and VLAN interface.
2. Configure a VPN routing instance that includes the IRB interface.

Configuring the PE1 Router for Routed Internet Traffic and MX Series Router NAT Functions

Complete the following tasks to configure the PE1 router at Site 1:

Configuring NAT on MX Series Routers for cCPE Services

The Junos OS provides carrier-grade NAT (CGN) for IPv4 and IPv6 networks. To enable the NAT service for the cCPE services:

1. Configure the properties for the MS-DPC and enable the extension provider service package application.

   ```
   [edit]
   edit chassis fpc 4 pic 0 adaptive-services service-package layer-3
   ```

2. Verify the configuration.

   ```
   user@host# show
   chassis {
   fpc 4 {
       pic 0 {
           adaptive-services {
               service-package layer-3;
           }
       }
   }
   }
   ```

Configuring the Service Interfaces for NAT

To configure the inside and outside service interfaces:

1. Configure the inside interface.

   ```
   [edit]
   user@host# set interfaces sp-2/0/0 unit 1 family inet
   user@host# set interfaces sp-2/0/0 unit 1 service-domain inside
   ```

2. Configure the outside interface.

   ```
   [edit]
   user@host# set interfaces sp-2/0/0 unit 2 family inet
   user@host# set interfaces sp-2/0/0 unit 2 service-domain outside
   ```

Defining the Service Rules

In this sample procedure, the service set contains only the NAT service. You can also add other services like stateful firewall. To define the service rules to be applied to traffic:

1. Configure the next-hop service set.

   a. Configure a name for the service set.

      ```
      [edit]
      user@host# edit services service-set acme-services
      ```

   b. Configure a name for the NAT rules.

      ```
      [edit services service-set acme-services]
      user@host# set nat-rules acme-nat-rule
      ```

   c. Define the next-hop service for the inside service interface.

      ```
      [edit services service-set acme-services]
      user@host# set next-hop-service inside-service-interface sp-2/0/0.1
      ```

   d. Define the next-hop service for the outside service interface.

      ```
      [edit services service-set acme-services]
      user@host# set next-hop-service outside-service-interface sp-2/0/0.2
      ```

2. Configure the public address pool and ports.

   a. Configure the public address pool name.

      ```
      user@host# top
      [edit]
      user@host# edit services nat pool acme-public-pool
      ```

   b. Specify the address or address prefix for NAT.

      ```
      [edit services nat pool acme-public-pool]
      user@host# set address public-address
      ```

   c. Configure the NAT port to be assigned automatically by the router.

      ```
      [edit services nat pool acme-public-pool]
      user@host# set port automatic
      ```

3. Configure the NAT rules.

a. Specify the name of the NAT rule.

```
user@host# top
[edit]
user@host# edit services nat rule acme-nat-rule
```

b. Specify the direction in which the rule match is applied.

```
[edit services nat rule acme-nat-rule]
user@host# set match-direction input
```

c. Define the NAT term actions.

```
[edit services nat rule acme-nat-rule]
user@host# set term translate then translated source-pool acme-public-pool
user@host# set translation-type napt-44
```

Configuring the Interface, Bridge Domain, and IRB Interface

To configure the interface, bridge domain, and IRB interface:

1. Configure the interface for VLAN tagging and flexible Ethernet services encapsulation.

```
[edit]
user@host# edit interfaces ge-1/2/3
user@host# set vlan-tagging
user@host#  set encapsulation flexible-ethernet-services
```

2. Configure the subscriber VLAN.

   a. Configure the logical interface.

```
[edit interfaces ge-1/2/3]
user@host# edit unit 105
```

   b. Configure the encapsulation for the logical interface.

```
[edit interfaces ge-1/2/3 unit 105 ]
user@host# set encapsulation vlan-bridge
```

   c. Bind an 802.1Q VLAN tag ID to the logical interface.

```
[edit interfaces ge-1/2/3 unit 105]
user@host# set vlan-id 105
```

   d. Configure the logical interface for Layer 2 bridging.

```
[edit interfaces ge-1/2/3 unit 105]
user@host# set family bridge
```

3. Configure the IRB interface.

   a. Configure the logical interface used for internal VPN traffic as the IRB interface.

```
user@host# top
[edit]
user@host# set interfaces irb unit 105
```

   b. Specify the private subnet of the VPN site on the IRB interface.

```
[edit interfaces irb unit 105]
user@host# set family inet address 192.168.1.1/24
```

4. Configure the bridge domain.

   a. Configure the bridge domain name.

      ```
      user@host# top
      [edit]
      user@host# edit bridge-domains bd-105
      ```

   b. Associate the subscriber's VLAN ID with the bridge domain.

      ```
      [edit bridge-domains bd-105]
      user@host# set vlan-id 105
      ```

   c. Specify the interface name for the bridge domain.

      ```
      [edit bridge-domains bd-105]
      user@host# set interface ge-1/2/3.105
      ```

   d. Specify the routing interface to include in the bridge domain.

      ```
      [edit bridge-domains bd-105]
      user@host# set routing-interface irb.105
      ```

5. Configure the routing instance.

   a. Configure the name of the routing instance.

      ```
      user@host# top
      [edit]
      user@host# edit routing-instances acme
      ```

   b. Configure the routing instance as a VRF instance.

      ```
      [edit routing-instances acme]
      user@host# set instance-type vrf
      ```

   c. Reference the IRB interface by specifying irb.*vlan-id*.

      ```
      [edit routing-instances acme]
      user@host# set interface irb.105
      ```

   d. Add the NAT inside interface to the routing instance by specifying the inside interface name and route distinguisher attached to the route.

      ```
      [edit routing-instances acme]
      user@host# set interface sp-2/0/0.1
      user@host# set route-distinguisher 65535:1
      ```

   e. Specify the VPN's community. VRF import and export policies are automatically generated.

      ```
      [edit routing-instances acme]
      user@host# set vrf-target target:65535:5
      ```

   f. Map the inner label of a packet to a specific VRF table. This enables examination of the encapsulated IP header.

      ```
      [edit routing-instances acme]
      user@host# set vrf-table-label
      ```

   g. Add a static route to send Internet traffic to the inside interface of NAT service.

      ```
      [edit routing-instances acme]
      user@host# set routing-options static route 0.0.0.0/0 next-hop sp-2/0/0.1
      ```


NOTE: You must advertise this route to remote PE routers through a VPN export policy.

Configuring the PE2 Router for Routed Internet Traffic and MX Series Router NAT Functions

PE2 is the adjacent PE router of the VPN site without Internet access. Complete the following tasks to configure PE2:

1. Configuring the Subscriber VLAN on page 124
2. Configuring the IRB Interface, Bridge Domain, and Routing Instance on page 124

Configuring the Subscriber VLAN

To configure the subscriber VLANs:

1. Configure the physical interface for VLAN tagging and flexible Ethernet services encapsulation.

   ```
   [edit]
   user@host# edit interfaces ge-1/2/3
   [edit interfaces ge-1/2/3]
   user@host# set vlan-tagging
   user@host# set encapsulation flexible-ethernet-services
   ```

2. Set up the VLAN for internal VPN traffic.

   a. Configure the logical interface.

      ```
      [edit interfaces ge-1/2/3]
      user@host# edit unit 105
      ```

   b. Configure the encapsulation for the logical interface.

      ```
      [edit interfaces ge-1/2/3 unit 105 ]
      user@host# set encapsulation vlan-bridge
      ```

   c. Bind an 802.1Q VLAN tag ID to the logical interface.

      ```
      [edit interfaces ge-1/2/3 unit 105]
      user@host# set vlan-id105
      ```

   d. Configure the logical interface for Layer 2 bridging.

      ```
      [edit interfaces ge-1/2/3 unit 105]
      user@host# set family bridge
      ```

Configuring the IRB Interface, Bridge Domain, and Routing Instance

To configure the IRB interface, bridge domain, and routing instance:

1. Configure the IRB interface.

   a. Configure the logical interface used for internal VPN traffic as the IRB interface.

      ```
      [edit]
      user@host# edit interfaces irb unit 105
      ```

b. Specify the private subnet of the VPN site on the IRB interface.

```
[edit interfaces irb unit 105]
user@host# set family inet address 192.168.2.1/24
```

2. Configure the bridge domain.

a. Configure the bridge domain name.

```
[edit]
user@host# edit bridge-domains bd-105
```

b. Associate the subscriber's VLAN ID with the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set vlan-id 105
```

c. Specify the interface to include in the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set interface ge-1/2/3.105
```

d. Specify the routing interface to include in the bridge domain.

```
[edit bridge-domains bd-105]
user@host# set routing-interface irb.105
```

3. Configure the routing instance.

a. Configure the name of the routing instance.

```
[edit]
user@host# edit routing-instances acme
```

b. Configure the routing instance as VRF.

```
[edit routing-instances acme]
user@host# set instance-type vrf
```

c. Reference the IRB interface by specifying irb.*vlan-id*.

```
[edit routing-instances acme]
user@host# set interface irb.105
```

d. Specify a route distinguisher attached to the route, enabling you to distinguish which VPN the route belongs to. Each routing instance must have a unique route distinguisher associated with it. The route distinguisher is used to place boundaries around a VPN so that the same IP address prefixes can be used in different VPNs without having them overlap. The format for the route distinguisher is *as-number:id*.

```
[edit routing-instances acme]
user@host# set route-distinguisher 65535:1
```

e. Specify the VPN's community. VRF import and export policies are automatically generated.

```
[edit routing-instances acme]
user@host# set vrf-target target:65535:5
```

f. Map the inner label of a packet to a specific VRF table. This enables examination of the encapsulated IP header.

```
[edit routing-instances acme]
```


user@host# **set vrf-table-label**

Verifying Routed Internet Traffic for NAT for cCPE Services

Purpose Verify that the subscriber's routing instance has a default route with a next hop of the inside service interface.

Action

- Display the subscriber's routing instance and verify that it has a default route with a next hop of the inside service interface:

  user@host> **show route table *table*.inet.0**

Related Documentation

- Understanding How to Run Carrier-Grade NAT (CGN) cCPE Services to Route Subscriber Internet Traffic on page 118
- Understanding How to Use cCPE Services to Route Internet Traffic to a Subscriber-Owned NAT Gateway on page 109
- Configuring Internet Access with VPNs Using CPE-Based Dual Ethernet (NAT Functions Provided by Subscriber-Owned Gateway) on page 111
- *Configuring Control and Data Cores*
- *Configuring the Junos OS to Enable Service Packages on Adaptive Services Interfaces*
- *Configuring Packages on the PIC*
- *Configuring Memory Settings*
- *Configuring Service Rules*
- *Service Set Properties*
- *Configuring the Address and Domain for Services Interfaces*
- *Configuring Service Sets to be Applied to Services Interfaces*
- *Configuring Pools of Addresses and Ports for Network Address Translation Overview*
- *Network Address Translation Rules Overview*
- *Layer 2 Bridging Interfaces Overview*
- *Configuring Layer 2 Bridging Interfaces*
- *Configuring a Bridge Domain*
- *Examples: Configuring Static Routes*
- *Filtering Packets in Layer 3 VPNs Based on IP Headers*
- *Routing Instances Overview*
- *Examples: Configuring Static Routes*

Verifying Routed Internet Traffic for NAT for cCPE Services

Purpose Verify that the subscriber's routing instance has a default route with a next hop of the inside service interface.

Action

- Display the subscriber's routing instance and verify that it has a default route with a next hop of the inside service interface:

  user@host> **show route table *table*.inet.0**

Related Documentation

- *Configuring Routed Internet Traffic Through a Subscriber-Owned NAT Gateway for cCPE Services*
- *Configuring Routed Internet Traffic and MX Series Router NAT Functions for cCPE Services*
- Understanding How to Use cCPE Services to Route Internet Traffic to a Subscriber-Owned NAT Gateway on page 109
- Understanding How to Run Carrier-Grade NAT (CGN) cCPE Services to Route Subscriber Internet Traffic on page 118

CHAPTER 11

Configuring Draft-Rosen Multicast VPNs with cCPE Services for the cCPE Application

Using Draft-Rosen Multicast VPNs with cCPE Services

Junos OS supports Layer 3 VPNs based on the Internet draft draft-rosen-rfc2547bis, BGP/MPLS VPNs. This Internet draft defines a mechanism by which service providers can use their IP backbones to provide Layer 3 VPN services to their customers. The sites that encompass a Layer 3 VPN are connected over the provider's existing public Internet backbone.

In a unicast environment for Layer 3 VPNs, all VPN states are contained within the PE routers. With multicast over Layer 3 VPNs, two PIM adjacencies are established: One between the CE router and the PE router through a VRF routing instance, the second between the main PE routers and their service provider core neighbors. The set of master PIM adjacencies throughout the service provider's network provides the forwarding paths, and eventually forms a rendezvous point (RP) multicast distribution tree. The tree is rooted at the RP contained within the service provider's network. As a result, core provider transit routers within the service provider's network must maintain multicast state information for the VPNs.

For draft-rosen multicast virtual private networks (MVPNs) to work correctly, there must be two types of rendezvous points. The VPN customer rendezvous point (VPN C-RP) is an RP that resides within a VPN, which connects the segments of a customer network. The service provider rendezvous point (SP-RP) resides within the service provider network itself. Because a PE router connects to both the customer network and the service provider network, the PE router acts as an SP-RP, a VPN C-RP, or both.

In MVPNs, the Layer 3 CPE must run the PIM protocol to form the multicast distribution tree in the customer network. Typically, C-RP is running in the Layer 3 CPE. Optionally, you can enable the bootstrap router (BSR), or auto-RP, to support C-RP redundancy. Moving to the cCPE architecture, the Layer 3 CPE is replaced with a Layer 2 CPE. IP multicast traffic can pass through the Layer 2 CPE, which replicates the multicast packets to all interfaces if IGMP snooping is not enabled. C-RP and BSR, or auto-RP functions, are moved to the cCPE architecture (the VRF routing instance in the adjacent PE router). The MX Series router supports C-RP, BSR, and auto-RP in VRF routing instances. If there is more than one Layer 2 interface in the cCPE bridge domain, IGMP snooping can be enabled in the bridge domain to ensure that IP multicast packets are replicated only to interfaces with receivers.

Related Documentation

- *Draft-Rosen Multicast VPNs Overview*
- Configuring Draft-Rosen Multicast VPNs with cCPE Services on page 131
- *Configuring Draft-Rosen Multicast VPNs*
- *Understanding Data MDTs*

Configuring Draft-Rosen Multicast VPNs with cCPE Services

This topic describes how to configure draft-rosen multicast VPNs in the cCPE environment. For this configuration, you need to configure PIM in sparse mode (PIM-SM) in the customer network. To ensure rendezvous point (RP) redundancy, BSR is enabled for RP election and failover. Alternatively, you can enable auto-RP instead of BSR. In the provider network, PIM sparse mode is used to establish a default multicast distribution tree (MDT), which connects all PE routers configured in the same VPN. Each VPN has one default MDT. The default MDT uses PIM Any Source Multicast. Initially, VPN multicast traffic is sent to every PE router in the VPN through the default MDT, including those PE routers that do not have receivers connected. If data MDT is configured (VRF in adjacent PE router), when the multicast traffic exceeds a traffic rate threshold, a data MDT is created. This MDT includes only PE routers with source or receivers of the attached multicast group. The provider network is transparent to customers, and the transition to cCPE has no impact on it.

For this configuration you will:

1. Configure a bridge domain, an IRB interface, and a VLAN interface.
2. Configure an additional unit on the loopback interface of the PE router. Assign an address from the VPN address space.
3. Configure a VPN routing instance that includes the IRB interface and the loopback interface.
4. Configure a bootstrap router and RP in the routing instance. The BSR and RP do not need to be configured in all routing instances of a VPN, and they do not need to reside in the same routing instance.
5. Enable PIM version 2 on the IRB interface and the loopback interface.

Before you begin, make sure you have completed the steps for the cCPE common configuration. See "Configuring the cCPE Common Configuration on MX Series Routers Using the Junos OS CLI" on page 16.

Then complete the following tasks to configure multicast VPNs for cCPE services:

1. Configure PE1. See "Configuring PE1 for Draft-Rosen Multicast VPNs and cCPE Services" on page 132.
2. Configure PE2. See "Configuring PE2 for Draft-Rosen Multicast VPNs and cCPE Services" on page 135.

Related Documentation

- Using Draft-Rosen Multicast VPNs with cCPE Services on page 129
- *Configuring Draft-Rosen Multicast VPNs*

Configuring PE1 for Draft-Rosen Multicast VPNs and cCPE Services

Complete the following tasks to configure the PE1 router for draft-rosen multicast VPNs and cCPE services:

1. Configuring the VLAN Interface, IRB Interface, and Bridge Domain for Draft-Rosen Multicast VPNs and cCPE Services on page 132
2. Configuring the Loopback Interface and Assigning a VPN Private Address on page 133
3. Configuring the Routing Instance on page 133
4. Configuring the BSR and RP in the Routing Instance on page 134
5. Enabling PIM Version 2 on the IRB and Loopback Interfaces on page 134

Configuring the VLAN Interface, IRB Interface, and Bridge Domain for Draft-Rosen Multicast VPNs and cCPE Services

To configure the VLAN interface, IRB interface, and bridge domain:

1. Configure the physical interface for VLAN tagging and flexible Ethernet services encapsulation.

   ```
   [edit]
   user@host# edit interfaces ge-1/0/1
   user@host# set vlan-tagging
   user@host# set encapsulation flexible-ethernet-services
   ```

2. Configure the subscriber VLAN.

   a. Configure the logical interface.

      ```
      [edit interfaces ge-1/0/1]
      user@host# edit unit 106
      ```

   b. Configure the encapsulation for the logical interface.

      ```
      [edit interfaces ge-1/2/3 unit 106]
      user@host# set encapsulation vlan-bridge
      ```

   c. Bind an 802.1Q VLAN tag ID to the logical interface.

      ```
      [edit interfaces ge-1/0/1 unit 106]
      user@host# set vlan-id 106
      ```

   d. Configure the logical interface for Layer 2 bridging.

      ```
      [edit interfaces ge-1/0/1 unit 106]
      user@host# set family bridge
      ```

3. Configure the IRB interface.

   a. Configure the logical interface used for internal VPN traffic as the IRB interface.

      ```
      [edit]
      user@host# edit interfaces irb unit 106
      ```

   b. Specify the private subnet of the VPN site on the IRB interface.

      ```
      [edit interfaces irb unit 106]
      user@host# set family inet address 10.132.12.1/24
      ```

4. Configure the bridge domain.

   a. Configure the bridge domain name and set the domain type to bridge.

      ```
      [edit]
      user@host# edit bridge-domains vlan-106 domain-type bridge
      ```

   b. Associate the subscriber's VLAN ID with the bridge domain.

      ```
      [edit bridge-domains vlan-106]
      user@host# set vlan-id 106
      ```

   c. Specify the interface name for the bridge domain.

      ```
      [edit bridge-domains vlan-106]
      user@host# set interface ge-1/0/1.106
      ```

   d. Specify the routing interface to include in the bridge domain.

      ```
      [edit bridge-domains vlan-106]
      user@host# set routing-interface irb.106
      ```

5. (Optional) If there are multiple Layer 2 interfaces in the same bridge domain, enable IGMP snooping in the bridge domain.

   ```
   [edit bridge-domains vlan-106]
   user@host# set protocols igmp-snooping
   ```

Configuring the Loopback Interface and Assigning a VPN Private Address

Configure an additional logical interface on the loopback interface of the PE router, and assign an IP address from the VPN address space.

- To configure the loopback interface:

  ```
  [edit ]
  user@host# set interfaces lo0 unit 2 family inet address 10.132.12.254/32
  ```

Configuring the Routing Instance

Configure a VPN routing instance that includes the IRB interface and the loopback interface. To configure the routing instance:

1. Configure the name of the routing instance.

   ```
   [edit]
   user@host# edit routing-instances acme
   ```

2. Configure the routing instance as a VRF instance.

   ```
   [edit routing-instances acme]
   user@host# set instance-type vrf
   ```

3. Reference the IRB interface in the routing instance by specifying irb.*vlan-id*.

   ```
   [edit routing-instances acme]
   user@host# set interface irb.106
   ```

4. Reference the loopback interface in the routing instance.

   ```
   [edit routing-instances acme]
   ```

```
user@host# set interface lo0.2
```

5. Specify a route distinguisher that is attached to the route, enabling you to distinguish which VPN the route belongs to. Each routing instance must have a unique route distinguisher associated with it.

```
[edit routing-instances acme]
user@host# set route-distinguisher 65535:4
```

6. Specify the VPN's community. VRF import and export policies are automatically generated.

```
[edit routing-instances acme]
user@host# set vrf-target target:65535:5
```

7. Map the inner label of a packet to a specific VRF table. This enables examination of encapsulated IP headers.

```
[edit routing-instances acme]
user@host# set vrf-table-label
```

8. Configure the group address for the Layer 3 VPN in the service provider's network.

```
[edit routing-instances acme]
user@host# set protocols pim vpn-group-address 239.1.1.1
```

Configuring the BSR and RP in the Routing Instance

To configure the BSR and RP:

1. Configure the BSR. Make sure to set different priorities on the PE routers.

```
[edit routing-instances acme]
user@host# edit protocols pim rp bootstrap family inet priority 2
```

2. Configure a candidate RP in the VRF. Make sure to set different priorities on the PE routers.

```
[edit routing-instances acme protocols pim rp]
user@host# set local family inet address 10.132.12.254
user@host# set local family inet priority 2
```

Enabling PIM Version 2 on the IRB and Loopback Interfaces

To enable PIM on the IRB and loopback interfaces:

- Enable PIM on the IRB interface and configure the PIM mode as sparse.

```
[edit routing-instances acme protocols pim]
user@host# set interface irb.106 mode sparse
user@host# set interface irb.106 version 2
user@host# set interface lo0.2 mode sparse
user@host# set interface lo0.2 version 2
```

Related Documentation

- Using Draft-Rosen Multicast VPNs with cCPE Services on page 129
- Configuring Draft-Rosen Multicast VPNs with cCPE Services on page 131
- *Example: Configuring Draft-Rosen MVPN Interoperability*

- *Configuring Draft-Rosen Multicast VPNs*

Configuring PE2 for Draft-Rosen Multicast VPNs and cCPE Services

Configuring the VLAN Interface, IRB Interface, and Bridge Domain for Draft-Rosen Multicast VPNs and cCPE Services

To configure the VLAN interface, IRB interface, and bridge domain:

1. Configure the physical interface for VLAN tagging and flexible Ethernet services encapsulation.

   ```
   [edit]
   user@host# edit interfaces ge-1/0/1
   user@host# set vlan-tagging
   user@host# set encapsulation flexible-ethernet-services
   ```

2. Configure the subscriber VLAN.

   a. Configure the logical interface.

      ```
      [edit interfaces ge-1/0/1]
      user@host# edit unit 106
      ```

   b. Configure the encapsulation for the logical interface.

      ```
      [edit interfaces ge-1/2/3 unit 106]
      user@host# set encapsulation vlan-bridge
      ```

   c. Bind an 802.1Q VLAN tag ID to the logical interface.

      ```
      [edit interfaces ge-1/0/1 unit 106]
      user@host# set vlan-id 106
      ```

   d. Configure the logical interface for Layer 2 bridging.

      ```
      [edit interfaces ge-1/0/1 unit 106]
      user@host# set family bridge
      ```

3. Configure the IRB interface.

   a. Configure the logical interface used for internal VPN traffic as the IRB interface.

      ```
      [edit]
      user@host# edit interfaces irb unit 106
      ```

   b. Specify the private subnet of the VPN site on the IRB interface.

      ```
      [edit interfaces irb unit 106]
      user@host# set family inet address 10.132.12.1/24
      ```

4. Configure the bridge domain.

   a. Configure the bridge domain name and set the domain type to bridge.

      ```
      [edit]
      user@host# edit bridge-domains vlan-106 domain-type bridge
      ```

   b. Associate the subscriber's VLAN ID with the bridge domain.

      ```
      [edit bridge-domains vlan-106]
      user@host# set vlan-id 106
      ```

   c. Specify the interface name for the bridge domain.

      ```
      [edit bridge-domains vlan-106]
      user@host# set interface ge-1/0/1.106
      ```

   d. Specify the routing interface to include in the bridge domain.

      ```
      [edit bridge-domains vlan-106]
      user@host# set routing-interface irb.106
      ```

5. (Optional) If there are multiple Layer 2 interfaces in the same bridge domain, enable IGMP snooping in the bridge domain.

   ```
   [edit bridge-domains vlan-106]
   user@host# set protocols igmp-snooping
   ```

Configuring the Loopback Interface and Assigning a VPN Private Address

Configure an additional logical interface on the loopback interface of the PE router, and assign an IP address from the VPN address space.

- To configure the loopback interface:

  ```
  [edit ]
  user@host# set interfaces lo0 unit 2 family inet address 10.132.12.254/32
  ```

Configuring the Routing Instance

Configure a VPN routing instance that includes the IRB interface and the loopback interface. To configure the routing instance:

1. Configure the name of the routing instance.

   ```
   [edit]
   user@host# edit routing-instances acme
   ```

2. Configure the routing instance as a VRF instance.

   ```
   [edit routing-instances acme]
   user@host# set instance-type vrf
   ```

3. Reference the IRB interface in the routing instance by specifying irb.*vlan-id*.

   ```
   [edit routing-instances acme]
   user@host# set interface irb.106
   ```

4. Reference the loopback interface in the routing instance.

   ```
   [edit routing-instances acme]
   ```

```
user@host# set interface lo0.2
```

5. Specify a route distinguisher that is attached to the route, enabling you to distinguish which VPN the route belongs to. Each routing instance must have a unique route distinguisher associated with it.

   ```
   [edit routing-instances acme]
   user@host# set route-distinguisher 65535:4
   ```

6. Specify the VPN's community. VRF import and export policies are automatically generated.

   ```
   [edit routing-instances acme]
   user@host# set vrf-target target:65535:5
   ```

7. Map the inner label of a packet to a specific VRF table. This enables examination of encapsulated IP headers.

   ```
   [edit routing-instances acme]
   user@host# set vrf-table-label
   ```

8. Configure the group address for the Layer 3 VPN in the service provider's network.

   ```
   [edit routing-instances acme]
   user@host# set protocols pim vpn-group-address 239.1.1.1
   ```

Configuring the BSR and RP in the Routing Instance

To configure the BSR and RP:

1. Configure the BSR. Make sure to set different priorities on the PE routers.

   ```
   [edit routing-instances acme]
   user@host# edit protocols pim rp bootstrap family inet priority 2
   ```

2. Configure a candidate RP in the VRF. Make sure to set different priorities on the PE routers.

   ```
   [edit routing-instances acme protocols pim rp]
   user@host# set local family inet address 10.132.12.254
   user@host# set local family inet priority 2
   ```

Enabling PIM Version 2 on the IRB and Loopback Interfaces

To enable PIM on the IRB and loopback interfaces:

- Enable PIM on the IRB interface and configure the PIM mode as sparse.

  ```
  [edit routing-instances acme protocols pim]
  user@host# set interface irb.106 mode sparse
  user@host# set interface irb.106 version 2
  user@host# set interface lo0.2 mode sparse
  user@host# set interface lo0.2 version 2
  ```

Related Documentation

- Using Draft-Rosen Multicast VPNs with cCPE Services on page 129
- Configuring Draft-Rosen Multicast VPNs with cCPE Services on page 131
- *Example: Configuring Draft-Rosen MVPN Interoperability*

- *Configuring Draft-Rosen Multicast VPNs*

Verifying Draft-Rosen Multicast-VPNs for cCPE

Purpose View the draft-rosen multicast-VPN configuration for cCPE services.

Action

- To display the C-RP status:

  user@host> **show pim rps extensive instance *vrf***

- To display the PIM bootstrap router status in the C Network:

  user@host> **show pim bootstrap instance *vrf***

- To display the PIM interfaces in the C Network:

  user@host> **show pim interfaces instance *vrf***

- To display the PIM neighbor information in the C Network:

  user@host> **show pim neighbors instance *vrf***

- To display the PIM joins in the C Network:

  user@host> **show pim join instance *vrf***

- To display the multicast forwarding cache entry—active multicast groups, source, and packet counters:

  user@host> **show multicast route instance *vrf* extensive**

- To display the IGMP status on the subscriber-facing interfaces:

  user@host> **show igmp interface interface**

- To display IGMP snooping on the subscriber-facing interfaces:

  user@host> **show igmp snooping membership**
  user@host> **show igmp snooping interface**

Related Documentation

- Using Draft-Rosen Multicast VPNs with cCPE Services on page 129
- Configuring Draft-Rosen Multicast VPNs with cCPE Services on page 131
- *Example: Configuring Draft-Rosen MVPN Interoperability*
- *Configuring Draft-Rosen Multicast VPNs*