



Junos[®] OS

Interprovider and Carrier-of-Carriers Feature Guide for Routing Devices

Release

14.1



Published: 2014-05-09

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

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

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

Junos[®] OS Interprovider and Carrier-of-Carriers Feature Guide for Routing Devices

14.1

Copyright © 2014, Juniper Networks, Inc.

All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	xi
	Documentation and Release Notes	xi
	Supported Platforms	xi
	Using the Examples in This Manual	xi
	Merging a Full Example	xii
	Merging a Snippet	xii
	Documentation Conventions	xiii
	Documentation Feedback	xv
	Requesting Technical Support	xv
	Self-Help Online Tools and Resources	xv
	Opening a Case with JTAC	xvi
Part 1	Overview	
Chapter 1	Introduction to Interprovider and Carrier-of-Carriers VPNs	3
	Traditional VPNs, Interprovider VPNs, and Carrier-of-Carriers VPNs	3
	Standard VPNs	4
	Interprovider and Carrier-of-Carriers VPNs	4
	Interprovider VPNs	5
	Linking VRF Tables Between Autonomous Systems	5
	Configuring Next Generation Layer 3 VPNs Options A, B, and C	5
	Configuring Multihop MP-EBGP Between AS Border Routers	6
	Carrier-of-Carriers VPNs	8
	Internet Service Provider as the Customer	9
	VPN Service Provider as the Customer	9
Part 2	Configuration	
Chapter 2	Configuring Interprovider and Carrier-of-Carriers VPNs	13
	Understanding Interprovider VPNs	13
	Configuring Carrier-of-Carriers VPNs for Customers That Provide Internet Service	14
	Configuring the Carrier-of-Carriers VPN Service Customer's CE Router	14
	Configuring MPLS	14
	Configuring BGP	14
	Configuring OSPF	15
	Configuring Policy Options	15
	Configuring the Carrier-of-Carriers VPN Service Provider's PE Routers	16
	Configuring MPLS	16
	Configuring BGP	16
	Configuring IS-IS	17

Configuring LDP	17
Configuring a Routing Instance	17
Configuring Policy Options	18
Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service	19
Configuring the Carrier-of-Carriers Customer's PE Router	19
Configuring MPLS	19
Configuring BGP	20
Configuring OSPF	20
Configuring LDP	20
Configuring VPN Service in the Routing Instance	21
Configuring Policy Options	21
Configuring the Carrier-of-Carriers Customer's CE Router	22
Configuring MPLS	22
Configuring BGP	22
Configuring OSPF and LDP	23
Configuring Policy Options	24
Configuring the Provider's PE Router	24
Configuring MPLS	24
Configuring a PE-Router-to-PE-Router BGP Session	24
Configuring IS-IS and LDP	25
Configuring Policy Options	25
Configuring a Routing Instance to Send Routes to the CE Router	26
Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics	26
Chapter 3	
Configuration Examples for Interprovider and Carrier-of-Carriers VPNs	29
Interprovider VPN Example—MP-EBGP Between ISP Peer Routers	29
Configuration for Router A	30
Configuration for Router B	30
Configuration for Router C	32
Configuration for Router D	33
Configuration for Router E	34
Configuration for Router F	36
Interprovider VPN Example—Multihop MP-EBGP with P Routers	36
Configuration for Router A	37
Configuration for Router B	37
Configuration for Router C	39
Configuration for Router D	40
Configuration for Router E	41
Configuration for Router F	43
Carrier-of-Carriers VPN Examples	43
Carrier-of-Carriers VPN Example—Customer Provides Internet Service	44
Configuration for Router A	44
Configuration for Router B	45
Configuration for Router C	45
Configuration for Router D	46
Configuration for Router E	47

	Configuration for Router F	48
	Configuration for Router G	49
	Configuration for Router H	49
	Configuration for Router I	50
	Configuration for Router J	51
	Configuration for Router K	52
	Configuration for Router L	53
	Carrier-of-Carriers VPN Example—Customer Provides VPN Service	53
	Configuration for Router A	54
	Configuration for Router B	54
	Configuration for Router C	56
	Configuration for Router D	56
	Configuration for Router E	57
	Configuration for Router F	59
	Configuration for Router G	59
	Configuration for Router H	59
	Configuration for Router I	61
	Configuration for Router J	62
	Configuration for Router K	63
	Configuration for Router L	64
	Multiple Instances for LDP and Carrier-of-Carriers VPNs	64
Chapter 4	Interprovider and Carrier-of-Carriers VPNs Configuration Statements . . .	65
	labeled-unicast (Protocols BGP VPN)	66
	per-group-label	67
	traffic-statistics (Protocols BGP)	67
Part 3	Administration	
Chapter 5	Interprovider and Carrier-of-Carriers VPNs Reference	71
	Example Terminology	71
	Supported Carrier-of-Carriers and Interprovider VPN Standards	72
Part 4	Index	
	Index	77

List of Figures

Part 1	Overview	
Chapter 1	Introduction to Interprovider and Carrier-of-Carriers VPNs	3
	Figure 1: Interprovider VPN Network Topology	5
	Figure 2: Carrier-of-Carriers VPN Architecture	8
Part 2	Configuration	
Chapter 3	Configuration Examples for Interprovider and Carrier-of-Carriers VPNs	29
	Figure 3: Network Topology for the Interprovider VPN Example	30
	Figure 4: Network Topology of Interprovider VPN Example—Multihop MP-EBGP	36
	Figure 5: Carrier-of-Carriers VPN Example Network Topology	44

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xiii
	Table 2: Text and Syntax Conventions	xiv
Part 1	Overview	
Chapter 1	Introduction to Interprovider and Carrier-of-Carriers VPNs	3
	Table 3: Comparison of Interprovider and Carrier-of-Carriers VPNs	9

About the Documentation

- Documentation and Release Notes on page xi
- Supported Platforms on page xi
- Using the Examples in This Manual on page xi
- Documentation Conventions on page xiii
- Documentation Feedback on page xv
- Requesting Technical Support on page xv

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:

- PTX Series
- MX Series
- T Series
- M 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.xml;
    }
  }
}
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.xml; }
```

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 xiii defines notice icons used in this guide.

Table 1: Notice Icons

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

Table 2 on page xiv 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
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric 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]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	}

GUI Conventions

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

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, 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

Overview

- [Introduction to Interprovider and Carrier-of-Carriers VPNs on page 3](#)

CHAPTER 1

Introduction to Interprovider and Carrier-of-Carriers VPNs

- [Traditional VPNs, Interprovider VPNs, and Carrier-of-Carriers VPNs on page 3](#)
- [Standard VPNs on page 4](#)
- [Interprovider and Carrier-of-Carriers VPNs on page 4](#)
- [Interprovider VPNs on page 5](#)
- [Carrier-of-Carriers VPNs on page 8](#)

Traditional VPNs, Interprovider VPNs, and Carrier-of-Carriers VPNs

As VPNs are deployed on the Internet, the customer of a VPN service provider might be another service provider rather than an end customer. The customer service provider depends on the VPN service provider to deliver a VPN transport service between the customer service provider's points of presence (POPs) or regional networks.

If the customer service provider's sites have different autonomous system (AS) numbers, then the VPN transit service provider supports carrier-of-carrier VPN service for the interprovider VPN service. If the customer service provider's sites have the same AS number, then the VPN transit service provider delivers a carrier-of-carriers VPN service.

The sections that follow provide an overview of traditional VPNs, interprovider and carrier-of-carriers VPNs, and the differences in how external and internal routes are handled in each of these environments.

In traditional IP routing architectures, there is a clear distinction between internal routes and external routes. From the perspective of an Internet service provider (ISP), internal routes include all the provider's internal links (including BGP next hops) and loopback interfaces. These internal routes are exchanged with other routing platforms in the ISP's network by means of an interior gateway protocol (IGP), such as OSPF or IS-IS. All routes learned at Internet peering points or from customer sites are classified as external routes and are distributed by means of an exterior gateway protocol (EGP) such as BGP. In traditional IP routing architectures, the number of internal routes is typically much smaller than the number of external routes.

Standard VPNs

The traditional distinction between internal routes and external routes also applies to VPN routing architectures. As shown in *Routers in a VPN*, the provider (P) routers maintain only the service provider's internal routes (to provider edge [PE] routers and other P routers); they do not maintain VPN routes. PE routers are the only devices in the provider network that are required to maintain external routes.

The BGP next hop connects the external routes to the internal routes in traditional VPNs:

- The BGP next hop is advertised with each external route in BGP advertisements.
- The route to the BGP next hop is an internal route that is advertised by the IGP.
- MPLS provides packet forwarding from the ingress PE router to the BGP next-hop egress PE router.

Interprovider and Carrier-of-Carriers VPNs

All interprovider and carrier-of-carriers VPNs share the following characteristics:

- Each interprovider or carrier-of-carriers VPN customer must distinguish between internal and external customer routes.
- Internal customer routes must be maintained by the VPN service provider in its PE routers.
- External customer routes are carried only by the customer's routing platforms, not by the VPN service provider's routing platforms.

The key difference between interprovider and carrier-of-carriers VPNs is whether the customer sites belong to the same AS or to separate ASs:

- [“Interprovider VPNs” on page 5](#)—The customer sites belong to different ASs. You need to configure EBGp to exchange the customer's external routes.
- [“Carrier-of-Carriers VPNs” on page 8](#)—The customer sites belong to the same AS. You need to configure IBGP to exchange the customer's external routes.

In general, each service provider in a VPN hierarchy is required to maintain its own internal routes in its P routers, and the internal routes of its customers in its PE routers. By recursively applying this rule, it is possible to create a hierarchy of VPNs.


The following are definitions of the types of PE routers specific to interprovider and carrier-of-carriers VPNs:

- The AS border router is located at the AS border and handles traffic leaving and entering the AS.
- The end PE router is the PE router in the customer VPN; it is connected to the CE router at the end customer's site.

Interprovider VPNs

Interprovider VPNs provide connectivity between separate ASs. This functionality might be used by a VPN customer who has connections to several different service providers, or different connections to the same service provider in different geographic regions, each of which has a different AS. [Figure 1 on page 5](#) illustrates the type of network topology used by an interprovider VPN.

Figure 1: Interprovider VPN Network Topology



The following sections describe the ways you can configure an interprovider VPN:

- [Linking VRF Tables Between Autonomous Systems on page 5](#)
- [Configuring Next Generation Layer 3 VPNs Options A, B, and C on page 5](#)
- [Configuring Multihop MP-EBGP Between AS Border Routers on page 6](#)

Linking VRF Tables Between Autonomous Systems

You can connect two separate ASs by simply linking the VPN routing and forwarding (VRF) table in the AS border router (ASBR) of one AS to the VRF table in the ASBR in the other AS. Each ASBR must include a VRF routing instance for each VPN configured in both service provider networks. You then configure an IP session between the two ASBRs. In effect, the ASBRs treat each other as customer edge (CE) routers.

Because of the complexity of the configuration, particularly with regard to scaling, this method is not recommended. The details of this configuration are not provided with documentation.

Configuring Next Generation Layer 3 VPNs Options A, B, and C

For next generation Layer 3 VPNs, the PE routers within an AS use multiprotocol external BGP (MP-EBGP) to distribute labeled VPN–Internet Protocol version 4 (IPv4) routes to an ASBR or to a route reflector of which the ASBR is a client. The ASBR uses multiprotocol external BGP (MP-EBGP) to distribute the labeled VPN-IPv4 routes to its peer ASBR in the neighboring AS. The peer ASBR then uses MP-IBGP to distribute labeled VPN-IPv4 routes to PE routers, or to a route reflector of which the PE routers are a client.

You can configure both unicast (Junos OS Release 9.5 and later) and multicast (Junos OS Release 12.1 and later) next generation Layer 3 VPNs across ASs. The Junos OS software supports next generation Layer 3 VPNs option A, option B, and option C:

- **Option A**—This is simple though less scalable interprovider VPN solution to the problem of providing VPN services to a customer that has different sites, not all of which can use the same service provider. In this implementation, the VPN routing and forwarding (VRF) table in the ASBR of one AS is linked to the VRF table in the ASBR in the other AS. Each ASBR must include a VRF instance for each VPN configured in both service provider networks. Then an IGP or BGP must be configured between the ASBRs.

Option B—For this interprovider VPN solution, the customer requires VPN services for different sites, yet the same service provider is not available for all of those sites. With option B, the ASBR routers keep all VPN-IPv4 routes in the routing information base (RIB), and the labels associated with the prefixes are kept in the forwarding information base (FIB). Because the RIB and FIB tables can take too much of the respective allocated memory, this solution is not very scalable for an interprovider VPN. If a transit service provider is used between service provider 1 and service provider 2, the transit service provider also has to keep all VPN-IPv4 routes in the RIB and the corresponding labels in the FIB. The ASBRs at the transit service provider have the same functionality as ASBRs at service provider 1 or service provider 2 in this solution. The PE routers within each AS use multiprotocol internal BGP (MP-IBGP) to distribute labeled VPN-IPv4 routes to an ASBR or to a route reflector of which the ASBR is a client. The ASBR uses MP-EBGP to distribute the labeled VPN-IPv4 routes to its peer ASBR router in the neighboring AS. The peer ASBR then uses MP-IBGP to distribute labeled VPN-IPv4 routes to PE routers, or to a route reflector of which the PE routers are a client.

Option C—For this interprovider VPN solution, the customer service provider depends on the VPN service provider to deliver a VPN transport service between the customer service provider's points of presence (POPs) or regional networks. This functionality might be used by a VPN customer who has connections to several different service providers, or different connections to the same service provider in different geographic regions, each of which has a different AS number. For option C, only routes internal to the service provider networks are announced between ASBRs. This is achieved by using the **family inet labeled-unicast** statements in the IBGP and EBGp configuration on the PE routers. Labeled IPv4 (not VPN-IPv4) routes are exchanged by the ASBRs to support MPLS. An MP-EBGP session between the end PE routers is used for the announcement of VPN-IPv4 routes. In this manner, VPN connectivity is provided while keeping VPN-IPv4 routes out of the core network.

**Related
Documentation**

- *Example: Configuring Interprovider Layer 3 VPN Option A*
- *Example: Configuring Interprovider Layer 3 VPN Option B*
- *Example: Configuring Interprovider Layer 3 VPN Option C*

Configuring Multihop MP-EBGP Between AS Border Routers

In this type of interprovider VPN configuration, P routers do not need to store all the routes in all the VPNs. Only the PE routers must have all the VPN routes. The P routers simply forward traffic to the PE routers—they do not store or process any information about the

packets' destination. The connections between the AS border routers in separate ASs forward traffic between the ASs, much as a label-switched path (LSP) works.

The following are the basic steps you take to configure an interprovider VPN in this manner:

1. Configure multihop EBGp redistribution of labeled VPN-IPv4 routes between the source and destination ASs.
2. Configure EBGp to redistribute labeled IPv4 routes from its AS to neighboring ASs.
3. Configure MPLS on the end PE routers of the VPNs.


Carrier-of-Carriers VPNs

The customer of a VPN service provider might be a service provider for the end customer. The following are the two main types of carrier-of-carriers VPNs (as described in RFC 4364):

- [“Internet Service Provider as the Customer” on page 9](#)—The VPN customer is an ISP that uses the VPN service provider’s network to connect its geographically disparate regional networks. The customer does not have to configure MPLS within its regional networks.
- [“VPN Service Provider as the Customer” on page 9](#)—The VPN customer is itself a VPN service provider offering VPN service to its customers. The carrier-of-carriers VPN service customer relies on the backbone VPN service provider for inter-site connectivity. The customer VPN service provider is required to run MPLS within its regional networks.

Figure 2 on page 8 illustrates the network architecture used for a carrier-of-carriers VPN service.

Figure 2: Carrier-of-Carriers VPN Architecture



807197

This topic covers the following:

- [Internet Service Provider as the Customer on page 9](#)
- [VPN Service Provider as the Customer on page 9](#)

Internet Service Provider as the Customer

In this type of carrier-of-carriers VPN configuration, ISP A configures its network to provide Internet service to ISP B. ISP B provides the connection to the customer wanting Internet service, but the actual Internet service is provided by ISP A.

This type of carrier-of-carriers VPN configuration has the following characteristics:

- The carrier-of-carriers VPN service customer (ISP B) does not need to configure MPLS on its network.
- The carrier-of-carriers VPN service provider (ISP A) must configure MPLS on its network.
- MPLS must also be configured on the CE routers and PE routers connected together in the carrier-of-carriers VPN service customer's and carrier-of-carriers VPN service provider's networks.

VPN Service Provider as the Customer

A VPN service provider can have customers that are themselves VPN service providers. In this type of configuration, also called a hierarchical or recursive VPN, the customer VPN service provider's VPN-IPv4 routes are considered external routes, and the backbone VPN service provider does not import them into its VRF table. The backbone VPN service provider imports only the customer VPN service provider's internal routes into its VRF table.

The similarities and differences between interprovider and carrier-of-carriers VPNs are shown in [Table 3 on page 9](#).

Table 3: Comparison of Interprovider and Carrier-of-Carriers VPNs

Feature	ISP Customer	VPN Service Provider Customer
Customer edge device	AS border router	PE router
IBGP sessions	Carry IPv4 routes	Carry external VPN-IPv4 routes with associated labels
Forwarding within the customer network	MPLS is optional	MPLS is required

PART 2

Configuration

- [Configuring Interprovider and Carrier-of-Carriers VPNs on page 13](#)
- [Configuration Examples for Interprovider and Carrier-of-Carriers VPNs on page 29](#)
- [Interprovider and Carrier-of-Carriers VPNs Configuration Statements on page 65](#)

CHAPTER 2

Configuring Interprovider and Carrier-of-Carriers VPNs

- [Understanding Interprovider VPNs on page 13](#)
- [Configuring Carrier-of-Carriers VPNs for Customers That Provide Internet Service on page 14](#)
- [Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service on page 19](#)
- [Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics on page 26](#)

Understanding Interprovider VPNs

Interprovider and carrier-of-carriers VPNs provide solutions for situations in which the customer of a VPN service provider might be another service provider rather than an end customer. The customer service provider depends on the virtual private network (VPN) service provider (SP) to deliver a VPN transport service between the customer service provider's points of presence (POPs) or regional networks.

For interprovider VPNs, the customer sites belong to different ASs. You need to configure external BGP (EBGP) to exchange the customer's external routes. There are several different methods for enabling interprovider VPNs based on RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*:

- Interprovider Layer 3 VPN Option A—Interprovider VRF-to-VRF connections at the AS boundary routers (ASBR) (not very scalable).
- Interprovider Layer 3 VPN Option B—Interprovider EBGP redistribution of labeled VPN-IPv4 routes from AS to neighboring AS (somewhat scalable).
- Interprovider Layer 3 VPN Option C—Interprovider multihop EBGP redistribution of labeled VPN-IPv4 routes between source and destination ASs, with EBGP redistribution of labeled IPv4 routes from AS to neighboring AS (very scalable).

Related Documentation

- [Example: Configuring Interprovider Layer 3 VPN Option A](#)
- [Example: Configuring Interprovider Layer 3 VPN Option B](#)
- [Example: Configuring Interprovider Layer 3 VPN Option C](#)

Configuring Carrier-of-Carriers VPNs for Customers That Provide Internet Service

You can configure a carrier-of-carriers VPN service for customers who want to provide basic Internet service. The carrier-of-carriers VPN service provider must configure MPLS in its network, although this configuration is optional for the carrier service customer. [Figure 2 on page 8](#) shows how the routers in this type of service interconnect.

To configure a carrier-of-carriers VPN, perform the tasks described in the following sections:

- [Configuring the Carrier-of-Carriers VPN Service Customer's CE Router on page 14](#)
- [Configuring the Carrier-of-Carriers VPN Service Provider's PE Routers on page 16](#)

Configuring the Carrier-of-Carriers VPN Service Customer's CE Router

The carrier-of-carriers VPN service customer's router acts as a CE router with respect to the service provider's PE router. The following sections describe how to configure the carrier-of-carriers VPN service customer's CE router:

- [Configuring MPLS on page 14](#)
- [Configuring BGP on page 14](#)
- [Configuring OSPF on page 15](#)
- [Configuring Policy Options on page 15](#)

Configuring MPLS

To configure MPLS on the customer's CE router, include the **mpls** statement:

```
mpls {  
  traffic-engineering bgp-igp;  
  interface interface-name;  
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

Configuring BGP

To configure a group to collate the customer's internal routes, include the **bgp** statement:

```
bgp {  
  group group-name {  
    type internal;  
    local-address address;  
    neighbor address;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

The customer's CE router must be able to send labels to the VPN service provider's router. Enable this by including the **labeled-unicast** statement in the configuration for the BGP group:

```
bgp {
  group group-name {
    export internal;
    peer-as as-number;
    neighbor address {
      family inet {
        labeled-unicast;
      }
    }
  }
}
```

You can include the **bgp** statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring OSPF

To configure OSPF on the customer's CE router, include the **ospf** statement:

```
ospf {
  area area-id {
    interface interface-name {
      passive;
    }
    interface interface-name;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring Policy Options

To configure policy options on the customer's CE router, include the **policy-statement** statement:

```
policy-statement statement-name {
  term term-name {
    from protocol [ospf direct ldp];
    then accept;
  }
}
```

```
term term-name {  
    then reject;  
}
```

You can include this statement at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

Configuring the Carrier-of-Carriers VPN Service Provider's PE Routers

The service provider's PE routers connect to the customer's CE routers and forward the customer's VPN traffic across the provider's network.

The following sections describe how to configure the carrier-of-carriers VPN service provider's PE routers:

- [Configuring MPLS on page 16](#)
- [Configuring BGP on page 16](#)
- [Configuring IS-IS on page 17](#)
- [Configuring LDP on page 17](#)
- [Configuring a Routing Instance on page 17](#)
- [Configuring Policy Options on page 18](#)

Configuring MPLS

To configure MPLS on the provider's PE routers, include the **mpls** statement:

```
mpls {  
    interface interface-name;  
    interface interface-name;  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring BGP

To configure a BGP session with the provider PE router at the other end of the provider's network, include the **bgp** statement:

```
bgp {  
    group group-name {  
        type internal;  
        local-address address;  
        family inet-vpn {  
            any;  
        }  
        neighbor address;
```



```
}  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring IS-IS

To configure IS-IS on the provider's PE routers, include the **isis** statement:

```
isis {  
  interface interface-name;  
  interface interface-name {  
    passive;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring LDP

To configure LDP on the provider's PE routers, include the **ldp** statement:

```
ldp {  
  interface interface-name;  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring a Routing Instance

To configure Layer 3 VPN service with the customer's CE router, include the **labeled-unicast** statement in the configuration for the routing instance so the PE router can send labels to the customer's CE router:

```
routing-instance-name {  
  instance-type vrf;  
  interface interface-name;  
  route-distinguisher address;  
  vrf-import policy-name;  
  vrf-export policy-name;  
  protocols {  
    bgp {  
      group group-name {  
        peer-as as-number;  
        neighbor address {  
          family inet {
```

```
        labeled-unicast;  
    }  
}  
}  
}
```

You can include these statements at the following hierarchy levels:

- [edit routing-instances]
- [edit logical-systems *logical-system-name* routing-instances]

Configuring Policy Options

To configure a policy statement to import routes from the customer's CE router, include the **policy-statement** statement:

```
policy-statement policy-name {  
  term term-name {  
    from {  
      protocol bgp;  
      community community-name;  
    }  
    then accept;  
  }  
  term term-name {  
    then reject;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

To configure a policy statement to export routes to the customer's CE router, include the **policy-statement** and **community** statements:

```
policy-statement policy-name {  
  term term-name {  
    from protocol bgp;  
    then {  
      community add community-name;  
      accept;  
    }  
  }  
  term term-name {  
    then reject;  
  }  
}  
community community-name members value;
```

You can include these statements at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service

You can configure a carrier-of-carriers VPN service for customers who want VPN service. [Figure 2 on page 8](#) shows how the routers in this type of service interconnect.

To configure the following routers in the customer's and provider's networks to enable carrier-of-carriers VPN service, you perform the steps in the following sections:

- [Configuring the Carrier-of-Carriers Customer's PE Router on page 19](#)
- [Configuring the Carrier-of-Carriers Customer's CE Router on page 22](#)
- [Configuring the Provider's PE Router on page 24](#)

Configuring the Carrier-of-Carriers Customer's PE Router

The carrier-of-carriers customer's PE router is connected to the end customer's CE router.

The following sections describe how to configure the carrier-of-carriers customer's PE router:

- [Configuring MPLS on page 19](#)
- [Configuring BGP on page 20](#)
- [Configuring OSPF on page 20](#)
- [Configuring LDP on page 20](#)
- [Configuring VPN Service in the Routing Instance on page 21](#)
- [Configuring Policy Options on page 21](#)

Configuring MPLS

To configure MPLS on the carrier-of-carriers customer's PE router, include the **mpls** statement:

```
mpls {  
  interface interface-name;  
  interface interface-name;  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring BGP

Include the **labeled-unicast** statement in the configuration for the IBGP session to the carrier-of-carriers customer's CE router (see [“Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service” on page 19](#)), and include the **family-inet-vpn** statement in the configuration for the IBGP session to the carrier-of-carriers PE router on the other side of the network:

```
bgp {
  group group-name {
    type internal;
    local-address address;
    neighbor address {
      family inet {
        labeled-unicast;
        resolve-vpn;
      }
    }
  }
  neighbor address {
    family inet-vpn {
      any;
    }
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring OSPF

To configure OSPF on the carrier-of-carriers customer's PE router, include the **ospf** statement:

```
ospf {
  area area-id {
    interface interface-name {
      passive;
    }
    interface interface-name;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring LDP

To configure LDP on the carrier-of-carriers customer's PE router, include the **ldp** statement:

```
ldp {
```

```
    interface interface-name;  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring VPN Service in the Routing Instance

To configure VPN service for the end customer's CE router on the carrier-of-carriers customer's PE router, include the following statements:

```
instance-type vrf;  
interface interface-name;  
route-distinguisher address;  
vrf-import policy-name;  
vrf-export policy-name;  
protocols {  
    bgp {  
        group group-name {  
            peer-as as-number;  
            neighbor address;  
        }  
    }  
}
```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Configuring Policy Options

To configure policy options to import and export routes to and from the end customer's CE router, include the **policy-statement** and **community** statements:

```
policy-statement policy-name {  
    term term-name {  
        from {  
            protocol bgp;  
            community community-name;  
        }  
        then accept;  
    }  
    term term-name {  
        then reject;  
    }  
}  
policy-statement policy-name {  
    term term-name {  
        from protocol bgp;  
        then {  
            community add community-name;  
        }  
    }  
}
```

```
        accept;
      }
    }
    term term-name {
      then reject;
    }
  }
  community community-name members value;
```

You can include these statements at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

Configuring the Carrier-of-Carriers Customer's CE Router

The carrier-of-carriers customer's CE router connects to the provider's PE router. Complete the instructions in the following sections to configure the carrier-of-carriers customers' CE router:

- [Configuring MPLS on page 22](#)
- [Configuring BGP on page 22](#)
- [Configuring OSPF and LDP on page 23](#)
- [Configuring Policy Options on page 24](#)

Configuring MPLS

In the MPLS configuration for the carrier-of-carriers customer's CE router, include the interfaces to the provider's PE router and to a P router in the customer's network:

```
mpls {
  traffic-engineering bgp-igp;
  interface interface-name;
  interface interface-name;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring BGP

In the BGP configuration for the carrier-of-carriers customer's CE router, configure a group that includes the **labeled-unicast** statement to extend VPN service to the PE router connected to the end customer's CE router:

```
bgp {
  group group-name {
    type internal;
    local-address address;
    neighbor address {
      family inet {
```

```

        labeled-unicast;
    }
}
}

```

You can include the **bgp** statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

To configure a group to send labeled internal routes to the provider's PE router, include the **bgp** statement:

```

bgp {
  group group-name {
    export internal;
    peer-as as-number;
    neighbor address {
      family inet {
        labeled-unicast;
      }
    }
  }
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring OSPF and LDP

To configure OSPF and LDP on the carrier-of-carriers customer's CE router, include the **ospf** and **ldp** statements:

```

ospf {
  area area-id {
    interface interface-name {
      passive;
    }
    interface interface-name;
  }
}
ldp {
  interface interface-name;
}

```

You can include these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring Policy Options

To configure the policy options on the carrier-of-carriers customer's CE router, include the **policy-statement** statement:

```
policy-statement policy-statement-name {  
  term term-name {  
    from protocol [ ospf direct ldp ];  
    then accept;  
  }  
  term term-name {  
    then reject;  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

Configuring the Provider's PE Router

The carrier-of-carriers provider's PE routers connect to the carrier customer's CE routers. Complete the instructions in the following sections to configure the provider's PE router:

- [Configuring MPLS on page 24](#)
- [Configuring a PE-Router-to-PE-Router BGP Session on page 24](#)
- [Configuring IS-IS and LDP on page 25](#)
- [Configuring Policy Options on page 25](#)
- [Configuring a Routing Instance to Send Routes to the CE Router on page 26](#)

Configuring MPLS

In the MPLS configuration, specify at least two interfaces—one to the customer's CE router and one to connect to the provider's PE router on the other side of the provider's network:

```
interface interface-name;  
interface interface-name;
```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

Configuring a PE-Router-to-PE-Router BGP Session

To configure a PE-router-to-PE-router BGP session on the provider's PE routers to allow VPN-IPv4 routes to pass between the PE routers, include the **bgp** statement:

```
bgp {  
  group group-name {
```



```
    type internal;
    local-address address;
    family inet-vpn {
        any;
    }
    neighbor address;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring IS-IS and LDP

To configure IS-IS and LDP on the provider's PE routers, include the **isis** and **ldp** statements:

```
isis {
    interface interface-name;
    interface interface-name {
        passive;
    }
}
ldp {
    interface interface-name;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Configuring Policy Options

To configure policy statements on the provider's PE router to export routes to and import routes from the carrier customer's network, include the **policy-statement** and **community** statements:

```
policy-statement statement-name {
    term term-name {
        from {
            protocol bgp;
            community community-name;
        }
        then accept;
    }
    term term-name {
        then reject;
    }
}
policy-statement statement-name {
    term term-name {
```

```
    from protocol bgp;
    then {
        community add community-name;
        accept;
    }
}
term term-name {
    then reject;
}
}
community community-name members value;
```

You can include these statements at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

Configuring a Routing Instance to Send Routes to the CE Router

To configure the routing instance on the provider's PE router to send labeled routes to the carrier customer's CE router, include the following statements:

```
instance-type vrf;
interface interface-name;
route-distinguisher value;
vrf-import policy-name;
vrf-export policy-name;
protocols {
    bgp {
        group group-name {
            peer-as as-number;
            neighbor address {
                family inet {
                    labeled-unicast;
                }
            }
        }
    }
}
```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics

You can configure BGP to gather traffic statistics for interprovider and carrier-of-carriers VPNs.

To configure BGP to gather traffic statistics for interprovider and carrier-of-carriers VPNs, include the **traffic-statistics** statement:

```
traffic-statistics {
```

```

file filename <world-readable | no-world-readable>;
interval seconds;
}

```

For a list of the hierarchy levels at which you can include this statement, see the summary section for this statement.



NOTE: Traffic statistics for interprovider and carrier-of-carriers VPNs are available only for IPv4. IPv6 is not supported.

If you do not specify a filename, the statistics are not written to a file. However, if you have included the **traffic-statistics** statement in the BGP configuration, the statistics are still available and can be accessed by means of the **show bgp group traffic-statistics group-name** command.

To account for traffic from each customer separately, separate labels must be advertised for the same prefix to the peer routers in different groups. To enable separate traffic accounting, you need to include the **per-group-label** statement in the configuration for each BGP group. By including this statement, statistics are collected and displayed that account for traffic sent by the peers of the specified BGP group.

If you configure the statement at the **[edit protocols bgp family inet]** hierarchy level, rather than configuring it for a specific BGP group, then the traffic statistics are shared with all BGP groups configured with the **traffic-statistics** statement but not configured with the **per-group-label** statement.

To account for traffic from each customer separately, include the **per-group-label** statement in the configuration for each BGP group:

```
per-group-label;
```

For a list of the hierarchy levels at which you can include this statement, see the summary section for this statement.

The following shows a sample of the output to the traffic statistics file:

```

Dec 19 10:39:54 Statistics for BGP group ext2 (Index 1) NLRI inet-labeled-unicast
Dec 19 10:39:54  FEC                Packets      Bytes      EgressAS   FECLabel
Dec 19 10:39:54  10.255.245.55          0           0           I       100160
Dec 19 10:39:54  10.255.245.57          0           0           I       100112
Dec 19 10:39:54  100.101.0.0            0           0          25       100080
Dec 19 10:39:54  100.102.0.0            0           0          25       100080
Dec 19 10:39:54  100.103.0.0          109        9592          25       100048
Dec 19 10:39:54  100.104.0.0          109        9592          25       100048
Dec 19 10:39:54  192.168.25.0           0           0           I       100064
Dec 19 10:39:54  Dec 19 10:39:54, read statistics for 5 FECs in 00:00:00 seconds
(10 queries) for BGP group ext2 (Index 1) NLRI inet-labeled-unicast

```


CHAPTER 3

Configuration Examples for Interprovider and Carrier-of-Carriers VPNs

- [Interprovider VPN Example—MP-EBGP Between ISP Peer Routers on page 29](#)
- [Interprovider VPN Example—Multihop MP-EBGP with P Routers on page 36](#)
- [Carrier-of-Carriers VPN Examples on page 43](#)
- [Carrier-of-Carriers VPN Example—Customer Provides Internet Service on page 44](#)
- [Carrier-of-Carriers VPN Example—Customer Provides VPN Service on page 53](#)
- [Multiple Instances for LDP and Carrier-of-Carriers VPNs on page 64](#)

Interprovider VPN Example—MP-EBGP Between ISP Peer Routers


In this example, all routes learned from the CE routers are sent over both service provider networks as VPN-IPv4 routes. The routes are initially learned by the PE routers (Router B and Router E) from the CE routers (Router A and Router F) and are announced by the PE routers to the AS border routers (Router C and Router D). The AS border routers are then configured with an MP-EBGP session, enabling them to pass the VPN-IPv4 routes with each other. When an AS border router—Router C for example—learns VPN-IPv4 routes from an IBGP PE, the following events occur:

1. Router C sets itself as the next hop for the route and creates a label for that route.
2. Router C advertises the VPN-IPv4 route to PE Router D in AS 10045.
3. Router D sets the next hop to itself, creates another label, and then forwards the label and the route to its IBGP PE router (Router E).

This example has scaling limitations because of restrictions on the number of labels each PE router needs to allocate at the AS border.

[Figure 3 on page 30](#) illustrates the network topology used in this VPN example.

Figure 3: Network Topology for the Interprovider VPN Example



For configuration information see the following sections:

- [Configuration for Router A on page 30](#)
- [Configuration for Router B on page 30](#)
- [Configuration for Router C on page 32](#)
- [Configuration for Router D on page 33](#)
- [Configuration for Router E on page 34](#)
- [Configuration for Router F on page 36](#)

Configuration for Router A

Configure a family **inet** EBGp session with Router B and export the direct routes:

```
[edit]
protocols {
  bgp {
    group to-provider {
      export attached;
      peer-as 10023;
      neighbor 192.168.198.2;
    }
  }
}
policy-options {
  policy-statement attached {
    from protocol direct;
    then accept;
  }
}
```

Configuration for Router B

Router A is configured as a CE router (using the **routing-instances** statement) in the configuration for Router B. Because they exchange VPN-IPv4 routes, Router D and Router C are configured as PE routers.

Configure Router B:

```
[edit]
protocols {
  rsvp {
    interface t3-0/0/0.0;
  }
  mpls {
    label-switched-path to-routerC {
      to 10.255.14.171;
      description "to-routerC for use with VPNs";
    }
    interface t3-0/0/0.0;
    interface so-1/2/0.0;
  }
  bgp {
    group to-ibgp {
      type internal;
      local-address 10.255.14.175;
      family inet-vpn {
        unicast;
      }
      neighbor 10.255.14.171;
    }
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface t3-0/0/0.0;
      interface lo0.0 {
        passive;
      }
    }
  }
}
routing-instances {
  vpna {
    instance-type vrf;
    interface so-1/2/0.0;
    route-distinguisher 10.255.14.175:9;
    vrf-import vpna-import;
    vrf-export vpna-export;
    protocols {
      bgp {
        group to-ce {
          peer-as 9;
          neighbor 192.168.198.1;
        }
      }
    }
  }
}
policy-options {
  policy-statement vpna-import {
    term 1 {
```

```

        from {
            protocol bgp;
            community vpna-comm;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
policy-statement vpna-export {
    term 1 {
        from protocol bgp;
        then {
            community add vpna-comm;
            accept;
        }
    }
    term 2 {
        then reject;
    }
}
community vpna-comm members target:100:1001;
}

```

Configuration for Router C

In the BGP protocol configuration for Router C, include the **keep all** statement. When this statement is included, BGP must store every route learned through BGP.

When you configure **keep all** or **keep none** and the peers support route refresh, the local speaker sends a refresh message and performs an import evaluation. For these peers, the sessions do not restart. To determine if a peer supports refresh, check for **Peer supports Refresh capability** in the output of the **show bgp neighbor** command.



CAUTION: If you configure **keep all** or **keep none** and the peer does not support session restart, the associated BGP sessions are restarted (flapped).

Configure two BGP sessions (configure **family inet-vpn** on both sessions):

- IBGP session to Router B (group **to-ibgp** in this example)
- EBGP session to Router D (group **to-ebgp-pe** in this example)

Interface **t3-0/2/0** is added at the **[edit protocols mpls]** hierarchy level, allowing BGP to announce routes with labels over the EBGP session.

Configure Router C:

```

[edit]
protocols {
    rsvp {
        interface t3-0/2/0.0;
    }
}

```



```

}
mpls {
  label-switched-path to-routerB {
    to 10.255.14.175;
    description "to-routerB for use with vpns";
  }
  interface t3-0/2/0.0;
  interface so-0/0/0.0;
}
bgp {
  keep all;
  group to-ibgp {
    type internal;
    local-address 10.255.14.171;
    family inet-vpn {
      unicast;
    }
    neighbor 10.255.14.175;
  }
  group to-ebgp-pe {
    type external;
    family inet-vpn {
      unicast;
    }
    neighbor 192.168.197.22 {
      peer-as 10045;
    }
  }
}
}
ospf {
  traffic-engineering;
  reference-bandwidth 4g;
  area 0.0.0.0 {
    interface t3-0/2/0.0;
    interface lo0.0 {
      passive;
    }
  }
}
}
}

```

Configuration for Router D

The configuration for Router D is almost identical to that of Router C:

```

[edit]
protocols {
  rsvp {
    interface fe-1/1/0.0;
  }
  mpls {
    label-switched-path to-E {
      to 10.255.14.177;
      description "to-routerE for vpna";
    }
    interface fe-1/1/0.0;
  }
}

```

```
    interface so-0/1/0.0;
  }
  bgp {
    keep all;
    group to-ibgp-pe {
      type internal;
      family inet-vpn {
        unicast;
      }
      neighbor 10.255.14.177;
    }
    group to-ebgp-pe {
      type external;
      family inet-vpn {
        unicast;
      }
      peer-as 10023;
      neighbor 192.168.197.21;
    }
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface fe-1/1/0.0;
      interface lo0.0 {
        passive;
      }
    }
  }
}
```

Configuration for Router E

The configuration for Router E is very similar to the configuration for Router B:

```
[edit]
protocols {
  rsvp {
    interface fe-1/1/2.0;
  }
  mpls {
    label-switched-path to-routerD {
      to 10.255.14.173;
      description "to-routerD for use with VPNa";
    }
    interface fe-1/1/2.0;
    interface so-1/2/0.0;
  }
  bgp {
    group to-ibgp-pe {
      type internal;
      local-address 10.255.14.177;
      family inet-vpn {
        unicast;
      }
    }
  }
}
```

```
        neighbor 10.255.14.173;
    }
}
ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
        interface fe-1/1/2.0;
        interface lo0.0 {
            passive;
        }
    }
}
}
routing-instances {
    vpn {
        instance-type vrf;
        interface so-1/2/0.0;
        route-distinguisher 10.255.14.177:11;
        vrf-import vpn-import;
        vrf-export vpn-export;
        protocols {
            bgp {
                group to-routerF-ce {
                    neighbor 192.168.198.14 {
                        peer-as 11;
                    }
                }
            }
        }
    }
}
}
policy-options {
    policy-statement vpn-import {
        term 1 {
            from {
                protocol bgp;
                community vpn-comm;
            }
            then accept;
        }
        term 2 {
            then reject;
        }
    }
    policy-statement vpn-export {
        term 1 {
            from protocol bgp;
            then {
                community add vpn-comm;
                accept;
            }
        }
        term 2 {
            then reject;
        }
    }
}
```

```

    }
    community vpna-comm members target:100:1001;
  }

```

Configuration for Router F

Configure Router F as a CE router; the configuration is similar to that for Router A:

```

[edit]
protocols {
  bgp {
    group to-provider {
      type external;
      export attached;
      neighbor 192.168.198.13 {
        peer-as 10045;
      }
    }
  }
}
policy-options {
  policy-statement attached {
    from protocol direct;
    then accept;
  }
}
}

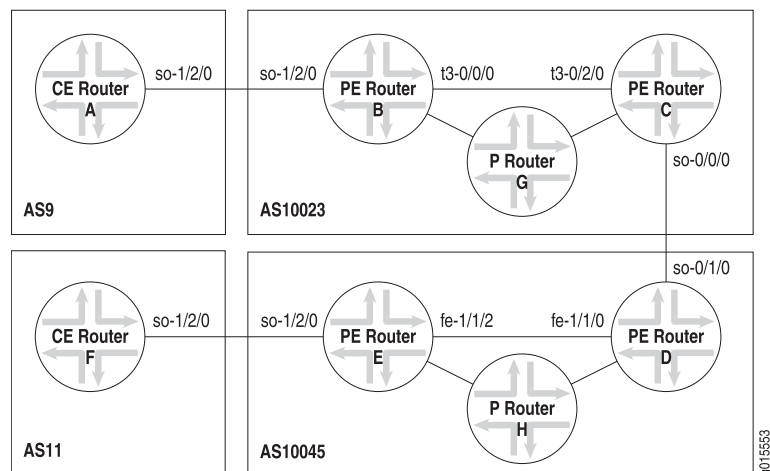
```

Interprovider VPN Example—Multihop MP-EBGP with P Routers

In this example, labeled IPv4 (not VPN-IPv4), routes are exchanged by the AS border routers (Router C and Router D) to provide MPLS connectivity between the PE routers. Router G and H are provider routers.

Figure 4 on page 36 illustrates the network topology used in this VPN example.

Figure 4: Network Topology of Interprovider VPN Example—Multihop MP-EBGP



Only routes internal to the service provider networks should be announced between Router C and Router D. Configure this by including the **family inet labeled-unicast** statement in the IBGP and EBGp configuration on the PE routers. When you set **family inet labeled-unicast**, the local router announces internal routes from inet.0 in the following manner:

- If a label exists for the route, the local router creates a label, performs a swap, and announces the route from inet.0 with the label.
- If a label does not exist for the route, the local router creates a label, performs a pop, and announces the route from inet.0 with the label.

Routes learned from the **labeled-unicast** session are placed into the inet.0 routing table.

In addition, you configure a multihop MP-EBGP session between the end PE routers (Router B and Router E). This additional MP-EBGP session allows the announcement of VPN-IPv4 routes, and allows you to maintain VPN connectivity while keeping VPN-IPv4 routes out of the core of the network.

For configuration information, see the following sections:

- [Configuration for Router A on page 37](#)
- [Configuration for Router B on page 37](#)
- [Configuration for Router C on page 39](#)
- [Configuration for Router D on page 40](#)
- [Configuration for Router E on page 41](#)
- [Configuration for Router F on page 43](#)

Configuration for Router A

The configuration for Router A in this example is identical to the configuration for Router A in [“Interprovider VPN Example—MP-EBGP Between ISP Peer Routers” on page 29](#). See [“Configuration for Router A” on page 30](#)

Configuration for Router B

Router A is configured as a CE router (using the **routing-instances** statement) in the configuration for Router B. Because they exchange VPN-IPv4 routes, Router C and Router D are configured as PE routers.

In the BGP group **to-ibgp**, include the **family inet labeled-unicast** statement to pass labeled IPv4 routes, and configure an EBGp multihop session to pass VPN-IPv4 routes:

```
[edit]
protocols {
  bgp {
    group to-ibgp {
      type internal;
      local-address 10.255.14.175;
      family inet {
        labeled-unicast {
          resolve-vpn;
```

```
    }
  }
  neighbor 10.255.14.171;
}
group to-remote-pe {
  multihop {
    ttl 10;
  }
  family inet-vpn {
    unicast;
  }
  neighbor 10.255.14.177 {
    peer-as 10045;
  }
}
mpls {
  label-switched-path to-routerC {
    to 10.255.14.171;
    description "to-routerC for use with VPNs";
  }
  interface t3-0/0/0.0;
  interface so-1/2/0.0;
}
ospf {
  traffic-engineering;
  reference-bandwidth 4g;
  area 0.0.0.0 {
    interface t3-0/0/0.0;
    interface lo0.0 {
      passive;
    }
  }
}
}
rsvp {
  interface t3-0/0/0.0;
}
}
routing-instances {
  vpna {
    instance-type vrf;
    interface so-1/2/0.0;
    route-distinguisher 10.255.14.175:9;
    vrf-import vpna-import;
    vrf-export vpna-export;
    protocols {
      bgp {
        group to-ce {
          peer-as 9;
          neighbor 192.168.198.1;
        }
      }
    }
  }
}
}
policy-options {
  policy-statement vpna-import {
```

```

term 1 {
  from {
    protocol bgp;
    community vpna-comm;
  }
  then accept;
}
term 2 {
  then reject;
}
}
policy-statement vpna-export {
  term 1 {
    from protocol bgp;
    then {
      community add vpna-comm;
      accept;
    }
  }
  term 2 {
    then reject;
  }
}
community vpna-comm members target:100:1001;
}
}

```

Configuration for Router C

Configure two BGP sessions (configure **family inet-vpn** on both sessions):

- IBGP session to Router B (group **to-ibgp** in this example)
- EBGP session to Router D (group **to-ebgp-pe** in this example)

Interface **t3-0/2/0** is added at the **[edit protocols mpls]** hierarchy level, allowing BGP to announce routes with labels over the EBGP session.

Configure Router C:

```

[edit]
protocols {
  bgp {
    group to-ibgp {
      type internal;
      local-address 10.255.14.171;
      family inet {
        labeled-unicast;
      }
      neighbor 10.255.14.175;
    }
    group to-ebgp-pe {
      type external;
      family inet {
        labeled-unicast;
      }
    }
  }
}

```

```
export internal;
neighbor 192.168.197.22 {
    peer-as 10045;
}
}
mpls {
    label-switched-path to-routerB {
        to 10.255.14.175;
        description "to-routerB for use with vpns";
    }
    interface t3-0/2/0.0;
    interface so-0/0/0.0;
    traffic-engineering bgp-igp;
}
rsvp {
    interface t3-0/2/0.0;
}
ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
        interface t3-0/2/0.0;
        interface lo0.0 {
            passive;
        }
    }
}
}
}
policy-options {
    policy-statement internal {
        term 1 {
            from protocol [ospf direct ldp];
            then accept;
        }
        term 2 {
            then reject;
        }
    }
}
}
```

Configuration for Router D

Configure Router D:

```
[edit]
protocols {
    bgp {
        group to-ibgp-pe {
            type internal;
            family inet {
                labeled-unicast;
            }
            neighbor 10.255.14.177;
        }
        group to-ebgp-pe {
```



```

    type external;
    family inet {
        labeled-unicast;
    }
    export internal;
    peer-as 10023;
    neighbor 192.168.197.21;
}
mpls {
    label-switched-path to-E {
        to 10.255.14.177;
        description "to-routerE for vpna";
    }
    interface fe-1/1/0.0;
    interface so-0/1/0.0;
    traffic-engineering bgp-igp;
}
ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
        interface fe-1/1/0.0;
        interface lo0.0 {
            passive;
        }
    }
}
}
rsvp {
    interface fe-1/1/0.0;
}
}
policy-options {
    policy-statement internal {
        term 1 {
            from protocol [ospf direct ldp];
            then accept;
        }
        term 2 {
            then reject;
        }
    }
}
}
}

```

Configuration for Router E

The configuration for Router E is very similar to the configuration for Router B:

```

[edit]
protocols {
    bgp {
        group to-ibgp-pe {
            type internal;
            local-address 10.255.14.177;
            family inet {
                labeled-unicast;
            }
        }
    }
}

```

```
    }
    neighbor 10.255.14.173;
  }
  group to-remote-pe {
    multihop {
      ttl 10;
    }
    family inet-vpn {
      unicast;
    }
    neighbor 10.255.14.175 {
      peer-as 10023;
    }
  }
}
mpls {
  label-switched-path to-routerD {
    to 10.255.14.173;
    description "to-routerD for use with VPNa";
  }
  interface fe-1/1/2.0;
  interface so-1/2/0.0;
}
ospf {
  traffic-engineering;
  reference-bandwidth 4g;
  area 0.0.0.0 {
    interface fe-1/1/2.0;
    interface lo0.0 {
      passive;
    }
  }
}
}
rsvp {
  interface fe-1/1/2.0;
}
}
routing-instances {
  vpn {
    instance-type vrf;
    interface so-1/2/0.0;
    route-distinguisher 10.255.14.177:11;
    vrf-import vpn-import;
    vrf-export vpn-export;
    protocols {
      bgp {
        group to-routerF-ce {
          neighbor 192.168.198.14 {
            peer-as 11;
          }
        }
      }
    }
  }
}
}
policy-options {
  policy-statement vpn-import {
```

```

term 1 {
  from {
    protocol bgp;
    community vpna-comm;
  }
  then accept;
}
term 2 {
  then reject;
}
}
policy-statement vpna-export {
  term 1 {
    from protocol bgp;
    then {
      community add vpna-comm;
      accept;
    }
  }
  term 2 {
    then reject;
  }
}
community vpna-comm members target:100:1001;
}
}

```

Configuration for Router F

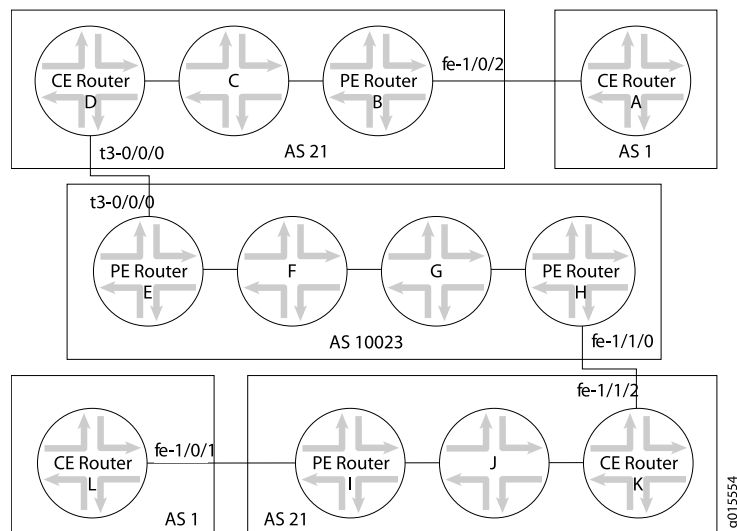
The configuration for Router F in this example is identical to the configuration for Router F in [“Interprovider VPN Example—MP-EBGP Between ISP Peer Routers”](#) on page 29. See [“Configuration for Router F”](#) on page 36.

Carrier-of-Carriers VPN Examples

A carrier-of-carriers service allows an Internet service provider (ISP) to connect to a transparent outsourced backbone at multiple locations.

[Figure 5](#) on page 44 shows the network topology in both carrier-of-carriers examples.

Figure 5: Carrier-of-Carriers VPN Example Network Topology



Carrier-of-Carriers VPN Example—Customer Provides Internet Service

In this example, the carrier customer is not required to configure MPLS and LDP on its network. However, the carrier provider must configure MPLS and LDP on its network.

For configuration information see the following sections:

- [Configuration for Router A on page 44](#)
- [Configuration for Router B on page 45](#)
- [Configuration for Router C on page 45](#)
- [Configuration for Router D on page 46](#)
- [Configuration for Router E on page 47](#)
- [Configuration for Router F on page 48](#)
- [Configuration for Router G on page 49](#)
- [Configuration for Router H on page 49](#)
- [Configuration for Router I on page 50](#)
- [Configuration for Router J on page 51](#)
- [Configuration for Router K on page 52](#)
- [Configuration for Router L on page 53](#)

Configuration for Router A

In this example, Router A represents an end customer. You configure this router as a CE device.

```
[edit]
protocols {
  bgp {
    group to-routerB {
```

```

        export attached;
        peer-as 21;
        neighbor 192.168.197.169;
    }
}
}
policy-options {
    policy-statement attached {
        from protocol direct;
        then accept;
    }
}
}

```

Configuration for Router B

Router B can act as the gateway router, responsible for aggregating end customers and connecting them to the network. If a full-mesh IBGP session is configured, you can use route reflectors.

```

[edit]
protocols {
    bgp {
        group int {
            type internal;
            local-address 10.255.14.179;
            neighbor 10.255.14.175;
            neighbor 10.255.14.181;
            neighbor 10.255.14.176;
            neighbor 10.255.14.178;
            neighbor 10.255.14.177;
        }
        group to-vpn-blue {
            peer-as 1;
            neighbor 192.168.197.170;
        }
    }
    ospf {
        area 0.0.0.0 {
            interface lo0.0 {
                passive;
            }
            interface fe-1/0/3.0;
            interface fe-1/0/2.0 {
                passive;
            }
        }
    }
}
}

```

Configuration for Router C

Configure Router C:

```

[edit]
protocols {
    bgp {

```

```
group int {
    type internal;
    local-address 10.255.14.176;
    neighbor 10.255.14.179;
    neighbor 10.255.14.175;
    neighbor 10.255.14.177;
    neighbor 10.255.14.178;
    neighbor 10.255.14.181;
}
}
ospf {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface fe-0/3/3.0;
        interface fe-0/3/0.0;
    }
}
}
```

Configuration for Router D

Router D is the CE router with respect to AS 10023. In a carrier-of-carriers VPN, the CE router must be able to send labels to the carrier provider; this is done with the **labeled-unicast** statement in group **to-isp-red**.

```
[edit]
protocols {
    mpls {
        interface t3-0/0/0.0;
    }
    bgp {
        group int {
            type internal;
            local-address 10.255.14.175;
            neighbor 10.255.14.179;
            neighbor 10.255.14.176;
            neighbor 10.255.14.177;
            neighbor 10.255.14.178;
            neighbor 10.255.14.181;
        }
        group to-isp-red {
            export internal;
            peer-as 10023;
            neighbor 192.168.197.13 {
                family inet {
                    labeled-unicast;
                }
            }
        }
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
    }
}
```

```

    }
    interface fe-0/3/0.0;
    interface t3-0/0/0.0 {
        passive;
    }
}
}
}
policy options {
    policy-statement internal {
        term a {
            from protocol [ ospf direct ];
            then accept;
        }
        term b {
            then reject;
        }
    }
}
}

```

Configuration for Router E

This configuration sets up the **inet-vpn** IBGP session with Router H and the PE router portion of the VPN with Router D. Because Router D is required to send labels in this example, configure the BGP session with the **labeled-unicast** statement within the virtual routing and forwarding (VRF) table.

```

[edit]
protocols {
    mpls {
        interface t3-0/2/0.0;
        interface at-0/1/0.0;
    }
    bgp {
        group pe-pe {
            type internal;
            local-address 10.255.14.171;
            family inet-vpn {
                any;
            }
            neighbor 10.255.14.173;
        }
    }
    isis {
        interface at-0/1/0.0;
        interface lo0.0 {
            passive;
        }
    }
    ldp {
        interface at-0/1/0.0;
    }
}
routing-instances {
    vpn-isp1 {

```

```
instance-type vrf;
interface t3-0/2/0.0;
route-distinguisher 10.255.14.171:21;
vrf-import vpn-isp1-import;
vrf-export vpn-isp1-export;
protocols {
  bgp {
    group to-isp1 {
      peer-as 21;
      neighbor 192.168.197.14 {
        family inet {
          labeled-unicast;
        }
      }
    }
  }
}
policy-options {
  policy-statement vpn-isp1-import {
    term a {
      from {
        protocol bgp;
        community vpn-isp1-comm;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement vpn-isp1-export {
    term a {
      from protocol bgp;
      then {
        community add vpn-isp1-comm;
        accept;
      }
    }
    term b {
      then reject;
    }
  }
  community vpn-isp1-comm members target:69:21;
}
```

Configuration for Router F

Configure Router F to act as a label-swapping router:

```
[edit]
protocols {
  isis {
    interface so-0/2/0.0;
    interface at-0/3/0.0;
```



```

        interface lo0.0 {
            passive;
        }
    }
    ldp {
        interface so-0/2/0.0;
        interface at-0/3/0.0;
    }
}

```

Configuration for Router G

Configure Router G to act as a label-swapping router:

```

[edit]
protocols {
    isis {
        interface so-0/0/0.0;
        interface so-1/0/0.0;
        interface lo0.0 {
            passive;
        }
    }
    ldp {
        interface so-0/0/0.0;
        interface so-1/0/0.0;
    }
}

```

Configuration for Router H

Router H acts as the PE router for AS 10023. The configuration that follows is similar to that for Router F:

```

[edit]
protocols {
    mpls {
        interface fe-1/1/0.0;
        interface so-1/0/0.0;
    }
    bgp {
        group pe-pe {
            type internal;
            local-address 10.255.14.173;
            family inet-vpn {
                any;
            }
            neighbor 10.255.14.171;
        }
    }
    isis {
        interface so-1/0/0.0;
        interface lo0.0 {
            passive;
        }
    }
}

```

```
ldp {
  interface so-1/0/0.0;
}
}
routing-instances {
  vpn-isp1 {
    instance-type vrf;
    interface fe-1/1/0.0;
    route-distinguisher 10.255.14.173:21;
    vrf-import vpn-isp1-import;
    vrf-export vpn-isp1-export;
    protocols {
      bgp {
        group to-isp1 {
          peer-as 21;
          neighbor 192.168.197.94 {
            family inet {
              labeled-unicast;
            }
          }
        }
      }
    }
  }
}
policy-options {
  policy-statement vpn-isp1-import {
    term a {
      from {
        protocol bgp;
        community vpn-isp1-comm;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement vpn-isp1-export {
    term a {
      from protocol bgp;
      then {
        community add vpn-isp1-comm;
        accept;
      }
    }
    term b {
      then reject;
    }
  }
  community vpn-isp1-comm members target:69:21;
}
```

Configuration for Router I

Configure Router I to connect to the basic Internet service customer (Router L):

```

[edit]
protocols {
  mpls {
    interface fe-1/0/1.0;
    interface fe-1/1/3.0;
  }
  bgp {
    group int {
      type internal;
      local-address 10.255.14.181;
      neighbor 10.255.14.177;
      neighbor 10.255.14.179;
      neighbor 10.255.14.175;
      neighbor 10.255.14.176;
      neighbor 10.255.14.178;
    }
    group to-vpn-green {
      peer-as 1;
      neighbor 192.168.197.198;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.0 {
        passive;
      }
      interface fe-1/0/1.0 {
        passive;
      }
      interface fe-1/1/3.0;
    }
  }
}

```

Configuration for Router J

Configure Router J as a label-swapping router:

```

[edit]
protocols {
  bgp {
    group int {
      type internal;
      local-address 10.255.14.178;
      neighbor 10.255.14.177;
      neighbor 10.255.14.181;
      neighbor 10.255.14.175;
      neighbor 10.255.14.176;
      neighbor 10.255.14.179;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.0 {
        passive;
      }
    }
  }
}

```

```
    }  
    interface fe-1/0/2.0;  
    interface fe-1/0/3.0;  
  }  
}
```

Configuration for Router K

Router K acts as the CE router at the end of the connection to the carrier provider. As in the configuration for Router D, include the **labeled-unicast** statement for the EBGP session:

```
[edit]  
protocols {  
  mpls {  
    interface fe-1/1/2.0;  
    interface fe-1/0/2.0;  
  }  
  bgp {  
    group int {  
      type internal;  
      local-address 10.255.14.177;  
      neighbor 10.255.14.181;  
      neighbor 10.255.14.178;  
      neighbor 10.255.14.175;  
      neighbor 10.255.14.176;  
      neighbor 10.255.14.179;  
    }  
    group to-isp-red {  
      export internal;  
      peer-as 10023;  
      neighbor 192.168.197.93 {  
        family inet {  
          labeled-unicast;  
        }  
      }  
    }  
  }  
  ospf {  
    area 0.0.0.0 {  
      interface lo0.0 {  
        passive;  
      }  
      interface fe-1/0/2.0;  
      interface fe-1/1/2.0 {  
        passive;  
      }  
    }  
  }  
}  
policy-options {  
  policy-statement internal {  
    term a {  
      from protocol [ ospf direct ];  
      then accept;  
    }  
    term b {
```

```

        then reject;
    }
}

```

Configuration for Router L

Configure Router L to act as the end customer for the carrier-of-carriers VPN service:

```

[edit]
protocols {
  bgp {
    group to-routerl {
      export attached;
      peer-as 21;
      neighbor 192.168.197.197;
    }
  }
}
policy-options {
  policy-statement attached {
    from protocol direct;
    then accept;
  }
}

```

Carrier-of-Carriers VPN Example—Customer Provides VPN Service

In this example, the carrier customer *must* run some form of MPLS (Resource Reservation Protocol [RSVP] or LDP) on its network to provide VPN services to the end customer. In the example below, Router B and Router I act as PE routers, and a functioning MPLS path is required between these routers if they exchange VPN-IPv4 routes.

For configuration information see the following sections:

- [Configuration for Router A on page 54](#)
- [Configuration for Router B on page 54](#)
- [Configuration for Router C on page 56](#)
- [Configuration for Router D on page 56](#)
- [Configuration for Router E on page 57](#)
- [Configuration for Router F on page 59](#)
- [Configuration for Router G on page 59](#)
- [Configuration for Router H on page 59](#)
- [Configuration for Router I on page 61](#)
- [Configuration for Router J on page 62](#)
- [Configuration for Router K on page 63](#)
- [Configuration for Router L on page 64](#)

Configuration for Router A

In this example, Router A acts as the CE router for the end customer. Configure a default **family inet** BGP session on Router A:

```
[edit]
protocols {
  bgp {
    group to-routerB {
      export attached;
      peer-as 21;
      neighbor 192.168.197.169;
    }
  }
}
policy-options {
  policy-statement attached {
    from protocol direct;
    then accept;
  }
}
```

Configuration for Router B

Because Router B is the PE router for the end customer CE router (Router A), you need to configure a routing instance (**vpna**). Configure the **labeled-unicast** statement on the IBGP session to Router D, and configure **family-inet-vpn** for the IBGP session to the other side of the network (see [Figure 5 on page 44](#)) with Router I:

```
[edit]
protocols {
  mpls {
    interface fe-1/0/2.0;
    interface fe-1/0/3.0;
  }
  bgp {
    group int {
      type internal;
      local-address 10.255.14.179;
      neighbor 10.255.14.175 {
        family inet {
          labeled-unicast {
            resolve-vpn;
          }
        }
      }
    }
    neighbor 10.255.14.181 {
      family inet-vpn {
        any;
      }
    }
  }
  ospf {
    area 0.0.0.0 {
```

```

        interface lo0.0 {
            passive;
        }
        interface fe-1/0/3.0;
    }
}
ldp {
    interface fe-1/0/3.0;
}
}
routing-instances {
    vpna {
        instance-type vrf;
        interface fe-1/0/2.0;
        route-distinguisher 10.255.14.179:21;
        vrf-import vpna-import;
        vrf-export vpna-export;
        protocols {
            bgp {
                group vpna-06 {
                    peer-as 1;
                    neighbor 192.168.197.170;
                }
            }
        }
    }
}
}
policy-options {
    policy-statement vpna-import {
        term a {
            from {
                protocol bgp;
                community vpna-comm;
            }
            then accept;
        }
        term b {
            then reject;
        }
    }
    policy-statement vpna-export {
        term a {
            from protocol bgp;
            then {
                community add vpna-comm;
                accept;
            }
        }
        term b {
            then reject;
        }
    }
}
community vpna-comm members target:100:1001;
}

```

Configuration for Router C

Configure Router C as a label-swapping router within the local AS:

```
[edit]
protocols {
  mpls {
    traffic-engineering bgp-igp;
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.0 {
        passive;
      }
      interface fe-0/3/3.0;
      interface fe-0/3/0.0;
    }
  }
  ldp {
    interface fe-0/3/0.0;
    interface fe-0/3/3.0;
  }
}
```

Configuration for Router D

Router D acts as the CE router for the VPN services provided by the AS 10023 network. In the BGP group configuration for group **int**, which handles traffic to Router B (10.255.14.179), you include the **labeled-unicast** statement. You also need to configure the BGP group **to-isp-red** to send labeled internal routes to the PE router (Router E).

```
[edit]
protocols {
  mpls {
    traffic-engineering bgp-igp;
    interface fe-0/3/0.0;
    interface t3-0/0/0.0;
  }
  bgp {
    group int {
      type internal;
      local-address 10.255.14.175;
      neighbor 10.255.14.179 {
        family inet {
          labeled-unicast;
        }
      }
    }
    group to-isp-red {
      export internal;
      peer-as 10023;
      neighbor 192.168.197.13 {
        family inet {
          labeled-unicast;
        }
      }
    }
  }
}
```



```

    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface fe-0/3/0.0;
  }
}
ldp {
  interface fe-0/3/0.0;
}
}
policy-options {
  policy-statement internal {
    term a {
      from protocol [ ospf direct ];
      then accept;
    }
    term b {
      then reject;
    }
  }
}
}

```

Configuration for Router E

Router E and Router H are PE routers. Configure a PE-router-to-PE-router BGP session to allow VPN-IPv4 routes to pass between these two PE routers. Configure the routing instance on Router E to send labeled routes to the CE router (Router D).

Configure Router E:

```

[edit]
protocols {
  mpls {
    interface t3-0/2/0.0;
    interface at-0/1/0.0;
  }
  bgp {
    group pe-pe {
      type internal;
      local-address 10.255.14.171;
      family inet-vpn {
        any;
      }
      neighbor 10.255.14.173;
    }
  }
  isis {
    interface at-0/1/0.0;
    interface lo0.0 {
      passive;
    }
  }
}

```

```
}
ldp {
  interface at-0/1/0.0;
}
}
policy-options {
  policy-statement vpn-isp1-import {
    term a {
      from {
        protocol bgp;
        community vpn-isp1-comm;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement vpn-isp1-export {
    term a {
      from protocol bgp;
      then {
        community add vpn-isp1-comm;
        accept;
      }
    }
    term b {
      then reject;
    }
  }
  community vpn-isp1-comm members target:69:21;
}
routing-instances {
  vpn-isp1 {
    instance-type vrf;
    interface t3-0/2/0.0;
    route-distinguisher 10.255.14.171:21;
    vrf-import vpn-isp1-import;
    vrf-export vpn-isp1-export;
    protocols {
      bgp {
        group to-isp1 {
          peer-as 21;
          neighbor 192.168.197.14 {
            as-override;
            family inet {
              labeled-unicast;
            }
          }
        }
      }
    }
  }
}
}
```

Configuration for Router F

Configure Router F to swap labels for routes running through its interfaces:

```
[edit]
protocols {
  isis {
    interface so-0/2/0.0;
    interface at-0/3/0.0;
    interface lo0.0 {
      passive;
    }
  }
  ldp {
    interface so-0/2/0.0;
    interface at-0/3/0.0;
  }
}
```

Configuration for Router G

Configure Router G:

```
[edit]
protocols {
  isis {
    interface so-0/0/0.0;
    interface so-1/0/0.0;
    interface lo0.0 {
      passive;
    }
  }
  ldp {
    interface so-0/0/0.0;
    interface so-1/0/0.0;
  }
}
```

Configuration for Router H

The configuration for Router H is similar to the configuration for Router E:

```
[edit]
protocols {
  mpls {
    interface fe-1/1/0.0;
    interface so-1/0/0.0;
  }
  bgp {
    group pe-pe {
      type internal;
      local-address 10.255.14.173;
      family inet-vpn {
        any;
      }
      neighbor 10.255.14.171;
    }
  }
}
```

```
    }
  }
  isis {
    interface so-1/0/0.0;
    interface lo0.0 {
      passive;
    }
  }
  ldp {
    interface so-1/0/0.0;
  }
}
routing-instances {
  vpn-isp1 {
    instance-type vrf;
    interface fe-1/1/0.0;
    route-distinguisher 10.255.14.173:21;
    vrf-import vpn-isp1-import;
    vrf-export vpn-isp1-export;
    protocols {
      bgp {
        group to-isp1 {
          peer-as 21;
          neighbor 192.168.197.94 {
            as-override;
            family inet {
              labeled-unicast;
            }
          }
        }
      }
    }
  }
}
policy-options {
  policy-statement vpn-isp1-import {
    term a {
      from {
        protocol bgp;
        community vpn-isp1-comm;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
  policy-statement vpn-isp1-export {
    term a {
      from protocol bgp;
      then {
        community add vpn-isp1-comm;
        accept;
      }
    }
    term b {
```

```

        then reject;
    }
}
community vpn-isp1-comm members target:69:21;
}

```

Configuration for Router I

Router I acts as the PE router for the end customer. The configuration that follows is similar to the configuration for Router B:

```

[edit]
protocols {
  mpls {
    interface fe-1/0/1.0;
    interface fe-1/1/3.0;
  }
  bgp {
    group int {
      type internal;
      local-address 10.255.14.181;
      neighbor 10.255.14.177 {
        family inet {
          labeled-unicast {
            resolve-vpn;
          }
        }
      }
      neighbor 10.255.14.179 {
        family inet-vpn {
          any;
        }
      }
    }
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.0 {
        passive;
      }
      interface fe-1/1/3.0;
    }
  }
  ldp {
    interface fe-1/1/3.0;
  }
}
routing-instances {
  vpna {
    instance-type vrf;
    interface fe-1/0/1.0;
    route-distinguisher 10.255.14.181:21;
    vrf-import vpna-import;
    vrf-export vpna-export;
    protocols {
      bgp {

```

```
        group vpn-a-0 {
            peer-as 1;
            neighbor 192.168.197.198;
        }
    }
}
}
policy-options {
    policy-statement vpn-a-import {
        term a {
            from {
                protocol bgp;
                community vpn-a-comm;
            }
            then accept;
        }
        term b {
            then reject;
        }
    }
    policy-statement vpn-a-export {
        term a {
            from protocol bgp;
            then {
                community add vpn-a-comm;
                accept;
            }
        }
        term b {
            then reject;
        }
    }
    community vpn-a-comm members target:100:1001;
}
```

Configuration for Router J

Configure Router J to swap labels for routes running through its interfaces:

```
[edit]
protocols {
    mpls {
        traffic-engineering bgp-igp;
    }
    ospf {
        area 0.0.0.0 {
            interface lo0.0 {
                passive;
            }
            interface fe-1/0/2.0;
            interface fe-1/0/3.0;
        }
    }
    ldp {
        interface fe-1/0/2.0;
```

```

        interface fe-1/0/3.0;
    }
}

```

Configuration for Router K

The configuration for Router K is similar to the configuration for Router D:

```

[edit]
protocols {
  mpls {
    traffic-engineering bgp-igp;
    interface fe-1/1/2.0;
    interface fe-1/0/2.0;
  }
  bgp {
    group int {
      type internal;
      local-address 10.255.14.177;
      neighbor 10.255.14.181 {
        family inet {
          labeled-unicast;
        }
      }
    }
    group to-isp-red {
      export internal;
      peer-as 10023;
      neighbor 192.168.197.93 {
        family inet {
          labeled-unicast;
        }
      }
    }
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.0 {
        passive;
      }
      interface fe-1/0/2.0;
    }
  }
  ldp {
    interface fe-1/0/2.0;
  }
}
policy-options {
  policy-statement internal {
    term a {
      from protocol [ ospf direct ];
      then accept;
    }
    term b {
      then reject;
    }
  }
}

```

```
}  
}
```

Configuration for Router L

In this example, Router L is the end customer's CE router. Configure a default family **inet** BGP session on Router L:

```
[edit]  
protocols {  
  bgp {  
    group to-l {  
      export attached;  
      peer-as 21;  
      neighbor 192.168.197.197;  
    }  
  }  
}  
policy-options {  
  policy-statement attached {  
    from protocol direct;  
    then accept;  
  }  
}
```

Multiple Instances for LDP and Carrier-of-Carriers VPNs

By configuring multiple LDP routing instances, you can use LDP to advertise labels in a carrier-of-carriers VPN from a core provider PE router to a customer carrier CE router. Having LDP advertise labels in this manner is especially useful when the carrier customer is a basic ISP and wants to restrict full Internet routes to its PE routers. By using LDP instead of BGP, the carrier customer shields its other internal routers from the Internet at large. Multiple-instance LDP is also useful when a carrier customer wants to provide Layer 3 VPN or Layer 2 VPN services to its customers.

For an example of how to configure multiple LDP routing instances for carrier-of-carriers VPNs, see the *Junos Feature Guide* on the product documentation page of the Juniper Networks website, located at <http://www.juniper.net/>.

CHAPTER 4

Interprovider and Carrier-of-Carriers VPNs Configuration Statements

- [labeled-unicast \(Protocols BGP VPN\) on page 66](#)
- [per-group-label on page 67](#)
- [traffic-statistics \(Protocols BGP\) on page 67](#)

labeled-unicast (Protocols BGP VPN)

Syntax	<pre>labeled-unicast { aggregate-label { community <i>community-name</i>; } explicit-null { connected-only; } per-group-label; resolve-vpn; traffic-statistics { file <i>filename</i> <world-readable no-world-readable>; interval <i>seconds</i>; } }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family (inet inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family (inet inet6)],</p> <p>[edit protocols bgp family (inet inet6)],</p> <p>[edit protocols bgp group <i>group-name</i> family (inet inet6)]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Advertise labeled routes from the inet.0 VPN, and place labeled routes into the inet.0 VPN. When the labeled-unicast statement is used, the local router automatically performs a next hop to self on all routes advertised into EBGp from IBGP and from IBGP to EBGp.</p>
Options	<p>resolve-vpn—(Optional) Store labeled routes in the inet.3 or inet6.3 routing table to resolve routes for a provider edge (PE) router located in a different autonomous system (AS). For a PE router to install a route in the virtual routing and forwarding (VRF) table, the next hop must resolve to a route stored in the inet.3 or inet6.3 routing table. This option is also used to configure inter-AS VPLS with MAC operations.</p> <p>The other statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Carrier-of-Carriers VPNs for Customers That Provide Internet Service on page 14 • Configuring Carrier-of-Carriers VPNs for Customers That Provide VPN Service on page 19 • Configuring Interoperability Between BGP Signaling and LDP Signaling in VPLS • Understanding Interprovider VPNs on page 13

per-group-label

Syntax	<code>per-group-label;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp family inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit protocols bgp family inet labeled-unicast], [edit protocols bgp group <i>group-name</i> family inet labeled-unicast]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Account for traffic from each customer separately by advertising separate labels for the same prefix to the peer routers in the BGP groups.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics on page 26

traffic-statistics (Protocols BGP)

Syntax	<pre>traffic-statistics { file <i>filename</i> <world-readable no-world-readable>; interval <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp family (inet inet6) labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family (inet inet6) labeled-unicast], [edit protocols bgp family (inet inet6) labeled-unicast], [edit protocols bgp group <i>group-name</i> family (inet inet6) labeled-unicast]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable the collection of traffic statistics for interprovider or carrier-of-carriers VPNs.
Options	<p>file <i>filename</i>—Specify a filename for the BGP labeled-unicast traffic statistics file. If you do not specify a filename, statistics are still collected but can only be viewed by using the <code>show bgp group traffic statistics <i>group-name</i></code> command.</p> <p>interval <i>seconds</i>—Specify how often BGP labeled-unicast traffic statistics are collected.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BGP to Gather Interprovider and Carrier-of-Carriers VPNs Statistics on page 26

PART 3

Administration

- [Interprovider and Carrier-of-Carriers VPNs Reference on page 71](#)

CHAPTER 5

Interprovider and Carrier-of-Carriers VPNs Reference

- [Example Terminology on page 71](#)
- [Supported Carrier-of-Carriers and Interprovider VPN Standards on page 72](#)

Example Terminology

B

bgp.l3vpn.0 The table on the provider edge (PE) router in which the VPN-IPv4 routes that are received from another PE router are stored. Incoming routes are checked against the **vrf-import** statements from all the VPNs configured on the PE router. If there is a match, the VPN–Internet Protocol version 4 (IPv4) route is added to the bgp.l3vpn.0 table. To view the bgp.l3vpn.0 table, issue the **show route table bgp.l3vpn.0** command.

M

MP-EBGP The multiprotocol external BGP (MP-EBGP) mechanism is used to export VPN-IPv4 routes across an autonomous system (AS) boundary. To apply this mechanism, use the **labeled-unicast** statement at the **[edit protocols bgp group group-name family inet]** hierarchy level.

R

- routing-instance-name.**
inet.0
- The routing table for a specific routing instance. For example, a routing instance called VPN-A has a routing table called VPN-A.inet.0. Routes are added to this table in the following ways:
 - They are sent from a customer edge (CE) router configured within the VPN-A routing instance.
 - They are advertised from a remote PE router that passes the **vrf-import** policy configured within VPN-A (to view the route, run the **show route** command). IPv4 (not VPN-IPv4) routes are stored in this table.

V

- vrf-export policy-name**
- An export policy configured on a particular routing instance on a PE router. It is required for the configuration of interprovider and carrier-of-carriers VPNs. It is applied to VPN-IPv4 routes (originally learned from locally connected CE routers as IPv4 routes), which are advertised to another PE router or route reflector.
- vrf-import policy-name**
- An import policy configured on a particular routing instance on a PE router. This policy is required for the configuration of interprovider and carrier-of-carriers VPNs. It is applied to VPN-IPv4 routes learned from another PE router or a route reflector.

Supported Carrier-of-Carriers and Interprovider VPN Standards

Junos OS substantially supports the following RFCs, which define standards for carrier-of-carriers and interprovider virtual private networks (VPNs).

- RFC 3107, *Carrying Label Information in BGP-4*
- RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*
Supported on MX Series routers with the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP.
- RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture*
Supported on MX Series routers with the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP.
- RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 5601, *Pseudowire (PW) Management Information Base (MIB)*
- RFC 5603, *Ethernet Pseudowire (PW) Management Information Base (MIB)*
- RFC 6368, *Internal BGP as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*

- Related Documentation**
- *Supported VPWS Standards*
 - *Supported Layer 2 VPN Standard*
 - *Supported Layer 3 VPN Standards*

- *Supported Multicast VPN Standards*
- *Supported VPLS Standards*
- *Supported BGP Standards*
- *Accessing Standards Documents on the Internet*

PART 4

Index

- [Index on page 77](#)

Index

Symbols

#, comments in configuration statements.....	xiv
(), in syntax descriptions.....	xiv
< >, in syntax descriptions.....	xiv
[], in configuration statements.....	xiv
{ }, in configuration statements.....	xiv
(pipe), in syntax descriptions.....	xiv

B

braces, in configuration statements.....	xiv
brackets	
angle, in syntax descriptions.....	xiv
square, in configuration statements.....	xiv

C

carrier-of-carriers VPNs	
overview.....	3
statistics.....	26
comments, in configuration statements.....	xiv
conventions	
text and syntax.....	xiii
curly braces, in configuration statements.....	xiv
customer support.....	xv
contacting JTAC.....	xv

D

documentation	
comments on.....	xv

F

font conventions.....	xiii
-----------------------	------

I

interprovider VPNs	
overview.....	3
statistics gathering.....	26

L

labeled-unicast statement.....	66
VPNs	
usage guidelines.....	15, 17, 20, 22, 26

Layer 3 VPNs	
next generation.....	5

M

manuals	
comments on.....	xv

N

next generation Layer 3 VPNs.....	5
-----------------------------------	---

P

parentheses, in syntax descriptions.....	xiv
per-group-label statement.....	67
usage guidelines.....	26

S

support, technical	See technical support
syntax conventions.....	xiii

T

technical support	
contacting JTAC.....	xv
traffic-statistics statement.....	67
VPNs	
usage guidelines.....	26

V

VPNs	
carrier-of-carriers	
supported software standards.....	72
interprovider	
supported software standards.....	72

