



CHAPTER 5

MPLS VPN Service Policies

This chapter describes how to use the Cisco IP Solution Center (ISC) GUI to define MPLS VPN Service Policies. It contains the following major sections:

- [Service Policy Overview, page 5-1](#)
- [Defining an MPLS VPN Service Policy, page 5-2](#)
- [Specifying PE and CE Interface Parameters, page 5-4](#)
- [Specifying the IP Address Scheme, page 5-8](#)
- [Specifying the Routing Protocol for a Service, page 5-11](#)
- [Defining VRF and VPN Information, page 5-29](#)
- [Enabling Template Association for a Policy, page 5-34](#)

Service Policy Overview

Provisioning an MPLS VPN begins with defining a service policy. A service policy can be applied to multiple PE-CE links in a single service request. A *network operator* defines service policies. A *service operator* uses a service policy to create service requests. Each service request contains a list of PE-CE links. When a service operator creates a service request, the operator sees only the policy information required to be completed. All the other necessary information is filled in by the service policy itself (as well as the Auto Discovery process).

Service Policy Editor

When you define a service policy for ISC, you are presented with a series of dialog boxes that allow you to specify the parameters for each major category required to complete an MPLS service request. The Service Policy editor presents three columns: **Attribute**, **Value**, and **Editable**:

- **Attribute**

The Attribute column displays the names of each parameter that you need to define for each major category (for example, IP addresses or routing protocols).

- **Value**

The Value column displays the fields and other selectable items that correspond to each parameter and option.

The type of dialog box that is invoked when you edit an attribute depends on the type of attribute. In some cases, the value is a simple string value or integer value, in which case a single text entry field appears. In other cases, the value is complex or consists of multiple values, such as an IP address. In these cases, a dialog box appears so you can specify the required values. The values you enter are validated; when invalid values are entered, you receive notification of the invalid values. In other cases, you will be presented with check boxes that will allow you to enable or disable a particular option.



Note In some cases, changing an attribute's value results in invalidating the values of related attributes. For example, changing the PE interface name can result in invalidating the PE encapsulation value. When this occurs, the service policy editor removes the invalid values and you will need to reset them appropriately.

There is a parent-child relationship between some attributes. In these cases, changing the value of a parent attribute can enable or disable the child attributes. For example, changing the value of the PE encapsulation could result in enabling or disabling the DLCI (data link connection identifier), VLAN ID, ATM circuit identifiers, and the tunnel source and destination address attributes.

- **Editable**

The Editable column allows the network operator to indicate the attributes that are likely to change across multiple service requests. When attributes are checked as editable, only those attributes will be made available to the service operator when creating or modifying service requests with that service request policy.

When an attribute category is set to be editable, all the related and child attributes are also editable attributes.

About IP Addresses in Cisco ISC

Within a VPN (or extranet), all IP addresses must be unique. Customer IP addresses are not allowed to overlap with provider IP addresses. Overlap is possible only when two devices cannot see each other; that is, when they are in isolated, non-extranet VPNs.

The ISC MPLS VPN software assumes that it has an IP address pool to draw addresses from. The only way to guarantee that the product can use these addresses freely is if they are provider IP addresses.

Predefining a unique section (or sections) of IP address space for the PE-CE links is the only way to ensure stable security. Thus, because of the security and maintenance issues, Cisco does not recommend using customer IP addresses on the PE-CE link.

Defining an MPLS VPN Service Policy

The remaining sections in this chapter provide an extended example of defining an MPLS service policy for a PE-CE link. This is to demonstrate the various steps involved in defining an MPLS service policy. The steps can be used as the basis for defining other types of MPLS VPN service policies. Additional types of MPLS VPN policies are described in other chapters in this guide.

To begin defining an MPLS VPN service policy for PE-CE link, perform the following steps.

Step 1 Click the **Service Design** tab.

Step 2 Choose **Policies**.

The Policies window appears.

Step 3 From the **Create** drop-down list, choose **MPLS Policy**.

The MPLS Policy Type dialog box appears, as shown in [Figure 5-1](#).

Figure 5-1 Defining the MPLS Service Policy

Attribute	Value
Policy Name *	mpls_pe_ce
Policy Owner *	<input checked="" type="radio"/> Customer <input type="radio"/> Provider <input type="radio"/> Global Policy
Customer *	Select
Policy Type *	<input checked="" type="radio"/> Regular: PE-CE <input type="radio"/> MVRFCE: PE-CE
CE Present *	<input checked="" type="checkbox"/>

Step 4 Enter a **Policy Name** for the MPLS policy.**Step 5** Choose the **Policy Owner**.

There are three types of MPLS policy ownership:

- Customer ownership
- Provider ownership
- Global ownership: Any service operator can make use of this MPLS policy.

This ownership has relevance when the ISC Role-Based Access Control (RBAC) comes into play. For example, an MPLS policy that is customer-owned can only be seen by operators who are allowed to work on this customer-owned policy.

Similarly, operators who are allowed to work on a provider's network can view, use, and deploy a particular provider-owned policy.



Note For Cable (PE-NoCE), policy ownership should be set to Provider.

Step 6 Click **Select** to choose the owner of the MPLS policy. (If you choose Global ownership, the Select function is not available.)

The Select Customer window or the Select Provider window appears and you can choose an owner of the policy and click **Select**.

Step 7 Choose the **Policy Type** of the MPLS policy.

There are two policy types for MPLS policies:

- Regular PE-CE: PE-to-CE link
- MVRFCE PE-CE: PE to CE link using the Multi-VRF feature for the PE

- Step 8** Check the **CE Present** check box if you want ISC to ask the service operator who uses this MPLS policy to provide a CE router and interface during service activation. The default is CE present in the service. If you do not check the **CE Present** check box, ISC asks the service operator, during service activation, only for the PE-CLE or the PE-POP router and customer-facing interface.
- Step 9** Click **Next**.
-

To continue with the example, see the following section, [Specifying PE and CE Interface Parameters](#), page 5-4.

Specifying PE and CE Interface Parameters

The MPLS Policy Interface dialog box appears, as shown in [Figure 5-2](#).



Tip

You do not have to choose a specific interface type for the PE and CE at this point. Notice that the fields are set by default to **Editable**. With the interface parameters set to **Editable**, the service operator can specify the exact interface type and format when he or she creates the service request.

If you want to specify the device interface information for this service policy when the service request is created, leave the fields as they are currently set by default, then click **Next**.

Figure 5-2 Specifying the PE UNI Security, and CE Interface Parameters

Attribute	Value	Editable
Reset All Attribute Editable Flags:		<input checked="" type="checkbox"/>
PE Information		
Interface Type:	ANY <input type="button" value="v"/>	
Interface Format:	<input type="text"/>	
Interface Description:	<input type="text"/>	<input checked="" type="checkbox"/>
Shutdown Interface:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Auto-Pick VLAN ID:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Use SVI:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
ETTH Support:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Standard UNI Port:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
UNI Security Information		
Disable CDP:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Filter BPDUs:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Use Existing ACL Name:	<input type="checkbox"/>	
UNI MAC Addresses:	<input type="button" value="Edit"/>	<input checked="" type="checkbox"/>
UNI Port Security:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Maximum MAC Address:	<input type="text"/> (1 - 5120)	<input checked="" type="checkbox"/>
Aging (in minutes):	<input type="text"/> (0 - 1440)	<input checked="" type="checkbox"/>
Violation Action:	PROTECT <input type="button" value="v"/>	<input checked="" type="checkbox"/>
Secure MAC Addresses:	<input type="button" value="Edit"/>	<input checked="" type="checkbox"/>
CE Information		
Interface Type:	ANY <input type="button" value="v"/>	
Interface Format:	<input type="text"/>	
Interface Description:	<input type="text"/>	<input checked="" type="checkbox"/>

To specify the PE, UNI Security, and CE interface information for this MPLS policy:

PE Interface Information

Step 1 Interface Type: From the drop-down list, choose the interface type for the PE.

Cisco IP Solution Center supports the following interface types (for both PEs and CEs):

- Any
- ATM (Asynchronous Transfer Mode)
- BRI (Basic Rate Interface)
- Ethernet
- Fast Ethernet
- FDDI (Fiber Distributed Data Interface)
- GE-WAN (Gigabit Ethernet WAN)
- Gigabit Ethernet

- HSSI (High Speed Serial Interface)
- Loopback
- MFR
- MultiLink
- PoS (Packet over Sonet)
- Port-Channel
- Serial
- Switch
- Tunnel
- VLAN

Step 2 Interface Format: Optionally, you can specify the slot number and port number for the PE interface.

Specify the format in the standard nomenclature: **slot number/port number** (for example, **1/0** indicates that the interface is located at slot 1, port 0).

This is especially useful to specify here if you know that the link will always go through a particular interface's slot/port location on all or most of the network devices in the service. If this parameter is left editable, it can be changed when the service operator creates the service request.

You can also specify the Interface Format as a Channelized Interface:

- **slot/subSlot/port** (for example, **2/3/4** indicates that the interface is located at Serial 2/3/4)
- **slot/subSlot/port/T1#:channelGroup#** (for example, **2/0/4/6:8** indicates that the interface is located at Serial 2/0/4/6:8)
- **slot/subSlot/port.STS-1Path/T1#:channelGroup#** (for example, **2/0/0.1/6:8** indicates that the interface is located at Serial 2/0/0.1/6:8)

Step 3 Interface Description: Optionally, you can enter a description of the PE interface.

Step 4 Shutdown Interface: When you check this check box, the specified PE interface is configured in a shut down state.

Step 5 Encapsulation: Choose the encapsulation used for the specified PE interface type.

When you choose an interface type, the Encapsulation field displays a drop-down list of the supported encapsulation types for the specified interface type.

[Table 5-1](#) shows the protocol encapsulations available for each of the supported interface types.

Table 5-1 Interface Types and Their Corresponding Encapsulations

Interface Type	Encapsulations
ATM	AAL5SNAP
BRI	Frame-Relay, Frame-Relay-ietf, HDLC (High-Level Data Link Control), PPP (Point-to-Point Protocol). Frame-Relay-ietf sets the encapsulation method to comply with the Internet Engineering Task Force (IETF) standard (RFC 1490). Use this method when connecting to another vendor's equipment across a Frame Relay network.
Ethernet	Default frame, dot1q (802.1Q)
Fast Ethernet	Default frame, ISL (Inter-Switch Link), dot1q (802.1Q)

Table 5-1 Interface Types and Their Corresponding Encapsulations (continued)

Interface Type	Encapsulations (continued)
FDDI (Fiber Distributed Data Interface)	None
Gigabit Ethernet	Default frame, ISL (Inter-Switch Link), dot1q (802.1Q)
Gigabit Ethernet WAN	Default frame, ISL (Inter-Switch Link), dot1q (802.1Q)
HSSI (High Speed Serial Interface)	Frame-Relay, Frame-Relay-ietf, HDLC (High-Level Data Link Control), PPP (Point-to-Point Protocol)
Loopback	None.
MFR	Frame-Relay, Frame-Relay-ietf, HDLC (High-Level Data Link Control), PPP (Point-to-Point Protocol).
MultiLink	PPP (Point-to-Point Protocol)
Port-Channel	Default frame, ISL (Inter-Switch Link), dot1q (802.1Q) NOTE: [Andrew to provide content]
POS (Packet Over Sonet)	Frame-Relay, HDLC (High-Level Data Link Control), PPP (Point-to-Point Protocol)
Serial	Frame-Relay, Frame-Relay-ietf, HDLC (High-Level Data Link Control), PPP (Point-to-Point Protocol)
Switch	AAL5SNAP
Tunnel	GRE (Generic Routing Encapsulation) - GRE is not supported in this release. -
VLAN	None

Step 6 Auto-Pick VLAN ID: Check this check box to have ISC automatically pick the VLAN ID.



Note If Auto-Pick VLAN ID is unchecked, you are prompted to enter the VLAN ID during the creation of the service request based on the policy.

Step 7 Use SVI: Check this check box to have ISC terminate VRF on SVI.

Step 8 ETTH Support: Check this check box to configure Ethernet-To-The-Home (ETTH). For an explanation of ETTH, see [Ethernet-To-The-Home \(ETTH\)](#), page 12-9.

Step 9 Standard UNI Port: Check this check box to access UNI Security Parameters:

UNI Security Information

Step 10 Disable CDP: Check this check box to disable CDP.

Step 11 Filter BPDU: Check this check box to filter BPDU.

Step 12 Use existing ACL Name: Check this check box to use existing ACL name.

Step 13 UNI MAC Addresses: Click **Edit** to modify or create a MAC address record.

Step 14 UNI Port Security: Check this check box to access UNI Port Security parameters:

- a. **Maximum MAC Address:** Enter a valid value.
- b. **Aging (in minutes):** Enter a valid value.

- c. **Violation Action:** From the drop-down list, choose one of the following:
 - PROTECT
 - RESTRICT
 - SHUTDOWN
- d. **Secure MAC Address:** Click **Edit** to modify or create a secure MAC address record.

CE Interface Information

- Step 15 Interface Type:** From the drop-down list, choose the interface type for the CE.
- Step 16 Interface Format:** Optionally, you can specify the slot number and port number for the CE interface.
- Step 17 Interface Description:** Optionally, you can enter a description of the CE interface.
- Step 18 Encapsulation:** Choose the encapsulation used for the specified CE interface type.
- Step 19** When satisfied with the interface settings, click **Next**.

To continue with the example, see the following section, [Specifying the IP Address Scheme, page 5-8](#).

Specifying the IP Address Scheme

The MPLS Policy Interface Address Selection window appears, as shown in [Figure 5-3](#).

Figure 5-3 Specifying the IP Address Scheme

Attribute	Value	Editable
PE-MVRFCPE Interface Address/Mask		
IP Numbering Scheme:	IPv4 Numbered	<input checked="" type="checkbox"/>
Extra MVRFCPE Loopback Required:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Automatically Assign IP Addresses:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
IP Address Pool:	Region Pool	<input checked="" type="checkbox"/>

To specify the IP address scheme you want to use for this service policy, perform the following steps.

- Step 1** Define the IP addressing scheme that is appropriate for the PE-CE link.

IP Numbering Scheme

You can choose from the following options.

- **IPv4 Numbered**

If you choose **IPv4 Numbered** and also check the **Automatically Assign IP Address** check box, ISC: MPLS checks for the presence of the corresponding IP addresses in the router's configuration file. If the addresses are present and they are in the same subnet, ISC uses those addresses (and does not allocate them from the address pool). If the IP addresses are not present in the configuration file, ISC picks IPv4 addresses from a /30 subnet point-to-point IP address pool.

- **IPv4 Unnumbered**

IPv4 addresses are drawn from the loopback IPv4 address pool. An unnumbered IPv4 address means that each interface “borrows” its address from another interface on the router (usually the loopback interface). Unnumbered addresses can only be used on point-to-point WAN links (such as Serial, Frame, and ATM), not on LAN links (such as Ethernet). If using IP unnumbered, then both the PE and CE must use the same IP unnumbered addressing scheme. When you choose **IPv4 Unnumbered**, ISC: MPLS creates a static route for the PE-CE link.

When you choose **IPv4 Unnumbered**, ISC: MPLS automatically creates a loopback interface (unless a loopback interface already exists with the correct attributes). For related information, see [Using Existing Loopback Interface Number, page 5-10](#).

- **IPv6 Numbered**

This addressing scheme is provided to support a 6VPE router. See [Chapter 4, “IPv6 and 6VPE Support in MPLS VPN”](#) for more information on IPv6 and 6VPE support in MPLS VPN management.



Note This option only appears if the policy type is a regular PE-CE policy.

- **IPv4+IPv6 Numbered**

In the case of a 6VPE device, the PE interface can be “dual stacked,” meaning it can contain both IPv4 and IPv6 addresses. In later steps, you will be able to enter the routing information independently for both IPv4 and IPv6. See [Chapter 4, “IPv6 and 6VPE Support in MPLS VPN”](#) for more information on IPv6 and 6VPE support in MPLS VPN management.



Note This option only appears if the policy type is a regular PE-CE policy.

Step 2 Indicate whether an extra loopback interface is required for the CE.

Extra CE Loopback Required

Even though a numbered IP address does not require a loopback address, ISC software provides the option to specify that an extra CE loopback interface is required. This option places an IP address on a CE router that is not tied to any physical interface.

If you enable **Extra CE Loopback Required**, you can enter the CE loopback address.

Step 3 Specify whether you want to automatically assign IP addresses.

Automatically Assign IP Address

If you choose **IPv4 Unnumbered** and also check the **Automatically Assign IP Address** check box, ISC picks two IP addresses from a /32 subnet point-to-point IP address pool.

If you choose **IPv4 Numbered** and also check the **Automatically Assign IP Address** check box, ISC checks for the presence of the corresponding IP addresses in the router’s configuration file. If the addresses are present and they are in the same subnet, ISC uses those addresses (and does not allocate them from the address pool). If the IP addresses are not present in the configuration file, ISC picks IP addresses from a /30 subnet point-to-point IP address pool.



Note This option is not supported for the **IPv6 Numbered** and **IPv4+IPv6 Numbered** address schemes.

Step 4 Specify the IP address pool and its associated Region for this service policy.

IP Address Pool

The IP Address Pool option gives the service operator the ability to have ISC automatically allocate IP addresses from the IP address pool attached to the Region. Prior to defining this aspect of the service policy, the Region must be defined and the appropriate IP address pools assigned to the Region.

You can specify IP address pool information for point-to-point (IP numbered) PE-CE links.

IP unnumbered addresses are drawn from the loopback IP address pool. An unnumbered IP address means that each interface “borrows” its address from another interface on the router (usually the loopback interface). Unnumbered addresses can only be used on point-to-point WAN links (such as Serial, Frame, and ATM), not on LAN links (such as Ethernet). If using IP unnumbered, then both the PE and CE must use the same IP unnumbered addressing scheme.



Note This option is not supported for the IPv6 Numbered and IPv4+IPv6 Numbered address schemes.

Step 5 When satisfied with the IP address scheme, click **Next**.

Using Existing Loopback Interface Number

On each PE, there is usually only one loopback interface number per VRF for interfaces using IP unnumbered addresses. However, if provisioning an interface using IP unnumbered addresses and manually assigned IP addresses, it is possible to have more than one loopback interface number under the same VRF. When using automatically-assigned IP addresses for provisioning IP unnumbered addresses, ISC associates the first loopback number with the same VRF name to the interface. If no loopback number already exists, ISC creates one.

If a service provider wants ISC to use an existing loopback interface number (for example, Loopback0), the service provider must modify the loopback interface description line in the configuration files for the pertinent routers (PE or CE).

To use the existing loopback interface number, you must modify the loopback interface description line so that it includes the keyword **VPN-SC**, as shown in the following example of a router configuration file.



Note When using an existing loopback interface number on a PE, an additional command line with the **ip vrf forwarding VRF_name** command must be included directly after the “description” line.

```
interface Loopback0
description by VPN-SC
ip vrf forwarding <VRF_name> ; This line is required on the PE only
ip address 209.165.202.129 255.255.255.224
```

You can use an existing loopback interface number only when the interface configuration meets these conditions: it must be a WAN serial interface using IP unnumbered addresses.

ISC selects loopback interface numbers by sequence. ISC uses the first loopback interface number that meets the requirement—for a CE, it is inclusion of the VPN-SC keyword; for a PE, it is the matching VRF name.

For example, if loopback1 and loopback2 include the VPN-SC keyword, but loopback3 does not, adding the VPN-SC keyword to loopback3 will not force ISC to choose loopback3 for the unnumbered interface when using automatically assigned addresses. Loopback1 will be chosen instead. The only way to choose a specific loopback interface number is to use a manually assigned IP address that matches the desired loopback interface number.

**Note**

Unlike standard interfaces, when loopback interfaces are provisioned in ISC, the resulting configuration file does not include a service request (SR) ID number. This is because multiple interfaces or service requests can use the same loopback interface.

To continue with the example, see the following section, [Specifying the Routing Protocol for a Service, page 5-11](#).

Specifying the Routing Protocol for a Service

You can now specify the routing protocol information for this service policy, as shown in [Figure 5-4](#).

**Note**

IPv4 and IPv6 routing are independent. The ISC GUI allows you to input the same or different routing protocols for IPv4 and IPv6, depending upon which addressing scheme you selected. Not all routing protocols are supported for IPv6. See [Chapter 4, “IPv6 and 6VPE Support in MPLS VPN”](#) for more information IPv6 and supported routing protocols.

The routing protocol you choose must run on both the PE and the CE. You can choose any one of the following protocols:

- **Static**—Specifies a static route (see [Static Protocol Chosen, page 5-12](#)).
- **RIP**—Routing Information Protocol (see [RIP Protocol Chosen, page 5-14](#)).
- **BGP**—Border Gateway Protocol (see [BGP Protocol Chosen, page 5-18](#)).
- **OSPF**—Open Shortest Path First (see [OSPF Protocol Chosen, page 5-21](#)).
- **EIGRP**—Enhanced Interior Gateway Routing Protocol (see [EIGRP Protocol Chosen, page 5-24](#)).
- **None**. Specifies parameters for cable services (see [None Chosen: Cable Services, page 5-28](#)).

To specify a routing protocol for the PE-CE link, perform the following steps.

Step 1 Choose the appropriate protocol from the Routing Protocol drop-down list.

**Note**

For IPv6, only Static, BGP, EIGRP and None are supported and are available in the drop-down list.

When you choose a particular routing protocol, the related parameters for that protocol are displayed.

Step 2 Enter the required information for the selected routing protocol, then click **Next**.

Step 3 Define the MPLS Policy VRF and VPN Selection parameters as described in [Defining VRF and VPN Information, page 5-29](#).

Redistribution of IP Routes

Route redistribution is the process of taking routing information from one source and importing that information into another source. Redistribution should be approached with caution. When you perform route redistribution, you lose information. Metrics must be arbitrarily reset. For example, if a group of RIP routes with a metric of five hops is redistributed into IGRP, there is no way to translate the five hop RIP metric into the composite metric of IGRP. You must arbitrarily choose a metric for the RIP routes as they are redistributed into IGRP. Also, when redistribution is performed at two or more points between two dynamic routing protocol domains, routing loops can occur.

CSC Support

To define a Service Policy with Carrier Supporting Carrier (CSC), choose the CSC Support check box from the MPLS Policy Editor - Routing Information. When CSC Support is checked, the CSC functionality is enabled to the MPLS VPN service. Provisioning CSC is explained in [Chapter 11](#), “Provisioning Carrier Supporting Carrier.”

Giving Only Default Routes to CE

When you enable the **Give only default routes to CE** option, you indicate whether the site needs *full routing* or *default routing*. Full routing is when the site must know specifically which other routes are present in the VPN. Default routing is when it is sufficient to send all packets that are not specifically for your site to the VPN.

If you choose this option, ISC configures **the default-info originate** command on the PE router under the running protocol (for RIP, OSPF, or EIGRP). For Static, ISC configures an **ip route 0.0.0.0 0.0.0.0 <out-going interface name>** command on the CE router.

A device can only have one default route. Therefore, the VPN can use a default route, but only on condition that the customer site does not already have a different one. The most common reason to already have a default route is that the site has an Internet feed that is independent of the VPN.

If the CE site already has Internet service, the CE can either route all packets to unknown destinations to the Internet or learn all the routes in the Internet. The obvious choice is to route all packets to unknown destinations to the Internet. If a site has an Internet feed, it might already have a default route. Under such conditions, setting the VPN as the default route is incorrect; the VPN should only route packets meant for other VPN sites.

Static Protocol Chosen

Static routing refers to routes to destinations that are listed manually in the router. Network reachability in this case is not dependent on the existence and state of the network itself. Whether a destination is up or down, the static routes remain in the routing table and traffic is still sent to that destination.

When you choose **Static** as the protocol, four options are enabled: **CSC Support**, **Give Only Default Routes to CE**, **Redistribute Connected (BGP only)**, and **Default Information Originate (BGP only)**, as shown in [Figure 5-4](#).

**Note**

Two other options (**AdvertisedRoutes** and **Default Routes - Routes to reach other sites**) are available when you create the service request. See [Setting Static Routing Protocol Attributes \(for IPv4 and IPv6\)](#), page 6-13.

Figure 5-4 Specifying the Static Routing Protocol

MPLS Policy Editor - Ipv4 Routing Information		
Attribute	Value	Editable
PE-CE Ipv4 Routing Information		
Routing Protocol	STATIC ▾	<input checked="" type="checkbox"/>
CsC Support:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Give Only Default Routes to CE:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Redistribute Connected (BGP only):	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Default Information Originate (BGP only):	<input type="checkbox"/>	<input checked="" type="checkbox"/>

211630

To specify Static as the routing protocol for the service policy, perform the following steps.

- Step 1 CSC Support:** To define a Service Policy with Carrier Supporting Carrier (CSC), choose the CSC Support check box from the MPLS Policy Editor - Routing Information.
- When CSC Support is checked, the CSC functionality is enabled to the MPLS VPN service. Provisioning CSC is explained in [Chapter 11, “Provisioning Carrier Supporting Carrier.”](#)
- Step 2 Give Only Default Routes to CE:** Specify whether this service policy should give only default routes to the CE when provisioning with static routes.
- When you enable the **Give only default routes to CE** option with static route provisioning on the PE-CE link, ISC creates a default route on the CE that points to the PE. The VRF static route to the CE site is redistributed into BGP to other sites in the VPN.
- When you choose this option, the default route (0.0.0.0/32) is automatically configured; the site contains no Internet feed or any other requirement for a default route. When the site encounters a packet that does not route locally, it can send the packet to the VPN.
- If you choose this option, ISC configures **the default-info originate** command on the PE router under the running protocol (for RIP, OSPF, or EIGRP). For Static, ISC configures an **ip route 0.0.0.0 0.0.0.0 <out-going interface name>** command on the CE router.
- Step 3 Redistribute Connected:** (BGP Only) Indicate whether this service policy should redistribute the connected routes to the other CEs in the VPN.
- When you enable the **Redistribute Connected** option, the connected routes (that is, the routes to the directly connected PEs or CEs) are distributed to all the other CEs in that particular VPN. This option is meant for IBGP if the routing protocol between PE-CE is a non-BGP protocol. For example, if the routing protocol is RIP, OSPF, EIGRP, or Static, the option is meant for the router BGP that is configured on the PE for the MPLS core. On the PE router, there is one router BGP process running at all times for MPLS. This option is also for BGP.

**Tip**

You must enable the **Redistribute Connected** option when joining the management VPN and you are also using IP numbered addresses.

- Step 4 Default Information Originate:** (BGP only) When you enable this option, ISC issues a **default-information-originate** command under the iBGP address family for the currently specified VRF.

The **Default Information Originate** option is required, especially in the hub and spoke topology because each spoke must be able to communicate with every other spoke (by injecting a default route in the hub PE to the spoke PEs).

- Step 5** When finished defining static routing for this service policy, click **Next**.

The MPLS Policy VRF and VPN Membership dialog box appears. To proceed, see [Defining VRF and VPN Information](#), page 5-29.

RIP Protocol Chosen

The Routing Information Protocol (RIP) is a distance-vector protocol that uses hop count as its metric. RIP is an interior gateway protocol (IGP), which means that it performs routing within a single autonomous system. RIP sends routing-update messages at regular intervals and when the network topology changes. When a router receives a routing update that includes changes to an entry, it updates its routing table to reflect the new route. The metric value for the path is increased by one, and the sender is specified as the next hop.

RIP routers maintain only the best route to a destination—that is, the route with the lowest possible metric value. After updating its routing table, the router immediately begins transmitting routing updates to inform other network routers of the change. These updates are sent independently of the regularly scheduled updates that RIP routers transmit.

To specify RIP as the routing protocol for the service policy, perform the following steps.

- Step 1** Choose **RIP** from the Routing Protocol drop-down list.

The RIP Routing Protocol dialog box appears, as shown in [Figure 5-5](#).

Figure 5-5 RIP Selected as the Routing Protocol

Attribute	Value	Editable
PE-CE Ipv4 Routing Information		
Routing Protocol	RIP	<input checked="" type="checkbox"/>
CsC Support:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Give Only Default Routes to CE:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Redistribute Static (BGP only):	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Redistribute Connected (BGP only):	<input type="checkbox"/>	<input checked="" type="checkbox"/>
RIP Metrics (BGP only):	(1-16)	<input checked="" type="checkbox"/>
Redistributed Protocols on PE:	Edit	<input checked="" type="checkbox"/>
Redistributed Protocols on CE:	Edit	<input checked="" type="checkbox"/>

- Step 2 CSC Support:** To define a Service Policy with Carrier Supporting Carrier (CSC), choose the CSC Support check box from the MPLS Policy Editor - Routing Information.

When CSC Support is checked, the CSC functionality is enabled to the MPLS VPN service. Provisioning CSC is explained in [Chapter 11, “Provisioning Carrier Supporting Carrier.”](#)

- Step 3 Give Only Default Routes to CE:** Specify whether you want to give only the default routes to the CE.

When an internetwork is designed hierarchically, default routes are a useful tool to limit the need to propagate routing information. Access-level networks, such as branch offices, typically have only one connection to headquarters. Instead of advertising all of an organization's network prefixes to a branch office, configure a default route. If a destination prefix is not in a branch office's routing table, forward the packet over the default route. The Cisco IP routing table displays the default route at the top of the routing table as the "Gateway of Last Resort." RIP automatically redistributes the 0.0.0.0 0.0.0.0 route.

If you choose this option, ISC configures **the default-info originate** command on the PE router under the running protocol (for RIP, OSPF, or EIGRP). For Static, ISC configures an **ip route 0.0.0.0 0.0.0.0 <out-going interface name>** command on the CE router.

When you enable the **Give Only Default Routes to CE** option for RIP, ISC creates a default RIP route on the PE; the default RIP route points to the PE and is sent to the CE. The provisioning request gives you the option of redistributing any other routing protocols in the customer network into the CE RIP routing protocol. The RIP routes on the PE to the CE site are redistributed into BGP to other VPN sites.

When you choose this option for RIP routing, the PE instructs the CE to send any traffic it cannot route any other way to the PE. Do *not* use this option if the CE site needs a default route for any reason, such as having a separate Internet feed.

- Step 4 Redistribute Static:** (BGP and RIP) Specify whether you want to redistribute static routes into the core BGP network.

When you enable the **Redistribute Static** option for RIP, the software imports the static routes into the core network (running BGP) and to the CE (running RIP).

- Step 5 Redistribute Connected:** (BGP only) Specify whether you want to redistribute the connected routes to the CEs in the VPN.

When you enable the **Redistribute Connected** option for BGP, the software imports the connected routes (that is, the routes to the directly connected PEs or CEs) to all the other CEs in that particular VPN.

When you enable the Redistribute Connected option, the connected routes (that is, the routes to the directly connected PEs or CEs) are distributed to all the other CEs in that particular VPN. This option is meant for IBGP if the routing protocol between PE-CE is a non-BGP protocol. For example, if the routing protocol is RIP, OSPF, EIGRP, or Static, the option is meant for the router `bgp` that is configured on the PE for the MPLS core. On the PE router, there is one router `bgp` process running at all times for MPLS. This option is also for BGP.

- Step 6 RIP Metrics:** (BGP only) Enter the appropriate RIP metric value. The valid metric values are **1** through **16**.

The metrics used by RIP are hop counts. The hop count for all directly connected interfaces is **1**. If an adjacent router advertises a route to another network with a hop count of 1, then the metric for that network is 2, since the source router must send a packet to that router to get to the destination network.

As each router sends its routing tables to its neighbors, a route can be determined to each network within the AS. If there are multiple paths within the AS from a router to a network, the router selects the path with the smallest hop count and ignores the other paths.

- Step 7 Redistributed Protocols on PE:** Specify whether you want to redistribute the routing protocols into the PE.

Redistribution allows routing information discovered through another routing protocol to be distributed in the update messages of the current routing protocol. With redistribution, you can reach all the points of your IP internetwork. When a RIP router receives routing information from another protocol, it updates all of its RIP neighbors with the new routing information already discovered by the protocol it imports redistribution information from.

To specify the protocols that RIP needs to import routing information to the PE:

- a. From the **Redistribute Protocols on PE** option, click **Edit**.

The PE Redistributed Protocol dialog box appears.

- b. Click **Add**.

The PE Redistributed Protocols dialog box appears.

- c. From the Protocol Type drop-down list, choose the protocol you want to import into the PE.

You can choose one of the following: **Static**, **OSPF**, or **EIGRP**.

- Redistribute Static. When you choose **Static** routes for redistribution into RIP, ISC imports the static routes into the PE that is running RIP.

There are no parameters or metrics required for redistributing Static routes into the PE.

- Redistribute OSPF (Open Shortest Path First). When you choose the **OSPF** protocol for redistribution into RIP, ISC imports the OSPF routes into the PE that is running RIP.

Parameter: OSPF process number

Metric: Any numeral from 1 to 16

- Redistribute EIGRP (Enhanced IGRP). When you choose the **EIGRP** protocol for redistribution into RIP, ISC imports the EIGRP routes into the PE that is running RIP.

Parameter: EIGRP autonomous system (AS) number

Metric: Any numeral from 1 to 16

- d. Choose the protocol you want to redistribute into RIP on the PE.

- e. Enter the appropriate parameter for the protocol selected.

- f. Click **Add**.

- g. Repeat these steps for any additional protocols you want to redistribute into RIP on the PE, then click **OK**.

Step 8 Redistribute Protocols on CE: Specify whether you want to redistribute the routing protocols into the CE.

To specify the protocols that RIP needs to import routing information to the CE:

- a. From the **Redistribute Protocols on CE** option, click **Edit**.

The CE Redistributed Protocol dialog box appears.

- b. Click **Add**.

The CE Redistributed Protocols dialog box appears.

- c. From the Protocol Type drop-down list, choose the protocol you want to import into the CE.

You can choose one of the following protocols: **Static**, **BGP**, **Connected (routes)**, **IGRP**, **OSPF**, **EIGRP**, or **IS-IS**.

- Redistribute Static. When you choose **Static** routes for redistribution into RIP, ISC imports the static routes into the CE that is running RIP.

There are no parameters required for redistributing Static routes into the CE.

- Redistribute BGP (Border Gateway Protocol). When you choose the **BGP** protocol for redistribution into RIP, ISC imports the BGP routes into the CE that is running RIP.

Parameter: BGP autonomous system (AS) number

- Redistribute Connected routes. When you choose the **Connected** routes for redistribution into RIP, ISC imports all the routes to the interfaces connected to the current router. Use the **Connected** option when you want to advertise a network, but you don't want to send routing updates into that network. Note that redistributing connected routes indiscriminately redistributes all connected routes into the routing domain.

Parameter: No parameter required

- Redistribute IGRP (Interior Gateway Routing Protocol). When you choose the **IGRP** (Interior Gateway Routing) protocol for redistribution into RIP, ISC imports the IGRP routes into the CE that is running RIP.

Parameter: IGRP autonomous system (AS) number

- Redistribute EIGRP (Enhanced IGRP). When you choose the **EIGRP** protocol for redistribution into RIP, ISC imports the EIGRP routes into the PE that is running RIP.

Parameter: EIGRP autonomous system (AS) number

- Redistribute OSPF (Open Shortest Path First). When you choose the **OSPF** protocol for redistribution into RIP, ISC imports the OSPF routes into the CE that is running RIP.

Parameter: OSPF process number

- Redistribute IS-IS (Intermediate System-to-Intermediate System). When you choose the **IS-IS** protocol for redistribution into RIP, ISC imports the IS-IS routes into the CE that is running RIP.

Parameter: IS-IS tag number

- d. Choose the protocol you want to redistribute into RIP on the CE.
- e. Enter the appropriate parameter for the selected protocol.
- f. Click **Add**.
- g. Repeat these steps for any additional protocols you want to redistribute into RIP on the CE, then click **OK**.

Step 9 When you are satisfied with the RIP protocol settings for this service policy, click **Next**.

The MPLS Policy VRF and VPN Membership dialog box appears. To proceed, see [Defining VRF and VPN Information, page 5-29](#).



Note

If a PE link is initially configured to use the RIP routing protocol and subsequently modified to use another routing protocol (or static routing), ISC does not remove all of the RIP CLI commands associated with the interface from the PE configuration file. Specifically, ISC does not remove the address family subcommands under the RIP command unless the VRF associated with the service request is removed. This is because ISC configures the RIP protocol using a network class (that is, network a.0.0.0) based under address-family. Later, if the routing protocol is changed, ISC does not remove any other services under the same network.

BGP Protocol Chosen

BGP (Border Gateway Protocol) operates over TCP (Transmission Control Protocol), using port 179. By using TCP, BGP is assured of reliable transport, so the BGP protocol itself lacks any form of error detection or correction (TCP performs these functions). BGP can operate between peers that are separated by several intermediate hops, even when the peers are not necessarily running the BGP protocol.

BGP operates in one of two modes: Internal BGP (iBGP) or External BGP (EBGP). The protocol uses the same packet formats and data structures in either case. IBGP is used between BGP speakers within a single autonomous system, while EBGP operates over inter-AS links.

To specify BGP as the routing protocol for the service policy, perform the following steps.

Step 1 Choose **BGP** from the Routing Protocol drop-down list.

The BGP Routing Protocol dialog box appears, as shown in [Figure 5-6](#).

Figure 5-6 BGP Selected as the Routing Protocol

Attribute	Value	Editable
PE-CE Ipv4 Routing Information		
Routing Protocol	BGP	<input checked="" type="checkbox"/>
CsC Support:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Redistribute Static (BGP only):	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Redistribute Connected (BGP only):	<input type="checkbox"/>	<input checked="" type="checkbox"/>
CE BGP AS ID:	(1-65535)	<input checked="" type="checkbox"/>
Neighbor Allow-AS in:	(1-10)	<input checked="" type="checkbox"/>
Neighbor AS Override:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Redistributed Protocols on CE:	Edit	<input checked="" type="checkbox"/>
Advertise Interval:	(1-600 Seconds)	<input checked="" type="checkbox"/>
Max Prefix Number:	(1-2147483647)	<input checked="" type="checkbox"/>
Max Prefix Threshold:	(1-100 %)	<input checked="" type="checkbox"/>
Max Prefix Warning Only:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Max Prefix Restart:	(1-65535 Minutes)	<input checked="" type="checkbox"/>

Step 2 CSC Support: To define a Service Policy with Carrier Supporting Carrier (CSC), check the CSC Support check box from the MPLS Policy Editor - Routing Information.

When CSC Support is checked, the CSC functionality is enabled to the MPLS VPN service. Provisioning CSC is explained in [Chapter 11, “Provisioning Carrier Supporting Carrier.”](#)

Step 3 Redistribute Static: (BGP only) Indicate whether you want to redistribute static routes into BGP.

If you are importing static routes into BGP, choose this check box.

Step 4 Redistribute Connected Routes: (BGP only) Indicate whether you want to redistribute the directly connected routes into BGP.

Enabling the **Redistribute Connected** option imports all the routes to the interfaces connected to the current router. Use the **Redistribute Connected** option when you want to advertise a network, but you don't want to send routing updates into that network. Note that redistributing connected routes indiscriminately redistributes all connected routes into the routing domain.

When you enable the **Redistribute Connected** option, the connected routes (that is, the routes to the directly connected PEs or CEs) are distributed to all the other CEs in that particular VPN. This option is meant for IBGP if the routing protocol between PE-CE is a non-BGP protocol. For example, if the

routing protocol is RIP, OSPF, EIGRP, or Static, the option is meant for the router BGP that is configured on the PE for the MPLS core. On the PE router, there is one router BGP process running at all times for MPLS. This option is also for BGP.

Step 5 CE BGP AS ID: Enter the BGP autonomous system (AS) number for the customer's BGP network.

The autonomous number assigned here to the CE must be different from the BGP AS number for the service provider's core network.

Step 6 Neighbor Allow-AS In: If appropriate, enter the **Neighbor Allow-AS-in** value.

When you enter a **Neighbor Allow-AS-in** value, you specify a maximum number of times (up to 10) that the service provider autonomous system (AS) number can occur in the autonomous system path.

Step 7 Neighbor AS Override: If required for this VPN, enable the **Neighbor AS Override** option.

The AS Override feature allows the MPLS VPN service provider to run the BGP routing protocol with a customer even if the customer is using the same AS number at different sites. This feature can be used if the VPN customer uses either a private or public autonomous system number.

When you enable the **Neighbor AS-Override** option, you configure VPN Solutions Center to reuse the same AS number on all the VPN's sites.

Step 8 Specify whether you want to redistribute routing protocols into the CE.

Redistributed Protocols on CE: The redistribution of routes into MP-iBGP is necessary only when the routes are learned through any means other than BGP between the PE and CE routers. This includes connected subnets and static routes. In the case of routes learned via BGP from the CE, redistribution is not required because it's performed automatically.

To specify the protocols that BGP needs to import routing information to the CE:

a. From the **Redistribute Protocols on CE** option, click **Edit**.

The CE Redistributed Protocol dialog box appears.

b. Click **Add**.

The CE Redistributed Protocols dialog box appears.

c. From the Protocol Type drop-down list, choose the protocol you want to import into the CE.

You can choose one of the following protocols: **Static**, **RIP**, **Connected (routes)**, **IGRP**, **OSPF**, **EIGRP**, or **IS-IS**.

- Redistribute Static. When you choose **Static** routes for redistribution into BGP, ISC imports the static routes into the CE that is running BGP.

Parameter: No parameter required

- Redistribute RIP (Routing Information Protocol). When you choose the **RIP** protocol for redistribution into BGP, Cisco ISC imports the RIP routes into the CE that is running BGP.

Parameter: No parameter required

- Redistribute Connected routes. When you choose the **Connected** routes for redistribution into BGP, ISC imports all the routes to the interfaces connected to the current router. Use the **Connected** option when you want to advertise a network, but you do not want to send routing updates into that network. Note that redistributing connected routes indiscriminately redistributes all connected routes into the routing domain.

Parameter: No parameter required

- Redistribute IGRP (Interior Gateway Routing Protocol). When you choose the **IGRP** protocol for redistribution into BGP, IP Solution Center imports the IGRP routes into the CE that is running BGP.

Parameter: IGRP autonomous system (AS) number

- Redistribute EIGRP (Enhanced IGRP). When you choose the **EIGRP** protocol for redistribution into BGP, ISC imports the EIGRP routes into the CE that is running BGP.

Parameter: EIGRP autonomous system (AS) number

- Redistribute OSPF (Open Shortest Path First). When you choose the **OSPF** protocol for redistribution into BGP, ISC imports the OSPF routes into the CE that is running BGP.

Parameter: OSPF process number

- Redistribute IS-IS (Intermediate System-to-Intermediate System). When you choose the **IS-IS** protocol for redistribution into BGP, ISC imports the IS-IS routes into the CE that is running BGP.

Parameter: IS-IS tag number

- Choose the protocol you want to redistribute into BGP on the CE.
- Enter the appropriate parameter for the selected protocol.
- Click **Add**.
- Repeat these steps for any additional protocols you want to redistribute into BGP on the PE, then click **OK**.

Step 9 Advertise Interval: Enter the eBGP advertisement interval.

The value is an integer ranging from 0 to 600, specifying the number of seconds of the advertisement interval. The default setting is 30 seconds for the eBGP peer, if it is not explicitly configured.

Step 10 Max Prefix Number: Enter the maximum number of prefixes that can be received from a neighbor.

The range is 1 to 2147483647. This feature allows a router to bring down a peer when the number of received prefixes from that peer exceeds the limit.

Step 11 Max Prefix Threshold: Enter a value that specifies at what percentage Max Prefix Number is configured.

The range is from 1 to 100 percent, with the default being 75 percent. When this threshold is reached, the router generates a warning message. For example, if the Max Prefix Number is 20 and the Max Prefix Threshold is 60, the router generates warning messages when the number of BGP learned routes from the neighbor exceeds 60 percent of 20, or 12 routes.

Step 12 Max Prefix Warning Only: Check this check box if you want to allow the router to generate a log message when the maximum prefix limit is exceeded, instead of terminating the peering session.

Step 13 Max Prefix Restart: Enter a value, in minutes, specifying when the router will automatically re-establish a peering session that has been brought down because the configured maximum prefix limit has been exceeded.

The range is from 1 to 65535. No intervention from the network operator is required when this feature is enabled. This feature attempts to re-establish a disabled peering session at the configured time interval that is specified. However, the configuration of the restart timer alone cannot change or correct a peer that is sending an excessive number of prefixes. The network operator will need to reconfigure the maximum prefix limit or reduce the number of prefixes that are sent from the peer. A peer that is configured to send too many prefixes can cause instability in the network, where an excessive number of prefixes are rapidly advertised and withdrawn. In this case, the Max Prefix Warning Only attribute can be configured to disable the restart capability, while the network operator corrects the underlying problem.

Step 14 When you are satisfied with the BGP protocol settings for this service policy, click **Next**.

The MPLS Policy VRF and VPN Membership dialog box appears. To proceed, see [Defining VRF and VPN Information](#), page 5-29.

OSPF Protocol Chosen

The MPLS VPN backbone is not a genuine OSPF area 0 backbone. No adjacencies are formed between PE routers—only between PEs and CEs. MP-iBGP is used between PEs, and all OSPF routes are translated into VPN IPv4 routes. Thus, redistributing routes into BGP does not cause these routes to become external OSPF routes when advertised to other member sites of the same VPN.

To specify OSPF as the routing protocol for the service policy, perform the following steps.

Step 1 Choose **OSPF** from the Routing Protocol drop-down list.

The OSPF Routing Protocol dialog box appears, as shown in [Figure 5-7](#).

Figure 5-7 OSPF Selected as the Routing Protocol

Attribute	Value	Editable
PE-CE Ipv4 Routing Information		
Routing Protocol	OSPF	<input checked="" type="checkbox"/>
CsC Support:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Metric to Redistribute OSPF into iBGP:	(0-4294967295)	<input checked="" type="checkbox"/>
Give Only Default Routes to CE:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Redistribute Static (BGP only):	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Redistribute Connected (BGP only):	<input type="checkbox"/>	<input checked="" type="checkbox"/>
OSPF Process ID on PE:	(1-65535)	<input checked="" type="checkbox"/>
OSPF Process ID on CE:	(1-65535)	<input checked="" type="checkbox"/>
OSPF Area Number or IP Address:	(0 - 4294967295 or a.b.c.d)	<input checked="" type="checkbox"/>
Redistributed Protocols on PE:	<input type="button" value="Edit"/>	<input checked="" type="checkbox"/>
Redistributed Protocols on CE:	<input type="button" value="Edit"/>	<input checked="" type="checkbox"/>

Step 2 **CSC Support:** To define a Service Policy with Carrier Supporting Carrier (CSC), choose the CSC Support check box from the MPLS Policy Editor - Routing Information.

When CSC Support is checked, the CSC functionality is enabled to the MPLS VPN service. Provisioning CSC is explained in [Chapter 11, “Provisioning Carrier Supporting Carrier.”](#)

Step 3 **Give Only Default Routes to CE:** Specify whether you want to give only the default routes to the CE.

When you enable the **Give only default routes to CE** option, you indicate whether the site needs *full routing* or *default routing*. Full routing is when the site must know specifically which other routes are present in the VPN. Default routing is when it is sufficient to send all packets that are not specifically for your site to the VPN.

If you choose this option, ISC configures the **default-info originate** command on the PE router under the running protocol RIP or EIGRP and the **default-info originate always** command on the PE router under the running protocol OSPF for Static and configures an **ip route 0.0.0.0 0.0.0.0 <out-going interface name>** command on the CE router.

Step 4 **Redistribute Static:** (BGP only) Indicate whether you want to redistribute static routes into OSPF.

If you are importing static routes into OSPF, check this check box.

Step 5 Redistribute Connected Routes: (BGP only) Indicate whether you want to redistribute the directly connected routes into OSPF.

Enabling the **Redistribute Connected** option imports all the routes to the interfaces connected to the current router. Use the **Redistribute Connected** option when you want to advertise a network, but you don't want to send routing updates into that network. Note that redistributing connected routes indiscriminately redistributes all connected routes into the routing domain.

This option is meant for IBGP if the routing protocol between PE-CE is a non-BGP protocol. For example, if the routing protocol is RIP, OSPF, EIGRP, or Static, the option is meant for the router bgp that is configured on the PE for the MPLS core. On the PE router, there is one router bgp process running at all times for MPLS. This option is also for BGP.

Step 6 OSPF Process ID on PE: Enter the OSPF process ID for the PE.

The OSPF process ID is a unique value assigned for each OSPF routing process within a single router—this process ID is internal to the PE only.

Step 7 OSPF Process ID on CE: Enter the OSPF process ID for the CE.

The OSPF process ID is a unique value assigned for each OSPF routing process within a single router—this process ID is internal to the CE only. You can enter this number either as any decimal number from 1 to 65535, or a number in dotted decimal notation.

Step 8 OSPF Process Area Number: Enter the OSPF process area number.

You can enter the OSPF area number for the PE either as any decimal number in the range specified, or a number in dotted decimal notation.

Step 9 Redistributed Protocols on PE: If necessary, specify the redistributed protocols into the PE.



Note

Restricting the amount of redistribution can be important in an OSPF environment. Whenever a route is redistributed into OSPF, it is done so as an external OSPF route. The OSPF protocol floods external routes across the OSPF domain, which increases the protocol's overhead and the CPU load on all the routers participating in the OSPF domain.

To specify the protocols that OSPF needs to import to the PE, follow these steps.

a. From the **Redistribute Protocols on PE** option, click **Edit**.

The PE Redistributed Protocol dialog box appears.

b. Click **Add**.

The PE Redistributed Protocols dialog box appears.

c. From the Protocol Type drop-down list, choose the protocol you want to import into the PE.

You can choose one of the following: **Static**, **EIGRP**, or **RIP**.

- Redistribute Static. When you choose **Static** routes for redistribution into OSPF, ISC imports the static routes into the PE that is running OSPF.

There are no parameters or metrics required for redistributing Static routes into the PE.

- Redistribute EIGRP (Enhanced IGRP). When you choose the **EIGRP** protocol for redistribution into OSPF, ISC imports the EIGRP routes into the PE that is running OSPF.

Parameter: EIGRP autonomous system (AS) number

Metric: Any numeral from 1 to 16777214

- Redistribute RIP. When you choose the **RIP** protocol for redistribution into OSPF, ISC imports the RIP routes into the PE that is running OSPF.

Parameter: No parameter required.

Metric: Any numeral from 1 to 16777214.

- Choose the protocol you want to redistribute into OSPF on the PE.
- Enter the appropriate parameter for the protocol selected.
- Click **Add**.
- Repeat these steps for any additional protocols you want to redistribute into OSPF on the PE, then click **OK**.

Step 10 Specify whether you want to redistribute the routing protocols into the CE.

Redistribute Protocols on CE: To specify the protocols that OSPF needs to import routing information to the CE, follow these steps.

- From the **Redistribute Protocols on CE** option, click **Edit**.

The CE Redistributed Protocol dialog box appears.

- Click **Add**.

The CE Redistributed Protocols dialog box appears.

- From the Protocol Type drop-down list, choose the protocol you want to import into the CE.

You can choose one of the following protocols: **Static**, **RIP**, **BGP**, **Connected (routes)**, **IGRP**, **EIGRP**, or **IS-IS**.

- Redistribute Static. When you choose **Static** routes for redistribution into OSPF, ISC imports the static routes into the CE that is running OSPF.

There are no parameters required for redistributing Static routes into the CE.

- Redistribute RIP. When you choose the **RIP** protocol for redistribution into OSPF, ISC imports the RIP routes into the CE that is running OSPF.

Parameter: No parameter required

- Redistribute BGP (Border Gateway Protocol). When you choose the **BGP** protocol for redistribution into OSPF, ISC imports the BGP routes into the CE that is running OSPF.

Parameter: BGP autonomous system (AS) number

- Redistribute Connected routes. When you choose the **Connected** routes for redistribution into OSPF, ISC imports all the routes to the interfaces connected to the current router. Use the **Connected** option when you want to advertise a network, but you don't want to send routing updates into that network. Note that redistributing connected routes indiscriminately redistributes all connected routes into the routing domain.

Parameter: No parameter required

- Redistribute IGRP (Interior Gateway Routing Protocol). When you choose the **IGRP** (Interior Gateway Routing) protocol for redistribution into OSPF, IP Solution Center imports the IGRP routes into the CE that is running OSPF.

Parameter: IGRP autonomous system (AS) number

- Redistribute EIGRP (Enhanced IGRP). When you choose the **EIGRP** protocol for redistribution into OSPF, ISC imports the EIGRP routes into the CE that is running OSPF.

Parameter: EIGRP autonomous system (AS) number

- Redistribute IS-IS (Intermediate System-to-Intermediate System). When you choose the **IS-IS** protocol for redistribution into OSPF, ISC imports the IS-IS routes into the CE that is running OSPF.

Parameter: IS-IS tag number

- d. Choose the protocol you want to redistribute into OSPF on the CE.
- e. Enter the appropriate parameter for the selected protocol.
- f. Click **Add**.
- g. Repeat these steps for any additional protocols you want to redistribute into OSPF on the CE, then click **OK**.

Step 11 When you are satisfied with the OSPF protocol settings for this service policy, click **Next**.

The MPLS Policy VRF and VPN Membership dialog box appears. To proceed, see [Defining VRF and VPN Information, page 5-29](#).

EIGRP Protocol Chosen

Enhanced IGRP (EIGRP) is a hybrid routing protocol that discovers a network like a distance vector protocol (namely IGRP), but maintains a topological database for rapid reconvergence. EIGRP supports variable length subnet masks and discontinuous subnets. When configured for IP, it automatically redistributes routes with IGRP processes defined in the same autonomous system. By default, EIGRP autosummarizes subnets at the classful network boundaries.

EIGRP performs the same metric accumulation as IGRP. However, if you examine the metric calculation between IGRP and EIGRP, you will see that the EIGRP value is much greater. If you divide the EIGRP metric by 256, you get the same IGRP metric value.

EIGRP allows all routers involved in a topology change to synchronize at the same time. Routers that are not affected by topology changes are not involved in the recomputation. The result is very fast convergence time.

To specify EIGRP as the routing protocol for the service policy, perform the following steps.

Step 1 Choose **EIGRP** from the Routing Protocol drop-down list.

The EIGRP Routing Protocol dialog box appears, as shown in [Figure 5-8](#).

Figure 5-8 EIGRP Selected as the Routing Protocol

MPLS Policy Editor - Ipv4 Routing Information

Attribute	Value	Editable
PE-CE Ipv4 Routing Information		
Routing Protocol	EIGRP	<input checked="" type="checkbox"/>
CsC Support:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Redistribute Static (BGP only):	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Redistribute Connected (BGP only):	<input type="checkbox"/>	<input checked="" type="checkbox"/>
EIGRP AS ID on PE:	(1-65535)	<input checked="" type="checkbox"/>
EIGRP AS ID on CE:	(1-65535)	<input checked="" type="checkbox"/>
Bandwidth Metric:	(1-4294967295)	<input checked="" type="checkbox"/>
Delay Metric:	(1-4294967295)	<input checked="" type="checkbox"/>
Reliability Metric:	(0-255)	<input checked="" type="checkbox"/>
Loading Metric:	(1-255)	<input checked="" type="checkbox"/>
MTU Metric:	(1-4294967295)	<input checked="" type="checkbox"/>
Redistributed Protocols on PE:	<input type="button" value="Edit"/>	<input checked="" type="checkbox"/>
Redistributed Protocols on CE:	<input type="button" value="Edit"/>	<input checked="" type="checkbox"/>

211633

Step 2 CSC Support: To define a Service Policy with Carrier Supporting Carrier (CSC), choose the CSC Support check box from the MPLS Policy Editor - Routing Information.

When CSC Support is checked, the CSC functionality is enabled to the MPLS VPN service. Provisioning CSC is explained in [Chapter 11, “Provisioning Carrier Supporting Carrier.”](#)

Step 3 Redistribute Static: (BGP only) If appropriate, enable the **Redistribute Static (BGP only)** option.

When you enable the Redistribute Static option for BGP, the software imports the static routes into the core network (running BGP).

Step 4 Redistribute Connected: (BGP only) If appropriate, enable the **Redistribute Connected (BGP only)** option.

When you enable the **Redistribute Connected** option, the connected routes (that is, the routes to the directly connected PEs or CEs) are distributed to all the other CEs in that particular VPN. This option is meant for IBGP if the routing protocol between PE-CE is a non-BGP protocol. For example, if the routing protocol is RIP, OSPF, EIGRP, or Static, the option is meant for the router BGP that is configured on the PE for the MPLS core. On the PE router, there is one router PCP process running at all times for MPLS. This option is also for BGP.



Note

Redistributing connected routes can be problematic because all the connected routes are redistributed indiscriminately into a specified routing domain. If you do not want all connected routes to be redistributed, use a *distribute-list out* statement to identify the specific connected routes that should be redistributed.

Step 5 EIGRP AS ID on PE: Enter the EIGRP autonomous system ID on the PE.

This is a unique 16-bit number.

Step 6 EIGRP AS ID on CE: Enter the EIGRP autonomous system ID on the CE.

This is a unique 16-bit number.

Step 7 Enter the values for the EIGRP metrics as described below.

EIGRP Metrics

EIGRP uses metrics in the same way as IGRP. Each route in the route table has an associated metric. EIGRP uses a composite metric much like IGRP, except that it is modified by a multiplier of 256. Bandwidth, Delay, Load, Reliability, and MTU are the submetrics. Like IGRP, EIGRP chooses a route based primarily on bandwidth and delay, or the composite metric with the lowest numerical value. When EIGRP calculates this metric for a route, it calls it the feasible distance to the route. EIGRP calculates a feasible distance to all routes in the network.

Bandwidth Metric: Bandwidth is expressed in units of Kilobits. It must be statically configured to accurately represent the interfaces that EIGRP is running on. For example, the default bandwidth of a 56-kbps interface and a T1 interface is 1,544 kbps.

Delay Metric: Delay is expressed in microseconds. It, too, must be statically configured to accurately represent the interface that EIGRP is running on. The delay on an interface can be adjusted with the **delay time_in_microseconds** interface subcommand.

Reliability Metric: Reliability is a dynamic number in the range of 1 to 255, where 255 is a 100 percent reliable link and 1 is an unreliable link.

Loading Metric: Load is the number in the range of 1 to 255 that shows the output load of an interface. This value is dynamic and can be viewed using the **show interfaces** command. A value of 1 indicates a minimally loaded link, whereas 255 indicates a link loaded 100 percent.

MTU Metric: The maximum transmission unit (MTU) is the recorded smallest MTU value in the path, usually 1500.



Note

Whenever you are influencing routing decisions in IGRP or EIGRP, use the Delay metric over Bandwidth. Changing bandwidth can affect other routing protocols, such as OSPF. Changing delay affects only IGRP and EIGRP.

Step 8 **Redistributed Protocols on PE:** If necessary, specify the redistributed protocols on the PE.

When configured for IP, it automatically redistributes routes with IGRP processes defined in the same autonomous system. By default, EIGRP autosummarizes subnets at the classful network boundaries.

To specify the protocols that EIGRP needs to import to the PE:

- a. From the **Redistribute Protocols on PE** option, click **Edit**.

The PE Redistributed Protocol dialog box appears.

- b. Click **Add**.

The PE Redistributed Protocols dialog box appears.

- c. From the Protocol Type drop-down list, choose the protocol you want to import into the PE.

You can choose one of the following: **Static**, **RIP**, or **OSPF**.

- **Redistribute Static.** When you choose **Static** routes for redistribution into EIGRP, ISC imports the static routes into the PE that is running OSPF.

There are no parameters or metrics required for redistributing Static routes into the PE.

- **Redistribute RIP.** When you choose the **RIP** protocol for redistribution into EIGRP, ISC imports the RIP routes into the PE that is running EIGRP.

Parameter: No parameter required

Metric: Any numeral from 1 to 16777214

- Redistribute OSPF (Open Shortest Path First). When you choose the **OSPF** protocol for redistribution into EIGRP, ISC imports the OSPF routes into the PE that is running EIGRP.

Parameter: OSPF process number

Metric: Any numeral from 1 to 16

- d. Choose the protocol you want to redistribute into EIGRP on the CE.
- e. Enter the appropriate parameter for the protocol selected.
- f. Click **Add**.
- g. Repeat these steps for any additional protocols you want to redistribute into EIGRP on the PE, then click **OK**.

Step 9 Redistribute Protocols on CE: Specify whether you want to redistribute the routing protocols into the CE.

To specify the protocols that EIGRP needs to import routing information to the CE:

- a. From the **Redistribute Protocols on CE** option, click **Edit**.

The CE Redistributed Protocol dialog box appears.

- b. Click **Add**.

The CE Redistributed Protocols dialog box appears.

- c. From the Protocol Type drop-down list, choose the protocol you want to import into the CE.

You can choose one of the following protocols: **Static**, **BGP**, **Connected (routes)**, **IGRP**, **RIP**, **OSPF**, or **IS-IS**.

- Redistribute Static. When you choose **Static** routes for redistribution into EIGRP, ISC imports the static routes into the CE that is running OSPF.

There are no parameters required for redistributing Static routes into the CE.

- Redistribute BGP (Border Gateway Protocol). When you choose the **BGP** protocol for redistribution into EIGRP, ISC imports the BGP routes into the CE that is running OSPF.

Parameter: BGP autonomous system (AS) number

- Redistribute Connected routes. When you choose the **Connected** routes for redistribution into EIGRP, ISC imports all the routes to the interfaces connected to the current router. Use the **Connected** option when you want to advertise a network, but you don't want to send routing updates into that network. Note that redistributing connected routes indiscriminately redistributes all connected routes into the routing domain.

When you enable the **Redistribute Connected** option, the connected routes (that is, the routes to the directly connected PEs or CEs) are distributed to all the other CEs in that particular VPN. This option is meant for IBGP if the routing protocol between PE-CE is a non-BGP protocol. For example, if the routing protocol is RIP, OSPF, EIGRP, or Static, the option is meant for the router BGP that is configured on the PE for the MPLS core. On the PE router, there is one router BGP process running at all times for MPLS. This option is also for BGP.

Parameter: No parameter required

- Redistribute IGRP (Interior Gateway Routing Protocol). When you choose the **IGRP** (Interior Gateway Routing) protocol for redistribution into EIGRP, IP Solution Center imports the IGRP routes into the CE that is running EIGRP.

Parameter: IGRP autonomous system (AS) number

- Redistribute RIP. When you choose the **RIP** protocol for redistribution into EIGRP, Cisco ISC imports the RIP routes into the CE that is running EIGRP.

Parameter: No parameter required

- Redistribute OSPF (Open Shortest Path First). When you choose the **OSPF** protocol for redistribution into EIGRP, ISC imports the OSPF routes into the CE that is running EIGRP.

Parameter: OSPF process number

- Redistribute IS-IS (Intermediate System-to-Intermediate System). When you choose the **IS-IS** protocol for redistribution into EIGRP, ISC imports the IS-IS routes into the CE that is running EIGRP.

Parameter: IS-IS tag number

- Choose the protocol you want to redistribute into EIGRP on the CE.
- Enter the appropriate parameter for the selected protocol.
- Click **Add**.
- Repeat these steps for any additional protocols you want to redistribute into EIGRP on the CE, then click **OK**.

Step 10 When you are satisfied with the EIGRP protocol settings for this service policy, click **Next**.

The MPLS Policy VRF and VPN Membership dialog box appears. To proceed, see [Defining VRF and VPN Information](#), page 5-29.

None Chosen: Cable Services

When operating a cable link, the link does not run a routing protocol. The **None** option in the service policy routing protocol dialog box is provided to allow for configuring a service over a cable link without having to unnecessarily specify a routing protocol.

If this service policy is for cable services, perform the following steps.

Step 1 Choose **None** from the list of routing protocols.

The following dialog box appears, as shown in [Figure 5-9](#).

Figure 5-9 No Routing Protocol Selected

Attribute	Value	Editable
PE-CE Ipv4 Routing Information		
Routing Protocol	NONE	<input checked="" type="checkbox"/>
CsC Support:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Redistribute Static (BGP only):	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Redistribute Connected (BGP only):	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Step 2 **CSC Support:** To define a Service Policy with Carrier Supporting Carrier (CSC), choose the CSC Support check box from the MPLS Policy Editor - Routing Information.

When CSC Support is checked, the CSC functionality is enabled to the MPLS VPN service. Provisioning CSC is explained in [Chapter 11, “Provisioning Carrier Supporting Carrier.”](#)

Step 3 **Redistribute Static:** If you want to distribute static routes into the provider core network (which runs BGP), check the **Redistribute Static (BGP only)** check box.

- Step 4 Redistribute Connected:** Because there is no routing protocol on the cable link, we recommend that you redistribute the connected routes to all the other CEs in the VPN. To do so, check the **Redistribute Connected (BGP only)** check box.

When you enable the **Redistribute Connected** option, the connected routes (that is, the routes to the directly connected PEs or CEs) are distributed to all the other CEs in that particular VPN. This option is meant for IBGP if the routing protocol between PE-CE is a non-BGP protocol. For example, if the routing protocol is RIP, OSPF, EIGRP, or Static, the option is meant for the router BGP that is configured on the PE for the MPLS core. On the PE router, there is one router BGP process running at all times for MPLS. This option is also for BGP.

- Step 5** When finished specifying the necessary settings, click **Next**.

The MPLS Policy VRF and VPN Membership dialog box appears. To proceed, see [Defining VRF and VPN Information](#), page 5-29.

Defining VRF and VPN Information

When you are finished defining the routing protocol(s) for the service policy, you must then specify the VRF and VPN information for this service policy. To do this, perform the following steps.

- Step 1** The MPLS Policy VRF and VPN Membership dialog box appears, as shown in [Figure 5-10](#).

Figure 5-10 Specifying the VRF Information

Attribute	Value	Editable
VRF Information		
Use VRF Object:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Export Map:	<input type="text"/>	<input checked="" type="checkbox"/>
Import Map:	<input type="text"/>	<input checked="" type="checkbox"/>
Maximum Routes:	<input type="text"/> (1-4294967295)	<input checked="" type="checkbox"/>
Maximum Route Threshold:	80 (1-100)	<input checked="" type="checkbox"/>
VRF Description:	<input type="text"/>	<input checked="" type="checkbox"/>
BGP Multipath Load Sharing:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Allocate New Route Distinguisher:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
VRF And RD Overwrite:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
VPN Selection		
PE VPN Membership:		<input checked="" type="checkbox"/>
Select	Customer	VPN
	Provider	CERC
		Is Hub
		<input type="button" value="Add"/> <input type="button" value="Delete"/>

- Step 2** If you want to set the VRF and VPN attributes via a previously defined VRF object, check the **Use VRF Object** check box.

For more information on this feature, see [Chapter 3, “Independent VRF Management.”](#) That chapter describes how to use independent VRF objects in MPLS VPN service policies and service requests.

If you are not using the VRF object feature, then define the VRF and VPN attributes as described in the following steps.

Step 3 Export Map: If necessary, enter the name of the export route map.

The name of the export route map you enter here must be the name of an existing export route map on the PE.

**Note**

The Cisco IOS supports only one export route map per VRF (therefore, there can be only one export route map per VPN).

When you use the ISC software to define a management VPN, ISC automatically generates an export route map for the management VPN. Because the Cisco IOS supports only one export route map per VRF and that route map is reserved for the management VPN, the Export Map field is not available if the VRF is part of the management VPN.

An export route map does not apply a filter; it can be used to override the default set of route targets associated with a route.

Step 4 Import Map: Enter the name of the import route map.

The name of the import route map you enter here must be the name of an existing import route map on the PE.

**Note**

The Cisco IOS supports only one import route map per VRF—therefore, there can be only one import route map per VPN.

An import route map does apply a filter. Therefore, if you want to exclude a particular route from the VRF on this PE, you can either set an export route map on the sending router to make sure it does not have any route targets that can be imported into the current VRF, or create an import route map on the PE to exclude the route.

Step 5 Maximum Routes: Specify the maximum number of routes that can be imported into the VRF on this PE.

**Note**

ISC will not allow provisioning of another value for Maximum Routes after it is configured with a value. Because a VRF might be used by multiple interfaces (links), after this value is configured for a link, it is recommended that you do not manually change it. ISC generates an error if you try to change the maximum routes value for an existing or new service request using this VRF.

Step 6 Maximum Route Threshold: Specify the threshold value for the number of maximum routes.

When the specified number of maximum routes is exceeded, ISC sends a warning message.

Step 7 VRF Description: Optionally, you can enter a description of the VRF for the current VPN.

Step 8 BGP Multipath Load Sharing: Check this check box to enable BGP multipath load sharing and maximum path configuration.

See [BGP Multipath Load Sharing and Maximum Path Configuration, page 5-32](#), for details on using this option.

Step 9 Allocate New Route Distinguisher: A route distinguisher (RD) is a 64-bit number appended to each IPv4 route that ensures that IP addresses that are unique in the VPN are also unique in the MPLS core. This extended address is also referred to as a VPN-IPv4 address.

When **Allocate New Route Distinguisher** is enabled, create a new VRF if there is no matching VRF configuration on that PE; otherwise, reuse it.

When **Allocate New Route Distinguisher** is disabled, find the first matching VRF configuration across the entire range of PEs, regardless of the PE. If this VRF is found on the PE being configured, reuse it. If it is not found on the PE, create it.



Note The service request might get a VRF that has already been configured on another PE router.

ISC automatically sets the route target (RT) and RD values, but you can assign your own values by checking the VRF and RD check box instead.



Note The **Allocate New Route Distinguisher** option is disabled if you enabled the unique route distinguisher feature when the VPN was created. For information, see [Enabling a Unique Route Distinguisher for a VPN, page 2-28](#).

Step 10 VRF and RD Overwrite: When you enable the **VRF and RD Overwrite** option, this dialog box presents two new fields, as shown in [Figure 5-11](#), that allow you to overwrite the default VRF name and route distinguisher values.



Caution

If not done correctly, changing the default values for the VRF name and the route distinguisher value can alter or disable service requests that are currently running. Please make these changes with caution and only when absolutely necessary.



Note The **VRF and RD Overwrite** option is disabled if you enabled the unique route distinguisher feature when the VPN was created. For information, see [Enabling a Unique Route Distinguisher for a VPN, page 2-28](#).

Figure 5-11 No Routing Protocol Selected

VRF And RD Overwrite	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
VRF Name:	<input type="text"/>	<input checked="" type="checkbox"/>
RD Value:	<input type="text"/>	<input checked="" type="checkbox"/>

a. **VRF Name:** Enter the new VRF name. It is recommended not to use special characters (' ` " < > () [] { } / \ & ^ ! ? ~ * % = , . + |), as this may cause misconfiguration of the VRF name for certain devices.

b. **RD Value:** Enter the new RD value.

Step 11 PE VPN Membership: In the check box, specify the VPN associated with this service policy.

The PE VPN Membership information includes the customer name, VPN name, service provider name, CE routing community name, and whether the CERC type is a hub-and-spoke CERC or a fully meshed CERC.

If the **Is Hub** check box is checked, it indicates that the CERC type is hub-and-spoke.

Using the **Add** and **Delete** buttons, you can add a VPN to this list or delete a VPN from this list.

Step 12 If you would like to enable template and data file support for the policy, click the **Next** button to access the Template Association window, and then refer to [Enabling Template Association for a Policy, page 5-34](#) details on working with templates and data files.

- Step 13** If you are satisfied with the VRF and VPN selections, click **Finish**.
The Policies window appears.

Now that you have defined a service policy for an MPLS PE-to-CE service, the service operator can now use this policy to create and deploy a service request for a PE-CE link. For details, see [Chapter 6, “MPLS VPN Service Requests.”](#)

BGP Multipath Load Sharing and Maximum Path Configuration

ISC supports the configuration of Border Gateway Protocol (BGP) multipath load sharing for external BGP (eBGP), internal BGP (iBGP), and external and internal BGP (eiBGP). As additional support for BGP multipath load sharing, MPLS also allows setting a unique route distinguisher (RD) per provider edge (PE) router for a virtual private network (VPN) and virtual route forwarding (VRF) table. The **BGP Multipath Load Sharing** option allows you to enable or disable BGP multipath load sharing, as shown in [Figure 5-12](#).

Figure 5-12 Multipath Configuration Options of the VRF and VPN Membership Window

BGP Multipath Load Sharing:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
BGP Multipath Action	iBGP	<input type="checkbox"/>
Maximum Paths [*] :	<input type="text"/> (1-32)	<input type="checkbox"/>
Import Paths:	<input type="text"/> (1-32)	<input type="checkbox"/>
Unequal Cost:	<input type="checkbox"/>	<input checked="" type="checkbox"/>

When the **BGP Multipath Load Sharing** check box is checked, additional fields are displayed for the BGP multipath action, maximum paths, import paths, and unequal cost routes.



Note

The additional fields appear dynamically in the GUI based on the **BGP Multipath Action** option you choose.

If there is no existing BGP multipath configuration, specifying multipath load sharing through these fields creates a new multipath BGP configuration for the VRF of the PE. If a BGP multipath configuration already exists, this action overwrites the existing configuration with the new multipath values. A remove option allows you to delete all existing BGP multipath configurations of a particular type for the VRF of the PE. If the **BGP Multipath Load Sharing** check box is unchecked, no BGP multipath actions are taken. See [Removing a Multipath Configuration, page 5-34](#), for how multipath settings defined in a service request can be removed.

The following sections describe each of the multipath scenarios, as determined by the type of BGP multipath selected in the **BGP Multipath Action** drop-down list. The options available in the drop-down list are:

- **eBGP**—Specifies the multipath configuration for eBGP. This is the default selection.
- **iBGP**—Specifies the multipath configuration for iBGP.
- **eiBGP**—Specifies the multipath configuration for both eBGP and iBGP. This option allows you to set a common shared value for maximum paths and import paths for both eBGP and iBGP.

- **eBGP+iBGP**—Specifies the multipath configuration for both eBGP and iBGP. This option allows you to set the maximum paths and import paths separately for both eBGP and iBGP.
- **Remove**—Deletes all existing BGP multipath configurations for the VRF of the PE.

Each of these scenarios is covered below.



Note

When creating service requests, in the MPLS Link Editor - VPN and VRF window, an additional BGP attribute called **Force Modify Shared Multipath Attributes** appears in the GUI when the **BGP Multipath Load Sharing** check box is checked. The purpose of this attribute is to enable forced modification of the shared VRF attributes used by other links. This field is not persisted. This attribute only appears when creating service requests, not when creating policies.

eBGP Multipath

When you select the **eBGP** option, the **Maximum Paths** and **Import Paths** fields appear. Where:

- **Maximum Paths**—Specifies the maximum number of routes to allow in the routing table.
- **Import Paths**—Specifies the number of redundant paths that can be configured as backup multipaths for a VRF.



Note

When setting up an eBGP multipath configuration, you must set a value for either **Maximum Paths** or **Import Paths**. Both fields cannot be blank.

iBGP Multipath

When you select the **iBGP** option, the **Maximum Paths**, **Import Paths**, and **Unequal Cost** fields appear. Where:

- **Maximum Paths**—Specifies the maximum number of routes to allow in the routing table. You must specify a value when setting up an iBGP multipath configuration.
- **Import Paths**—Specifies the number of redundant paths that can be configured as backup multipaths for a VRF.
- **Unequal Cost**—Enables/disables unequal-cost multipath. Unequal-cost multipath allows traffic to be distributed among multiple unequal-cost paths to provide greater overall throughput and reliability.

eiBGP Multipath

When you select the **eiBGP** option, the **Maximum Paths** and **Import Paths** fields appear. Where:

- **Maximum Paths**—Specifies the maximum number of routes to allow in the routing table. You must specify a value when setting up an eiBGP multipath configuration.
- **Import Paths**—Specifies the number of redundant paths that can be configured as backup multipaths for a VRF.

eiBGP+iBGP Multipath

When you select the **eiBGP+iBGP** option, the **Maximum Paths**, **Import Paths**, and **Unequal Cost** fields appear. Where:

- **Maximum Paths**—Specifies the maximum number of routes to allow in the routing table. The number of routes can be specified separately for eBGP and iBGP.
- **Import Paths**—Specifies the number of redundant paths that can be configured as backup multipaths for a VRF. The number of paths can be specified separately for eBGP and iBGP.

- **Unequal Cost**—Enables/disables unequal-cost multipath. Unequal-cost multipath allows traffic to be distributed among multiple unequal-cost paths to provide greater overall throughput and reliability.

**Note**

The support for multipath load sharing requires unique route distinguishers (RDs) for each PE router for a VPN (VRF). This is to prevent the same RDs from being allocated to different customers. This allows the use of the same RD for the same VRF. That is, all sites in the PE can have the same unique RD. The unique RD feature is optional. It is enabled at both a global VPN level or a service request level. To enable the unique RD per PE for a VPN, the Create VPN window contains a new **Enable Unique Route Distinguisher** field. For more information on using this feature, see [Enabling a Unique Route Distinguisher for a VPN, page 2-28](#).

Removing a Multipath Configuration

A multipath configuration can be removed by selecting the **Remove** option in drop-down list of the BGP Multipath Action attribute. The Remove option removes the multipath configuration for the VRF on the PE, if it is previously configured.

If a service request is saved with a multipath configuration and the configuration has to be removed, you should use the Remove option.

**Note**

A multipath configuration cannot be removed by simply unchecking the BGP Multipath Load Sharing check box. It must be removed by setting the BGP Multipath Action attribute to Remove, and then saving the service request. You should uncheck the BGP Multipath Load Sharing check box only after removing the multipath configuration.

Enabling Template Association for a Policy

The ISC template feature gives you a means to download free-format CLIs to devices configured for links within an MPLS service request. If you enable templates, you can use templates and data files to download commands that are not currently supported by ISC.

- Step 1** To enable template association for the policy, click the **Next** button in MPLS Policy Editor - VRF and VPN Membership window.

The Template Association window appears. In this window, you can enable template support and, optionally, associate templates and data files with the policy. For instructions about associating templates with policies and how to use the features in this window, see [Appendix B, “Working with Templates and Data Files.”](#) For additional information about using templates in ISC, also see the *Cisco IP Solution Center Infrastructure Reference, 5.1*.

- Step 2** When you have completed setting up templates and data files for the policy per the instructions in the appendix, click **Finish** in the Template Association window to close it.

The Policies window appears.

Now that you have defined a service policy for an MPLS PE-to-CE service, the service operator can now use this policy to create and deploy a service request for a PE-CE link. For details, see [Chapter 6, “MPLS VPN Service Requests.”](#)

