



MPLS VPN Service Policies

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

- [Service Policy Overview, page 3-1](#)
- [Creating Service Policies, page 3-7](#)
- [Creating an MPLS Service Policy for a PE-to-CE Link, page 3-9](#)

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).

Creating an MPLS VPN in ISC

At its simplest, a virtual private network (VPN) is a collection of sites that share the same routing table. A VPN is also a framework that provides private IP networking over a public infrastructure such as the Internet. In IP Solution Center (ISC), a VPN is a set of customer sites that are configured to communicate through a VPN service. A VPN is defined by a set of administrative policies.

A VPN is a network in which two sites can communicate over the provider's network in a private manner; that is, no site outside the VPN can intercept their packets or inject new packets. The provider network is configured such that only one VPN's packets can be transmitted through that VPN—that is, no data can come in or out of the VPN unless it is specifically configured to allow it. There is a physical connection from the provider edge network to the customer edge network, so authentication in the conventional sense is not required.

To create a new VPN in ISC: MPLS, follow these steps:

-
- Step 1** Log into ISC.
 - Step 2** Choose the **Service Inventory** tab.
 - Step 3** Choose **Inventory and Connection Manager**.

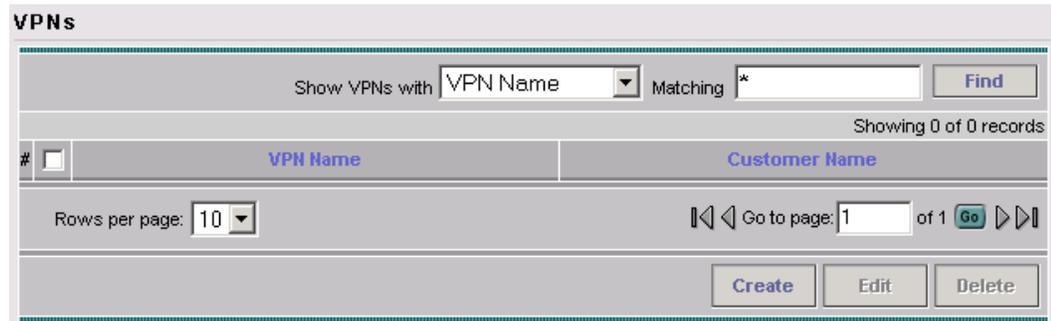
The Inventory and Connection Manager window appears (see [Figure 3-1](#)).

Figure 3-1 Creating an MPLS VPN in ISC



- Step 4** From the Inventory and Connection Manager, choose **VPNs**.
The VPNs dialog box appears (see [Figure 3-2](#)).

Figure 3-2 Viewing Existing VPNs or Creating a New VPN



- Step 5** From the VPNs dialog box, click **Create**.
The Create VPN dialog box appears (see [Figure 3-3](#)).

Figure 3-3 Creating an MPLS VPN in ISC

Create VPN

Name*:

Customer*:

MPLS Attributes

Create Default CE Routing Community:

Enable Multicast:

Data MDT Size:

Data MDT Threshold: (1 - 4294967 bits/sec)

CE Routing Communities:

VPLS Attributes

Enable VPLS:

Service Type:

Topology:

Note: * - Required Field

Step 6 *Name:* Enter the name of the VPN:

Step 7 *Customer:* To select the customer associated with this VPN:

- a. Click **Select**.

The Select Customer dialog box appears (see [Figure 3-4](#)).

Figure 3-4 Selecting a Customer for the VPN

Select Customer - Microsoft Internet Explorer

Show Customers with Customer Name matching

Showing 1 - 3 of 3 records

#	Customer Name
1.	<input checked="" type="radio"/> CUST1
2.	<input type="radio"/> Customer_Ford
3.	<input type="radio"/> DiscoveredL2Customer

Rows per page: Go to page: of 1

- b. From the list of customers, select the appropriate customer, then click **Select**.

Step 8 *Create Default CE Routing Community:* To create a default CE routing community, select the **Create Default CE Routing Community** check box and select a provider.

Step 9 *Enable Multicast:*

An IP address that starts with the binary prefix *1110* is identified as a *multicast group address*. There can be more than one sender and receiver at any time for a given multicast group address. The senders send their data by setting the group address as the destination IP address. It is the responsibility of the network to deliver this data to all the receivers in the network who are listening to that group address.

**Note**

Before you can create a VPN with multicast enabled, you must define one or more multicast resource pools. See [Create a Multicast Pool, page 2-34](#), for further information.

- a. To enable multicast VPN routing, check the **Enable Multicast** check box.
ISC enables two additional fields required to configure multicast routing (see [Figure 3-5](#)).

Figure 3-5 Selecting a Customer for the VPN

MPLS Attributes	
Create Default CE Routing Community:	<input checked="" type="checkbox"/> PROV1
Enable Multicast:	<input checked="" type="checkbox"/>
Data MDT Size:	0
Data MDT Threshold:	0 (1 - 4294967 bits/sec)
CE Routing Communities:	<div style="border: 1px solid gray; width: 100%; height: 40px;"></div> <div style="text-align: right;"> <input type="button" value="Select"/> <input type="button" value="Remove"/> </div>

To implement multicast routing, ISC employs the concept of a *multicast domain* (MD), which is a set of VRFs associated with interfaces that can send multicast traffic to each other. A VRF contains VPN routing and forwarding information for unicast. A *multicast VRF* contains multicast routing and forwarding information and supports multicast routing.

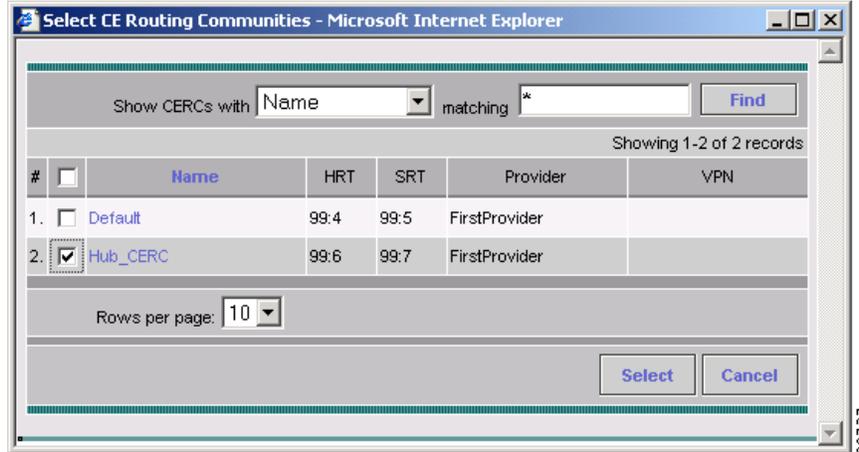
- b. *Data MDT Size*: From the drop-down list, select the data MDT size.
MDT refers to a *multicast distribution tree* (MDT). The MDT defined here carries multicast traffic from customer sites associated with the multicast domain.
- c. *Data MDT Threshold*: Enter the bandwidth threshold for the data multicast distribution tree.
The *data MDT* contains a range of multicast group addresses and a bandwidth threshold. Thus, whenever a CE behind a multicast-VRF exceeds that bandwidth threshold while sending multicast traffic, the PE sets up a new data MDT for the multicast traffic from that source. The PE informs the other PEs about this data MDT and, if they have receivers for the corresponding group, the other PEs join this data MDT.

Step 10 *CE Routing Communities*: If you do not choose to enable the default CERC, you can select a customized CERC that you have already created in ISC (see [Defining CE Routing Communities, page 3-6](#)):

- a. From the CE Routing Communities pane, click **Select**.

The Select CE Routing Communities dialog box appears (see [Figure 3-6](#)).

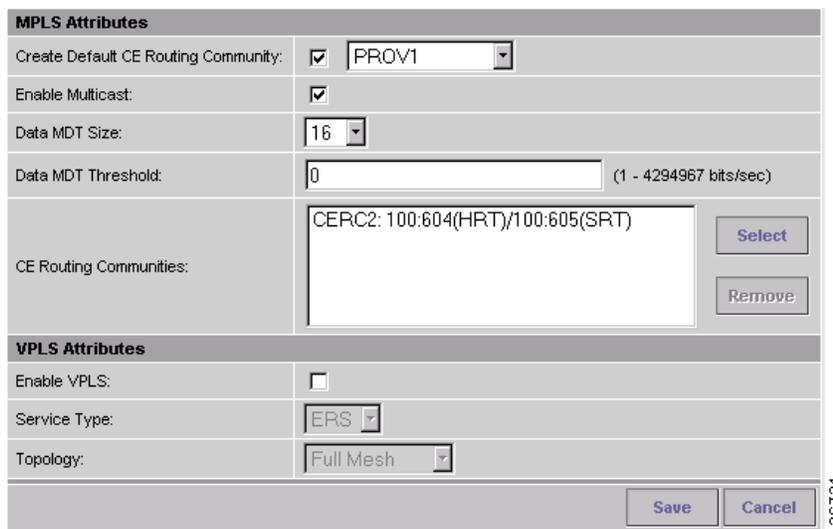
Figure 3-6 Selecting a CERC



- b. Click the check box for the CERC you want used for this VPN, then click **Select**.

You return to the Create VPN dialog box, where the new CERC selection appears, along with its *hub route target (HRT)* and spoke route target (SRT) values (see Figure 3-7).

Figure 3-7 New CERC Selected



- Step 11 Enable VPLS** (optional) Select this check box to enable VPLS.
- Step 12 Service Type** (optional) Select the VPLS service type from the drop-down menu: **ERS** (Ethernet Relay Service) or **EWS** (Ethernet Wire Service).
- Step 13 Topology** (optional) Select the VPLS topology from the drop-down menu: **Full Mesh** (each CE will have direct connections to every other CE) or **Hub and Spoke** (only the Hub CE has connection to each Spoke CE and the Spoke CEs do not have direct connection to each other).
- Step 14** When satisfied with the settings for this VPN, click **Save**.

You have successfully created a VPN, as shown in the Status display in the lower left corner of the VPNs dialog box.

Defining CE Routing Communities

When you create a VPN, the ISC software creates one default CE routing community (CERC) for you. But if your network topology and configuration require customized CERC definitions, you can define CERCs customized for your network.



Tip

Customized CERCs should be defined only in consultation with the VPN network administrator.

To build complex topologies, it is necessary to break down the required connectivity between CEs into groups, where each group is either fully meshed, or has a hub-and-spoke pattern. A CE can be in more than one group at a time, so long as each group has one of the two basic configuration patterns.

Each subgroup in the VPN needs its own CERC. Any CE that is only in one group just joins the corresponding CERC (as a spoke if necessary). If a CE is in more than one group, then you can use the Advanced Setup choice during provisioning to add the CE to all the relevant groups in one service request. Given this information, ISC does the rest, assigning route target values and VRF tables to arrange the precise connectivity the customer requires.

To define a new CERC:

- Step 1** Select the **Service Inventory** tab.
- Step 2** Choose **Inventory and Connection Manager**.
The Inventory and Connection Manager window appears.
- Step 3** Choose **CE Routing Communities**.
The CE Routing Communities dialog box appears (see [Figure 3-8](#)).

Figure 3-8 CE Routing Communities Defined for This VPN

#	<input type="checkbox"/>	Name	HRT	SRT	Provider	VPN
1.	<input type="checkbox"/>	Default	99:0	99:1	FirstProvider	AcmeIncVPN
2.	<input type="checkbox"/>	Default	99:2	99:3	FirstProvider	WidgetsIncVPN
3.	<input type="checkbox"/>	Default	99:4	99:5	FirstProvider	

Showing 1-3 of 3 records

Rows per page: 10

Create Edit Delete

116329

- Step 4** From the CE Routing Communities dialog box, click **Create**.
The Create CE Routing Community dialog box appears (see [Figure 3-9](#)).

Figure 3-9 Defining a New CE Routing Community

- Step 5** Complete the CERC fields as required for the VPN:
- c. *Provider:* To specify the service provider associated with this CERC, click **Select**.
The Select Provider dialog box appears.
 - d. Choose the name of the service provider, then click **Select**.
 - e. *Name:* Enter the name of the CERC.
 - f. *CERC Type:* Specify the CERC type: *Hub and Spoke* or *Fully Meshed*.
 - g. *Auto-Pick Route Target Values:* Choose to either let ISC automatically set the route target (RT) values or set the RT values manually.
By default, the **Auto-pick route target values** check box is checked. If you uncheck the check box, you can enter the Route Target values manually.

**Caution**

If you choose to bypass the **Auto-pick route target values** option and set the route target (RT) values manually, note that the RT values cannot be edited after they have been defined in the Cisco IP Solution Center software.

- Step 6** When you have finished entering the information in the Create CE Routing Community dialog box, click **Save**.

Creating Service Policies

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).

The 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 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.

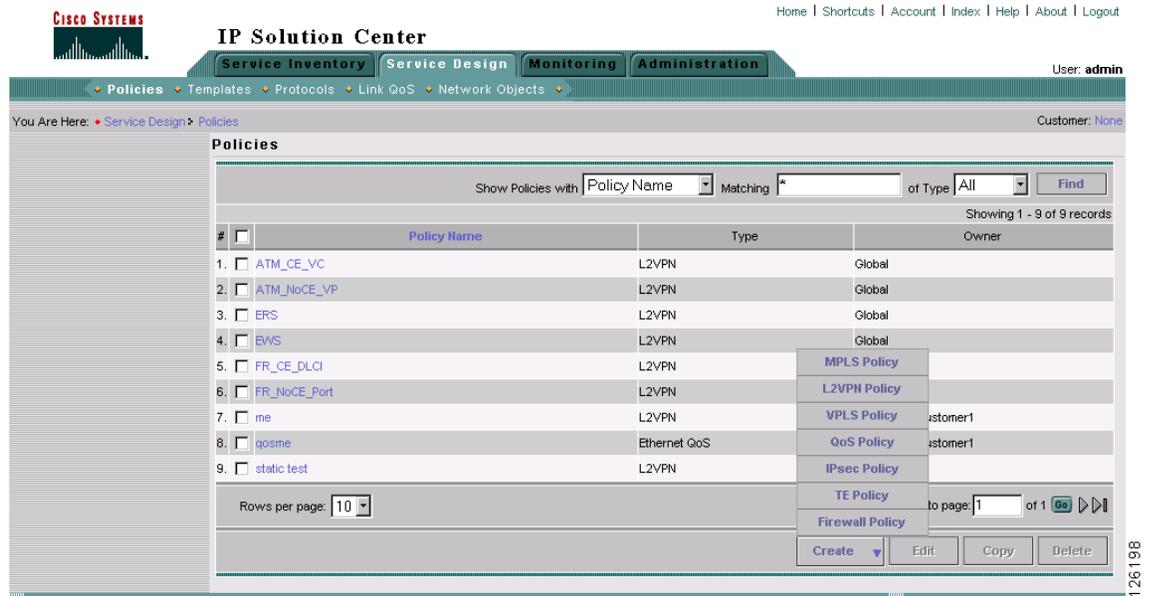
Creating an MPLS Service Policy for a PE-to-CE Link

To create an MPLS service policy for a PE-to-CE link, follow these steps:

- Step 1** Log into ISC.
- Step 2** Select the **Service Design** tab.
- Step 3** Choose **Policies**.

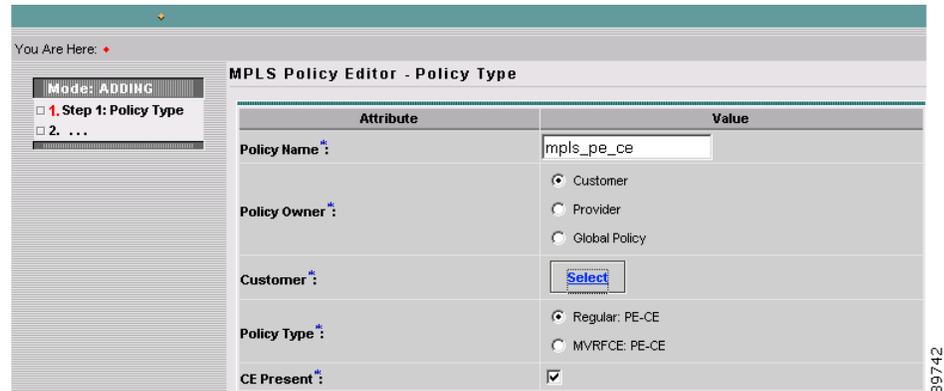
The Policies window appears (see [Figure 3-10](#)).

Figure 3-10 Creating a New Service Policy



- Step 4** From the **Create** drop-down list, choose **MPLS Policy**.
The MPLS Policy Type dialog box appears (see [Figure 3-11](#)).

Figure 3-11 Defining the MPLS Service Policy



-
- Step 1** Enter a **Policy Name** for the MPLS policy.
- Step 2** 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.
- Step 3** 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 4** 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 5** Select 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 select 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 6** Click **Next**.
-

Specifying the PE and CE Interface Parameters

The MPLS Policy Interface dialog box appears (see [Figure 3-12](#)).



Tip

You do not have to select 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 3-12 Specifying the PE and CE Interface Parameters

MPLS Policy Editor - Interface

Attribute	Value	Editable
Reset All Attribute Editable Flags:		<input checked="" type="checkbox"/>
PE Information		
Interface Type:	ANY	
Interface Format:		
Interface Description:		<input checked="" type="checkbox"/>
Shutdown Interface:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Auto-Pick VLAN ID:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Link Speed:	None	<input checked="" type="checkbox"/>
Link Duplex:	None	<input checked="" type="checkbox"/>
ETTH Support:	<input type="checkbox"/>	<input checked="" type="checkbox"/>
CE Information		
Interface Type:	ANY	
Interface Format:		
Interface Description:		<input checked="" type="checkbox"/>

- Step 2 of 5 -

< Back Next > Finish Cancel

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

PE Interface Information

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

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 enable this check box, the specified PE interface is configured in a shut down state.

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

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

Table 3-1 shows the protocol encapsulations available for each of the supported interface types.

Table 3-1 *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)
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)

Table 3-1 Interface Types and Their Corresponding Encapsulations (continued)

Interface Type	Encapsulations (continued)
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:* Enable this check box to have ISC automatically pick the VLAN ID.

Step 7 *Link Speed:* Enter a Link Speed (optional) of 10, 100, 1000, or auto.

Step 8 *Link Duplex:* Enter a Line Duplex (optional) of full, half, or auto.

Step 9 *ETTH Support:* Enable this check box to configure Ethernet-To-The-Home (ETTH). See [Ethernet-To-The-Home, page 10-9](#) for an explanation of ETTH.

CE Interface Information

Step 10 *Interface Type:* From the drop-down list, select the interface type for the CE.

Step 11 *Interface Format:* Optionally, you can specify the slot number and port number for the CE interface.

Step 12 *Interface Description:* Optionally, you can enter a description of the CE interface.

Step 13 *Encapsulation:* Select the encapsulation used for the specified CE interface type.

Step 14 When satisfied with the interface settings, click **Next**.

Specifying the IP Address Scheme

The MPLS Policy Interface Address Selection dialog box appears (see [Figure 3-13](#)). This lets you specify the IP address scheme you want to use for this service policy.

Figure 3-13 Specifying the IP Address Scheme

Attribute	Value	Editable
PE-CE Interface Addresses/Mask		
IP Numbering Scheme:	IP Numbered	<input checked="" type="checkbox"/>
Extra CE 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"/>

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

IP Numbering Scheme

A point-to-point link between two routers can be either a *numbered* IP address or an *unnumbered* IP address. The service provider must determine whether to use numbered or unnumbered IP addresses for the PE-CE link. Defining the link to use unnumbered addresses can save precious IP addresses because many interfaces can borrow the same IP address.

You can choose among two options: **IP Numbered** or **IP Unnumbered**.

- **IP Numbered**

If you select **IP numbered** and choose to not use automatically assigned IP addresses, you can enter the IP addresses for the PE interface and CE interface in the fields provided. Entering the IP addresses in these fields forces the MPLS VPN software to use the indicated addresses.

If you choose **IP numbered** and also enable 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 IP addresses from a /30 subnet point-to-point IP address pool.

- **IP Unnumbered**

IP 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. When you choose **IP unnumbered**, ISC: MPLS creates a static route for the PE-CE link.

When you choose **IP unnumbered**, ISC: MPLS automatically creates a loopback interface (unless a loopback interface already exists with the correct attributes). For related information, see [Using an Existing Loopback Interface Number, page 3-15](#).

If you select **IP unnumbered** and choose to not use automatically assigned IP addresses, you can enter the IP addresses for the PE interface and CE interface in the fields provided. Entering the IP addresses in these fields forces the ISC: MPLS software to use the indicated addresses.

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 **IP unnumbered** and also enable 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 **IP numbered** and also enable 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.

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.

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

Using an 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 select loopback3 for the unnumbered interface when using automatically assigned addresses. Loopback1 will be chosen instead. The only way to select 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.

Specifying the Routing Protocol for a Service

You can now specify the routing protocol information for this service policy (see [Figure 3-14](#)).

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 3-17](#)).
- **RIP.** Routing Information Protocol (see [RIP Protocol Chosen, page 3-18](#)).
- **BGP.** Border Gateway Protocol (see [BGP Protocol Chosen, page 3-22](#)).
- **OSPF.** Open Shortest Path First (see [OSPF Protocol Chosen, page 3-25](#)).
- **EIGRP.** Enhanced Interior Gateway Routing Protocol (see [EIGRP Protocol Chosen, page 3-29](#)).
- **None.** Specifies parameters for cable services (see [None Chosen: Cable Services, page 3-34](#)).

To specify a routing protocol for the PE-CE link:

-
- Step 1** Select the appropriate protocol from the Routing Protocol drop-down list.
- When you select 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 the Service Policy VRF and VPN Information, page 3-35](#).
-

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 select 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 9, "Provisioning Carrier Supporting Carrier."](#)

Giving Only 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 select 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 may 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 select **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)** (see [Figure 3-14](#)).



Note

Two other options (**AdvertisedRoutes** and **Default Routes - Routes to reach other sites**) are available when you create the service request. See [Static Routing Protocols, page 4-12](#).

Figure 3-14 Specifying the Static Routing Protocol

MPLS Policy Editor - Routing Information		
Attribute	Value	Editable
PE-CE 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"/>

- 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 9, “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 select 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 select 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 the Service Policy VRF and VPN Information](#), page 3-35.

RIP Protocol Chosen

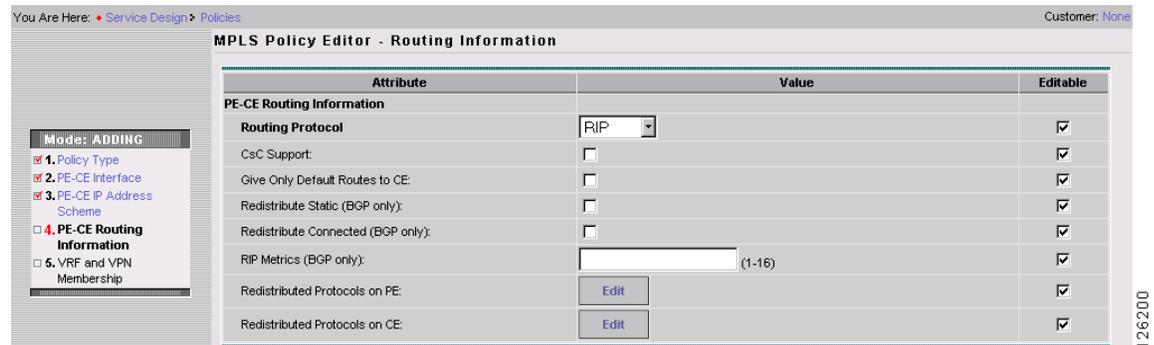
The Routing Information Protocol (RIP) is a distance-vector protocol that uses hop count as its metric. RIP is an interior gateway protocol (IGP), which means that it performs routing within a single autonomous system. RIP sends routing-update messages at regular intervals and when the network topology 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.

Step 1 To specify RIP as the routing protocol for the service policy, choose **RIP** from the Routing Protocol drop-down list.

The RIP Routing Protocol dialog box appears (see [Figure 3-15](#)).

Figure 3-15 RIP Selected as the Routing Protocol



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 9](#), “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 select 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:

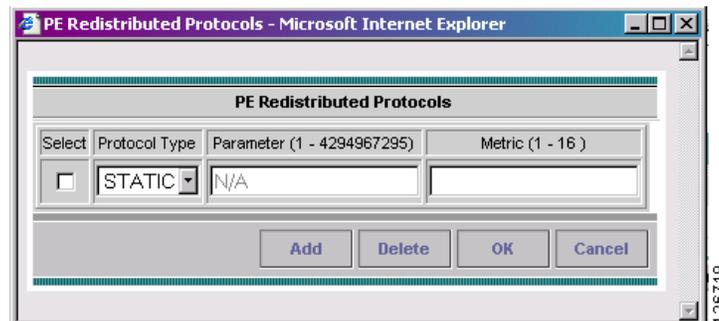
- a. From the **Redistribute Protocols on PE** option, click **Edit**.

The PE Redistributed Protocol dialog box appears.

- b. Click **Add**.

The following dialog box appears (see [Figure 3-16](#)).

Figure 3-16 *Selecting Protocols to Redistribute into the PE*



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

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

- *Redistribute Static*

When you select **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 select 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 select 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*

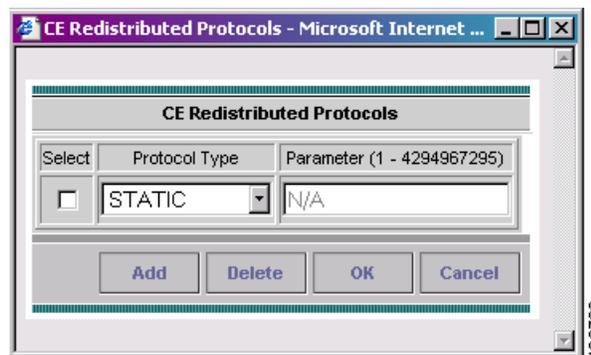
- Select the protocol you want to redistribute into RIP on the PE.
- Enter the appropriate parameter for the protocol selected.
- Click **Add**.
- Repeat these steps for any additional protocols you want to redistribute into RIP on the PE, then click **OK**.

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:

- From the **Redistribute Protocols on CE** option, click **Edit**.
The CE Redistributed Protocol dialog box appears.
- Click **Add**.
The following dialog box appears (see [Figure 3-17](#)).

Figure 3-17 *Selecting Protocols to Redistribute into the CE*



- From the Protocol Type drop-down list, select 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 select **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 select 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 select 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 select the **IGRP** (Interior Gateway Routing) protocol for redistribution into RIP, IP Solution Center imports the IGRP routes into the CE that is running RIP.

Parameter: *IGRP autonomous system (AS) number*

- *Redistribute EIGRP (Enhanced IGRP)*

When you select 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 select 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 select 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*

- Select the protocol you want to redistribute into RIP 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 RIP on the CE, then click **OK**.

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

To complete this service policy, go to [Defining the Service Policy VRF and VPN Information, page 3-35](#).

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.

Step 1 To specify BGP as the routing protocol for the service policy, choose **BGP** from the Routing Protocol drop-down list.

The BGP Routing Protocol dialog box appears (see [Figure 3-18](#)).

Figure 3-18 BGP Selected as the Routing Protocol

Attribute	Value	Editable
PE-CE Routing Information		
Routing Protocol	BGP	<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"/>
CE BGP AS ID:	<input type="text"/> (1-65535)	<input checked="" type="checkbox"/>
Neighbor Allow-AS in:	<input type="text"/> (1-10)	<input checked="" type="checkbox"/>
Neighbor AS Override:	<input type="checkbox"/>	<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 9](#), “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 select 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 4 *Redistribute Static (BGP Only)*: Indicate whether you want to redistribute static routes into BGP. If you are importing static routes into BGP, select this check box.

Step 5 *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 6 *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 7 *Neighbor Allow-AS In:* If appropriate, enter the **Neighbor Allow-AS-in** value.

When you enter a **Neighbor AllowAS-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 8 *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 9 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:

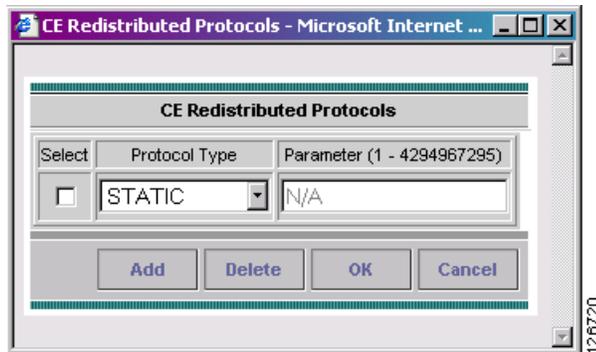
a. From the **Redistribute Protocols on CE** option, click **Edit**.

The CE Redistributed Protocol dialog box appears.

b. Click **Add**.

The following dialog box appears (see [Figure 3-19](#)).

Figure 3-19 *Selecting Protocols to Redistribute into the CE*



c. From the Protocol Type drop-down list, select 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 select **Static** routes for redistribution into BGP, ISC imports the static routes into the CE that is running BGP.

Parameter: *No parameter required*

- *Redistribute RIP (Border Gateway Protocol)*

When you select 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 select 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 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 select the **IGRP** (Interior Gateway Routing) 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 select 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 select 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 select 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*

- d. Select the protocol you want to redistribute into BGP 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 BGP on the PE, then click **OK**.

Step 10 When you're satisfied with the BGP protocol settings for this service policy, click **Next**.

To complete this service policy, go to [Defining the Service Policy VRF and VPN Information, page 3-35](#).

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.

Step 1 To specify OSPF as the routing protocol for the service policy, choose **OSPF** from the Routing Protocol drop-down list.

The OSPF Routing Protocol dialog box appears (see [Figure 3-20](#)).

Figure 3-20 OSPF Selected as the Routing Protocol

You Are Here: [Service Design](#) > [Policies](#) Customer: None

MPLS Policy Editor - Routing Information

Attribute	Value	Editable
PE-CE Routing Information		
Routing Protocol	OSPF	<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"/>
OSPF Process ID on PE:	<input type="text"/> (1-65535)	<input checked="" type="checkbox"/>
OSPF Process ID on CE:	<input type="text"/> (1-65535)	<input checked="" type="checkbox"/>
OSPF Area Number or IP Address:	<input type="text"/> (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"/>

126202

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 9](#), “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 select 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, select 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:

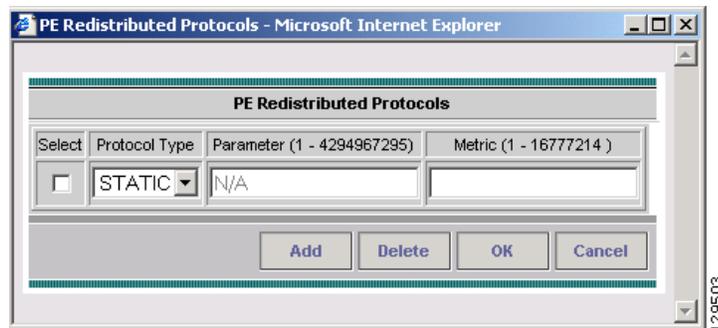
- a. From the **Redistribute Protocols on PE** option, click **Edit**.

The PE Redistributed Protocol dialog box appears.

- b. Click **Add**.

The following dialog box appears (see [Figure 3-21](#)).

Figure 3-21 *Selecting Protocols to Redistribute into the PE*



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

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

- *Redistribute Static*

When you select **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 select 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 select 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

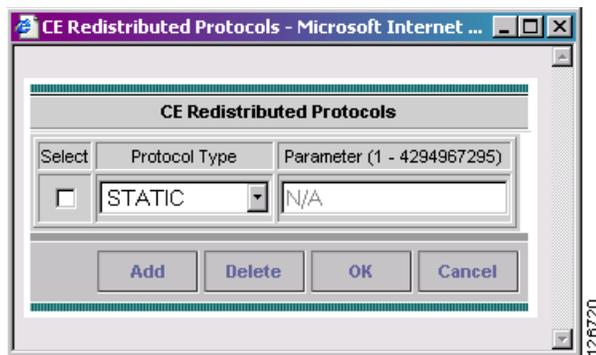
- d. Select the protocol you want to redistribute into OSPF on the PE.
- e. Enter the appropriate parameter for the protocol selected.
- f. Click **Add**.
- g. Repeat these steps for any additional protocols you want to redistribute into OSPF on the PE, then click **OK**.

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:

- a. From the **Redistribute Protocols on CE** option, click **Edit**.
The CE Redistributed Protocol dialog box appears.
- b. Click **Add**.
The following dialog box appears (see [Figure 3-22](#)).

Figure 3-22 Selecting Protocols to Redistribute into the CE



- c. From the Protocol Type drop-down list, select 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 select **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 select 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 select 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 select 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 select 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 select 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 select 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*

- Select the protocol you want to redistribute into OSPF 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 OSPF on the CE, then click **OK**.

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

To complete this service policy, go to [Defining the Service Policy VRF and VPN Information, page 3-35](#).

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 auto-summarizes 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.

Step 1 To specify EIGRP as the routing protocol for the service policy, choose **EIGRP** from the Routing Protocol drop-down list.

The EIGRP Routing Protocol dialog box appears (see [Figure 3-23](#)).

Figure 3-23 EIGRP Selected as the Routing Protocol

Attribute	Value	Editable
PE-CE Routing Information		
Routing Protocol	EIGRP	<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"/>
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"/>

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 9](#), “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 `bgp` 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 auto-summarizes subnets at the classful network boundaries.

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

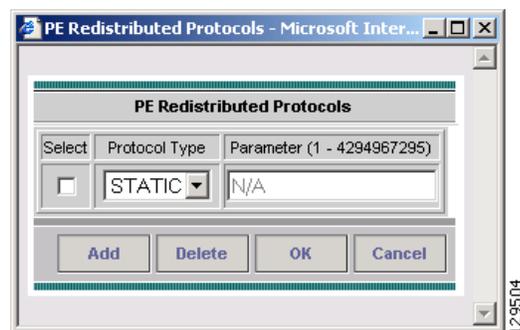
- a. From the **Redistribute Protocols on PE** option, click **Edit**.

The PE Redistributed Protocol dialog box appears.

- b. Click **Add**.

The following dialog box appears (see [Figure 3-24](#)).

Figure 3-24 Selecting Protocols to Redistribute into the PE



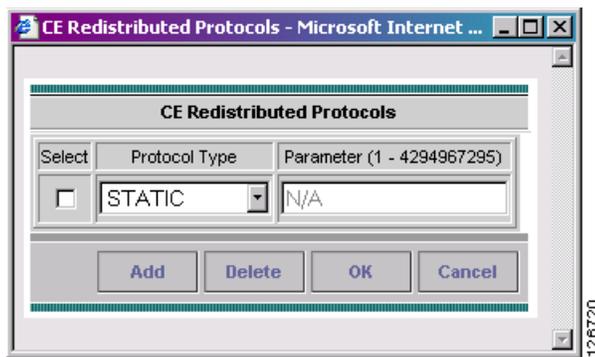
- c. From the Protocol Type drop-down list, select the protocol you want to import into the PE. You can choose one of the following: **Static**, **RIP**, or **OSPF**.
 - *Redistribute Static*
When you select **Static** routes for redistribution into EIGRP, ISC imports the static routes into the PE that is running OSPF.
There are no parameters or metrics required for redistributing Static routes into the PE.
 - *Redistribute RIP*
When you select 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 select 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. Select 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 following dialog box appears (see [Figure 3-25](#)):

Figure 3-25 *Selecting Protocols to Redistribute into the CE*



- c. From the Protocol Type drop-down list, select 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 select **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 select 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 select 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 select 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 select 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 select 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 select 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*

- d. Select the protocol you want to redistribute into EIGRP 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 EIGRP on the CE, then click **OK**.

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

To complete this service policy, go to [Defining the Service Policy VRF and VPN Information, page 3-35](#).

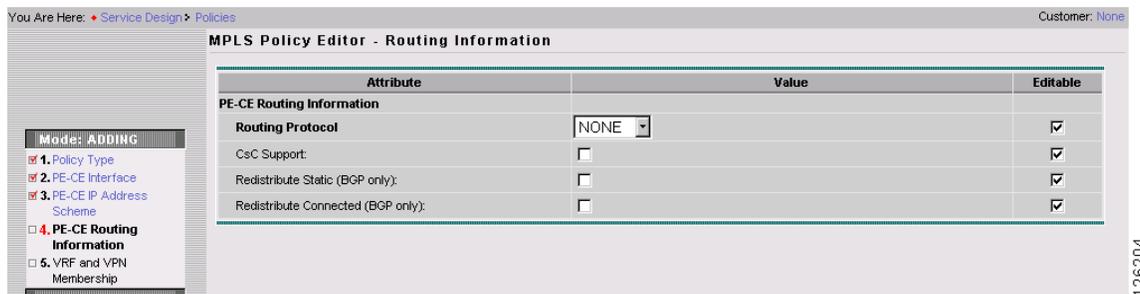
None Chosen: Cable Services

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.

Step 1 If this service policy is for cable services, choose **None** from the list of routing protocols.

The following dialog box appears (see [Figure 3-26](#)):

Figure 3-26 No Routing Protocol Selected



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 9, "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**.

Defining the Service Policy VRF and VPN Information

When you are finished defining the routing protocol(s) for this service policy, you must then specify the VRF information.

The MPLS Policy VRF and VPN Membership dialog box appears (see [Figure 3-27](#)).

Figure 3-27 Specifying the VRF Information

MPLS Policy Editor - VRF and VPN Membership

Mode: ADDING

- 1. Policy Type
- 2. PE-CE Interface
- 3. PE-CE IP Address Scheme
- 4. PE-CE Routing Information
- 5. VRF and VPN Membership

Attribute	Value	Editable			
VRF Information					
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:	<input type="text"/> 80 (1-100)	<input checked="" type="checkbox"/>			
VRF Description:	<input type="text"/>	<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"/>			
Template Association					
Template Enable:	<input checked="" type="checkbox"/>				
VPN Selection					
PE VPN Membership:		<input checked="" type="checkbox"/>			
Select	Customer	VPN	Provider	CERC	Is Hub
<input type="checkbox"/>	AcmeInc	AcmeIncVPN	FirstProvider	Default	<input checked="" type="checkbox"/>

Add Delete

To specify the VRF and VPN information for this service policy:

Step 1 *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 2 *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 3 *Maximum Routes*: Specify the maximum number of routes that can be imported into the VRF on this PE.

Step 4 *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 5 *VRF Description*: Optionally, you can enter a description of the VRF for the current VPN. By default, this field states: *Provisioned by ISC*.

Step 6 *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*.

ISC automatically sets the route target (RT) and RD values, but you can assign your own values by checking the *Allocate new route distinguisher* checkbox.

Step 7 *VRF and RD Overwrite*: When you enable the **VRF and RD Overwrite** option, this dialog box presents two new fields (see [Figure 3-28](#)) 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.

Figure 3-28 *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"/>

89754

a. *VRF Name*: Enter the new VRF name.

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

Step 8 *Template Enable*: This option determines whether the network devices configured for links within an MPLS service request can be associated with MPLS templates.

When you enable the **Template Enable** option, the next dialog box that appears will let the service operator select the templates to be associated with the MPLS link.

Step 9 *PE VPN Membership*: In the **Select** 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 10 When satisfied with the VRF and VPN selections, click **Finish**.

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 4, “MPLS VPN Service Requests.”](#)

