Cisco VPN Solutions Center: MPLS Solution Provisioning Guide

Software Release 2.2

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Text Part Number: 78-14530-01
Customer Order Number: 7814530=

Cisco VPN Solutions Center: MPLS Solution Provisioning Guide

CONTENTS

About This Guide

This guide describes how to use the Cisco VPN Solutions Center: MPLS Solution 2.2 software.

This guide is designed for Service Provider network managers and operators who are responsible for provisioning MPLS VPNs for their customers. The network manager and operators should be familiar with the following topics:

- Basic concepts and terminology used in internetworking
- Multiprotocol Label Switching (MPLS) and virtual private network (VPN) terms and technology
- Network topologies and protocols

This chapter describes how the guide is organized, how to obtain additional documentation, and how to contact Cisco's Technical Assistance Center.

How This Guide Is Organized

The chapters of this guide are as follows:

Chapter 1	Introduction to Cisco MPLS VPN Technology	Provides an overview of the major concepts that structure the Cisco VPN Solutions Center: MPLS Solution, a network management system that defines and monitors virtual private network (VPN) services for service providers.
Chapter 2	Setting Up Devices in the VPN Solutions Center MPLS Environment	Describes the tasks the network administrator must complet before using VPN Solutions Center and provides detailed instructions on setting up the various devices in the network prior to starting the software.
Chapter 3	Starting and Stopping the VPN Solutions Center Software	Provides the instructions for starting and shutting down the software.
Chapter 4	Setting Up Networks and Importing Configurations Into VPN Solutions Center	Provides instructions for the network administrator on setting up the software equivalent of the Service Provider's MPLS network in VPN Solutions Center software. This chapter also describes how to import PE and CE configuration files into VPNSC.

Chapter 5	Creating MPLS VPNs and Administering Service Request Profiles	Focuses on the procedures that a network administrator uses to a) create MPLS VPNs in VPNSC, and b) create and administer Service Request Profiles
Chapter 6	Provisioning MPLS VPN Service Requests	Provides a detailed description of how network operators create and deploy MPLS VPN service requests.
Chapter 7	Monitoring MPLS VPN Performance	Provides a tutorial on collecting VPN accounting data, monitoring service related performance, and configuring traps.
Chapter 8	The VPNSC Management Network	This chapter describes the VPNSC management subnet and how to implement it in the provider network. It also provides the fundamental concepts and considerations, as well as our recommendations, for administering customer edge routers (CEs) in a service provider environment.
Chapter 9	Provisioning MPLS VPN Cable Services	Provides an introduction an overview of the MPLS VPN cable feature and then shows how to provision cable services through VPN Solutions Center.
Chapter 10	Provisioning with the VPN Solutions Center Template Manager	Describes how to use the Template Manager to define standard templates to generate Cisco IOS configurations for common provisioning tasks, such as common IPv4, QoS, and VPN provisioning. The Template Manager provides fast, flexible, and extensible Cisco IOS command generation capability.
Chapter 11	Provisioning Multi-VRF CEs in VPN Solutions Center	This chapter describes how to define a Multi-VRF CE device and how to use the Template Manager to provision the provider-facing and customer-facing interfaces on a Multi-VRF CE.
Chapter 12	Spanning Multiple Autonomous Systems	Provides a technical overview of how routing is accomplished between multiple autonomous systems, then describes how to configure VPNs in the VPN Solutions Center software to communicate between multiple autonomous systems.
Chapter 13	Provisioning MPLS VPNs on the Cisco Series DSL Switch	Provides an overview of IP DSL Switch configuration and a detailed description of the four VPNSC templates for provisioning the DSL switch.
Chapter 14	Repository Administration	Provides information on converting a 1.x Repository to the 2.x format as well as converting a 2.0 Repository to the 2.x format. It also discussed how to back up and restore the VPNSC Repository, and how to export and import the Repository.
Chapter 15	VPNSC: MPLS Solution Troubleshooting Guide	Provides troubleshooting information in a question and answer format to help you resolve problems you may encounter when deploying MPLS VPNs.

Appendixes

Appendix A	VPN Solutions Center Configuration File Examples	Provides examples of configuration files, including a cable network example and an example showing the migration process for numbered access list entries.
Appendix B	Cisco VPNSC: MPLS Solution Command Reference	Provides information on the set of Cisco IOS commands specifically for configuring MPLS VPNs.

Glossary

Index

Related Information

Cisco documentation and additional literature are available in a CD-ROM package, which ships with your product. The Documentation CD-ROM, a member of the Cisco Connection Family, is updated monthly. Therefore, it might be more up to date than printed documentation. To order additional copies of the Documentation CD-ROM, contact your local sales representative or call customer service. The CD-ROM package is available as a single package or as an annual subscription.

Sending Feedback on the Documentation

If you are reading Cisco product documentation on the World Wide Web, you can submit comments electronically:

1. Click **Feedback** in the toolbar.
2. Select **Documentation**.
3. Complete the form, click **Submit**, and send it to Cisco Systems.

You can e-mail your comments to **bug-doc@cisco.com**.

To submit your comments by mail, for your convenience many documents contain a response card behind the front cover. Otherwise, you can mail your comments to the following address:

Cisco Systems, Inc.
Document Resource Connection
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Related Documentation

This *Provisioning Guide* and the following documents comprise the VPNSC: MPLS Solution 2.2 documentation set:

- *Release Notes for Cisco VPN Solutions Center 2.2*
- *Cisco VPN Solutions Center Installation Guide*
- *Cisco VPN Solutions Center: MPLS Solution User Reference*
- *Cisco VPN Solutions Center: MPLS Solution API Programmer Guide*

Obtaining Documentation

These sections explain how to obtain documentation from Cisco Systems.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at this URL:

http://www.cisco.com

Translated documentation is available at this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which is shipped with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual subscription.

Ordering Documentation

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace:

  http://www.cisco.com/cgi-bin/order/order_root.pl

- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:

  http://www.cisco.com/go/subscription

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, U.S.A.) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit comments electronically on Cisco.com. In the Cisco Documentation home page, click the **Fax** or **Email** option in the "Leave Feedback" section at the bottom of the page.

You can e-mail your comments to bug-doc@cisco.com.

You can submit your comments by mail by using the response card behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Document Resource Connection
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain online documentation, troubleshooting tips, and sample configurations from online tools by using the Cisco Technical Assistance Center (TAC) Web Site. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC Web Site.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com is a highly integrated Internet application and a powerful, easy-to-use tool that provides a broad range of features and services to help you with these tasks:

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

If you want to obtain customized information and service, you can self-register on Cisco.com. To access Cisco.com, go to this URL:

http://www.cisco.com

Technical Assistance Center

The Cisco Technical Assistance Center (TAC) is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two levels of support are available: the Cisco TAC Web Site and the Cisco TAC Escalation Center.

Cisco TAC inquiries are categorized according to the urgency of the issue:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration.
- Priority level 3 (P3)—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- Priority level 2 (P2)—Your production network is severely degraded, affecting significant aspects of business operations. No workaround is available.
- Priority level 1 (P1)—Your production network is down, and a critical impact to business operations will occur if service is not restored quickly. No workaround is available.

The Cisco TAC resource that you choose is based on the priority of the problem and the conditions of service contracts, when applicable.

Cisco TAC Web Site

You can use the Cisco TAC Web Site to resolve P3 and P4 issues yourself, saving both cost and time. The site provides around-the-clock access to online tools, knowledge bases, and software. To access the Cisco TAC Web Site, go to this URL:

http://www.cisco.com/tac

All customers, partners, and resellers who have a valid Cisco service contract have complete access to the technical support resources on the Cisco TAC Web Site. The Cisco TAC Web Site requires a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to this URL to register:

http://www.cisco.com/register/

If you are a Cisco.com registered user, and you cannot resolve your technical issues by using the Cisco TAC Web Site, you can open a case online by using the TAC Case Open tool at this URL:

http://www.cisco.com/tac/caseopen

If you have Internet access, we recommend that you open P3 and P4 cases through the Cisco TAC Web Site.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses priority level 1 or priority level 2 issues. These classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer automatically opens a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to this URL:

http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Before calling, please check with your network operations center to determine the level of Cisco support services to which your company is entitled: for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). When you call the center, please have available your service agreement number and your product serial number.

Introduction to Cisco MPLS VPN Technology

Technology Overview

The Cisco VPN Solutions Center: MPLS Solution, a modular suite of network and service management applications, is a network management system that defines and monitors virtual private network (VPN) services for service providers. VPN Solutions Center allows service providers to provision and manage intranet and extranet VPNs.

The product provides the aspect of operations management that addresses flow-through provisioning, service auditing, and Service Level Agreement (SLA) measurement of IP-based MPLS VPN environments. Multiprotocol Label Switching (MPLS) is an emerging industry standard upon which tag switching is based.

VPN Solutions Center is a scalable, provider-focused VPN technology that allows service providers to plan, provision, and manage for IP VPN services according to a customer's service level agreement. This product complements Cisco's MPLS-based VPN solutions by simplifying the provisioning, service assurance, and billing processes, thereby reducing the cost of deploying and operating VPN services.

VPN Solutions Center does not contain a billing application, but the product enables billing by providing the usage data on services that a billing engine can process.

VPN Solutions Center focuses on provisioning, auditing, and monitoring the links between the customer's routers through the provider's network. VPNSC deals only with the provider's edge routers and the customer's edge routers.

Figure 1-1 VPN Solutions Center: MPLS Solution in the Service Provider Network

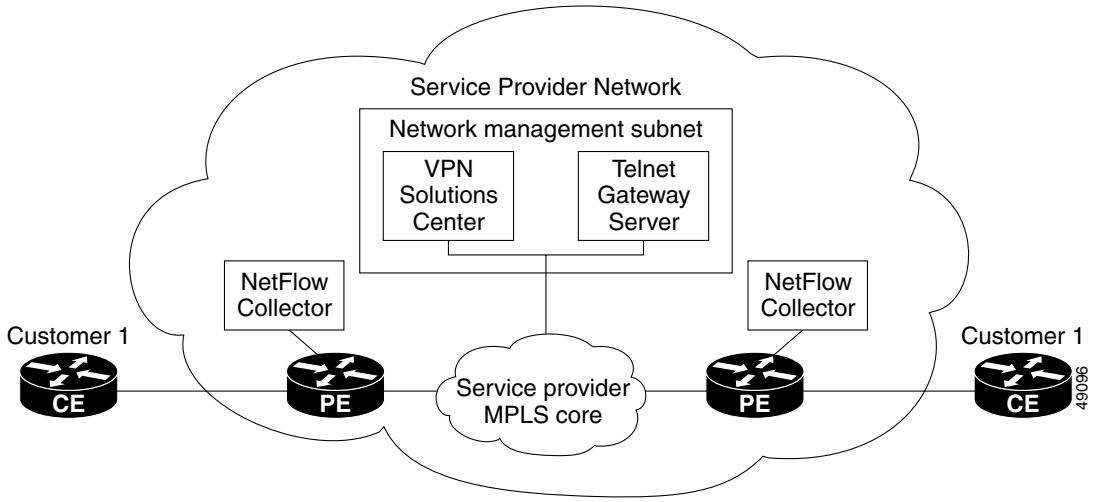

As shown in Figure 1-1, a customer edge router (CE) is connected to a provider edge router (PE) in such a way that the customer's traffic is encapsulated and transparently sent to other CEs, thus creating a virtual private network. CEs advertise routes to the VPN for all the devices in their site. The VPN Solutions Center provisioning engine accesses the configuration files on both the CE and PE to compute the necessary changes to those files that are required to support the service on the PE-CE link.

Using the VPN Solutions Center (VPNSC) software, service providers can do the following:

- Provision IP-based MPLS VPN services
- Generate audit reports for service requests
- Perform data collection to measure SLA performance
- Evaluate service usage for each VPN

An MPLS VPN consists of a set of sites that are interconnected by means of an MPLS provider core network. At each site, there are one or more CEs, which attach to one or more PEs. PEs use the Border Gateway Protocol-Multiprotocol (MP-BGP) to dynamically communicate with each other.

It is not required that the set of IPv4 addresses used in any two VPNs be mutually exclusive because the PEs translate IPv4 addresses into IPv4 VPN entities by using MP-BGP with extended community attributes.

The set of IP addresses used in a VPN, however, must be exclusive of the set of addresses used in the provider network. Every CE must be able to address the PEs to which it is directly attached. Thus, the IP addresses of the PEs must not be duplicated in any VPN.

The Customer's and Provider's View of the Network

From the customer's point of view, they see their internal routers communicating with their customer edge routers (CEs) from one site to another through a VPN managed by the service provider (see Figure 1-2).

Figure 1-2 The Customer's View of the Network

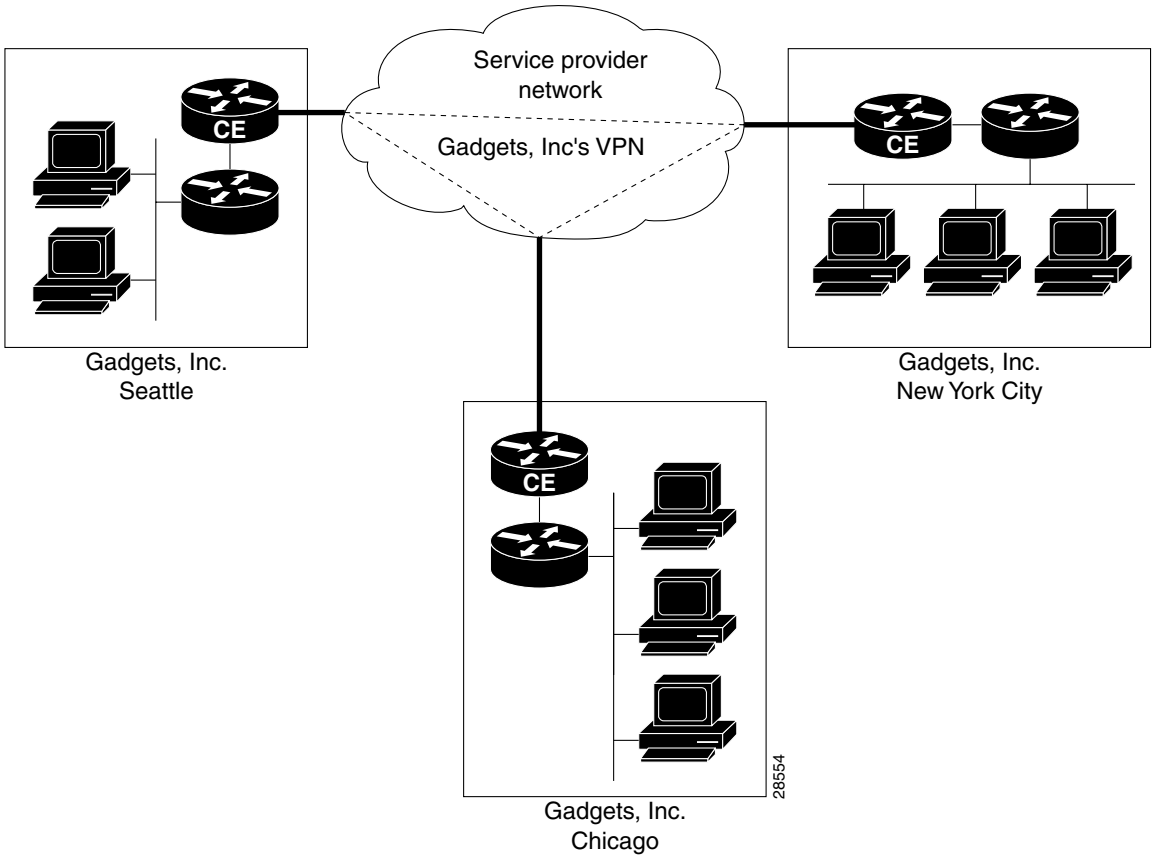

This simple view of the customer's network is the advantage of employing VPNs: the customer experiences direct communication to their sites as though they had their own private network, even though their traffic is traversing a public network infrastructure and they are sharing that infrastructure with other businesses.

The service provider's view of the network is naturally very different, as shown in Figure 1-3. This illustration shows two different customers, with each customer having a single VPN. A customer can, however, have multiple VPNs.

Figure 1-3 Service Provider's View of the Network

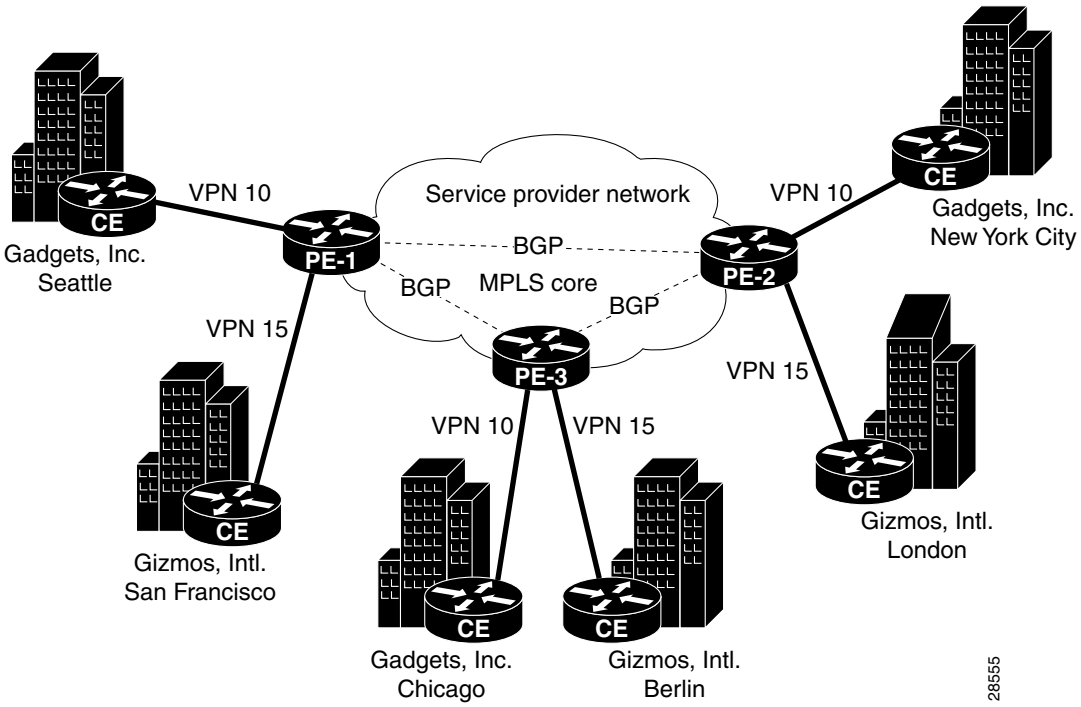

About PEs

At the edge of the provider network are provider edge routers (PEs). Within the provider network are other provider routers as needed (often designated as P routers) that communicate with each other and the PEs via the Border Gateway Protocol-Multiprotocol (MP-BGP). Note that in this model, the service provider need only provision the links between the PEs and CEs.

PEs maintain separate routing tables called VPN routing and forwarding tables (VRFs). The VRFs contain the routes for directly connected VPN sites only. (For more information about VRFs, see the "VPN Routing and Forwarding Tables (VRFs)" section on page 1-15). PEs exchange VPN-IPv4 updates through MP-iBGP sessions. These updates contain VPN-IPv4 addresses and labels. The PE originating the route is the next hop of the route. PE addresses are referred to as host routes into the core interior gateway protocol.

Benefits

MPLS-based VPNs provide the following benefits:

- A platform for rapid deployment of additional value-added IP services, including intranets, extranets, voice, multimedia, and network commerce
- Privacy and security equal to Layer-2 VPNs by constraining the distribution of a VPN's routes to only those routers that are members of that VPN, and by using MPLS for forwarding
- Seamless integration with customer intranets
- Increased scalability with thousands of sites per VPN and hundreds of thousands of VPNs per service provider

- IP Class of Service (CoS) with support for multiple classes of service within a VPN, as well as priorities among VPNs
- Easy management of VPN membership and rapid deployment of new VPNs
- Scalable any-to-any connectivity for extended intranets and extranets that encompass multiple businesses

About MPLS VPNs

At its simplest, a virtual private network (VPN) is a collection of sites that share the same routing table. A VPN is also a network in which customer connectivity to multiple sites is deployed on a shared infrastructure with the same administrative policies as a private network.The path between two systems in a VPN, and the characteristics of that path, may also be determined (wholly or partially) by policy. Whether a system in a particular VPN is allowed to communicate with systems not in the same VPN is also a matter of policy.

In MPLS VPN, a VPN generally consists of a set of sites that are interconnected by means of an MPLS provider core network, but it is also possible to apply different policies to different systems that are located at the same site. Policies can also be applied to systems that dial in; the chosen policies would be based on the dial-in authentication processes.

A given set of systems can be in one or more VPNs. A VPN can consist of sites (or systems) that are all from the same enterprise (intranet), or from different enterprises (extranet); it may consist of sites (or systems) that all attach to the same service provider backbone, or to different service provider backbones.

Figure 1-4 VPNs Sharing Sites

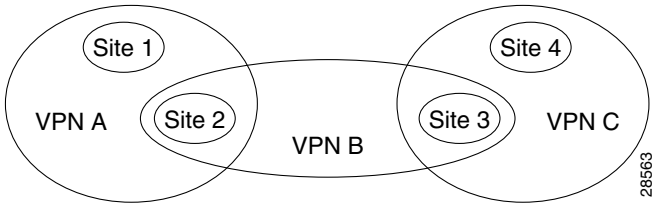

MPLS-based VPNs are created in Layer 3 and are based on the peer model, which makes them more scalable and easier to build and manage than conventional VPNs. In addition, value-added services, such as application and data hosting, network commerce, and telephony services, can easily be targeted and deployed to a particular MPLS VPN because the service provider backbone recognizes each MPLS VPN as a secure, connectionless IP network.

The MPLS VPN model is a true peer VPN model that enforces traffic separations by assigning unique VPN route forwarding tables (VRFs) to each customer's VPN. Thus, users in a specific VPN cannot see traffic outside their VPN. Traffic separation occurs without tunneling or encryption because it is built directly into the network. (For more information on VRFs, see the "VPN Routing and Forwarding Tables (VRFs)" section on page 1-15.)

The service provider's backbone is comprised of the PE and its provider routers. MPLS VPN provides the ability that the routing information about a particular VPN be present *only* in those PE routers that attach to that VPN.

Characteristics of MPLS VPNs

MPLS VPNs have the following characteristics:

- Multiprotocol Border Gateway Protocol-Multiprotocol (MP-BGP) extensions are used to encode customer IPv4 address prefixes into unique VPN-IPv4 Network Layer Reachability Information (NLRI) values.

  NLRI refers to a destination address in MP-BGP, so NLRI is considered "one routing unit." In the context of IPv4 MP-BGP, NLRI refers to a network prefix/prefix length pair that is carried in the BGP4 routing updates.

- Extended MP-BGP community attributes are used to control the distribution of customer routes.
- Each customer route is associated with an MPLS label, which is assigned by the provider edge router that originates the route. The label is then employed to direct data packets to the correct egress customer edge router.

  When a data packet is forwarded across the provider backbone, two labels are used. The first label directs the packet to the appropriate egress PE; the second label indicates how that egress PE should forward the packet.

- Cisco MPLS CoS and QoS mechanisms provide service differentiation among customer data packets.
- The link between the PE and CE routers uses standard IP forwarding.

  The PE associates each CE with a per-site forwarding table that contains only the set of routes available to that CE.

Principal Technologies

There are four principal technologies that make it possible to build MPLS-based VPNs:

- Multiprotocol Border Gateway Protocol (MP-BGP) between PEs carries CE routing information
- Route filtering based on the VPN route target extended MP-BGP community attribute
- MPLS forwarding carries packets between PEs (across the service provider backbone)
- Each PE has multiple VPN routing and forwarding instances (VRFs)

Intranets and Extranets

If all the sites in a VPN are owned by the same enterprise, the VPN is a corporate *intranet*. If the various sites in a VPN are owned by different enterprises, the VPN is an *extranet*. A site can be in more than one VPN. Both intranets and extranets are regarded as VPNs.

While the basic unit of interconnection is the site, the MPLS VPN architecture allows a finer degree of granularity in the control of interconnectivity. For example, at a given site, it may be desirable to allow only certain specified systems to connect to certain other sites. That is, certain systems at a site may be members of an intranet as well as members of one or more extranets, while other systems at the same site may be restricted to being members of the intranet only.

A CE router can be in multiple VPNs, although it can only be in a single site. When a CE router is in multiple VPNs, one of these VPNs is considered its primary VPN. In general, a CE router's primary VPN is the intranet that includes the CE router's site. A PE router may attach to CE routers in any number of different sites, whether those CE routers are in the same or in different VPNs. A CE router may, for robustness, attach to multiple PE routers. A PE router attaches to a particular VPN if it is a router adjacent to a CE router that is in that VPN.

Security Requirements for MPLS VPNs

This section discusses the security requirements for MPLS VPN architectures. This section concentrates on protecting the core network against attacks from the "outside," that is, the Internet and connected VPNs. Protection against attacks from the "inside," that is, when an attacker has logical or physical access to the core network is not discussed here, since any network can be attacked with access from the inside.

Address Space and Routing Separation

Between two non-intersecting VPNs of an MPLS VPN service, it is assumed that the address space between different VPNs is entirely independent. This means, for example, that two non-intersecting VPNs must be able to both use the 10/8 network without any interference. From a routing perspective, this means that each end system in a VPN has a unique address, and all routes to this address point to the same end system. Specifically:

- Any VPN must be able to use the same address space as any other VPN.
- Any VPN must be able to use the same address space as the MPLS core.
- Routing between any two VPNs must be independent.
- Routing between any VPN and the core must be independent.

Address Space Separation

From a security point of view, the basic requirement is to avoid that packets destined to a host a.b.c.d within a given VPN reach a host with the same address in another VPN or the core.

MPLS allows distinct VPNs to use the same address space, which can also be private address space. This is achieved by adding a 64-bit route distinguisher (RD) to each IPv4 route, making VPN-unique addresses also unique in the MPLS core. This "extended" address is also called a *VPN-IPv4 address*. Thus customers of an MPLS service do not need to change current addressing in their networks.

In the case of using routing protocols between CE and PE routers (for static routing this is not an issue), there is one exception—the IP addresses of the PE routers the CE routers are peering with. To be able to communicate to the PE router, routing protocols on the CE routers must configure the address of the peer router in the core. This address must be unique from the CE router's perspective. In an environment where the service provider manages also the CE routers as CPE (customer premises equipment), this can be made invisible to the customer.

Routing Separation

Routing separation between the VPNs can also be achieved. Every PE router maintains a separate Virtual Routing and Forwarding instance (VRF) for each connected VPN. Each VRF on the PE router is populated with routes from one VPN, through statically configured routes or through routing protocols that run between the PE and the CE router. Since every VPN results in a separate VRF, there are no interferences between the VPNs on the PE router.

Across the MPLS core to the other PE routers, this routing separation is maintained by adding unique VPN identifiers in multi-protocol BGP, such as the route distinguisher (RD). VPN routes are exclusively exchanged by MP-BGP across the core, and this BGP information is not redistributed to the core network, but only to the other PE routers, where the information is kept again in VPN-specific VRFs. Thus routing across an MPLS network is separate per VPN.

Given addressing and routing separation across an MPLS core network, MPLS offers in this respect the same security as comparable Layer 2 VPNs, such as ATM or Frame Relay. It is not possible to intrude into other VPNs through the MPLS core, unless this has been configured specifically.

Hiding the MPLS Core Structure

The internal structure of the MPLS core network (PE and Provider router devices) should not be visible to outside networks (either the Internet or any connected VPN). While a breach of this requirement does not lead to a security problem itself, it is generally advantageous when the internal addressing and network structure remains hidden to the outside world. The ideal is to not reveal any information of the internal network to the outside. This applies equally to the customer networks as to the MPLS core.

Denial-of-service attacks against a core router, for example, are much easier to carry out if an attacker knows the IP address. Where addresses are not known, they can be guessed, but when the MPLS core structure is hidden, attacks are more difficult to make. Ideally, the MPLS core should be as invisible to the outside world as a comparable Layer 2 infrastructure (for example, Frame Relay or ATM).

In practice, a number of additional security measures have to be taken, most of all *extensive packet filtering*. MPLS does not reveal unnecessary information to the outside, not even to customer VPNs. The addressing in the core can be done with either private addresses or public addresses. Since the interface to the VPNs, as well as potentially to the Internet, is BGP, there is no need to reveal any internal information. The only information required in the case of a routing protocol between a PE and CE is the address of the PE router. If this is not desired, you can configure static routing between the PE and CE. With this measure, the MPLS core can be kept completely hidden.

To ensure reachability across the MPLS cloud, customer VPNs will have to advertise their routes as a minimum to the MPLS core. While this could be seen as too open, the information known to the MPLS core is not about specific hosts, but networks (routes); this offers some degree of abstraction. Also, in a VPN-only MPLS network (that is, no shared Internet access), this is equal to existing Layer 2 models, where the customer has to trust the service provider to some degree. Also in a Frame Relay or ATM network, routing information about the VPNs can be seen on the core network.

In a VPN service with shared Internet access, the service provider typically announces the routes of customers that wish to use the Internet to his upstream or peer providers. This can be done via a network address translation (NAT) function to further obscure the addressing information of the customers' networks. In this case, the customer does not reveal more information to the general Internet than with a general Internet service. Core information is not revealed at all, except for the peering addresses of the PE router) that hold the peering with the Internet.

In summary, in a pure MPLS VPN service, where no Internet access is provided, the level of information hiding is as good as on a comparable Frame Relay or ATM network—no addressing information is revealed to third parties or the Internet. If a customer chooses to access the Internet via the MPLS core, he will have to reveal the same addressing structure as for a normal Internet service. NAT can be used for further address hiding.

If an MPLS network has no interconnections to the Internet, this is equal to Frame Relay or ATM networks. With Internet access from the MPLS cloud, the service provider has to reveal at least one IP address (of the peering PE router) to the next provider, and thus the outside world.

Resistance to Attacks

It is not possible to directly intrude into other VPNs. However, it is possible to attack the MPLS core, and try to attack other VPNs from there. There are two basic ways the MPLS core can be attacked:

- Attacking the PE routers directly.

- Attacking the signaling mechanisms of MPLS (mostly routing)

There are two basic types of attacks: *denial-of-service (DoS) attacks*, where resources become unavailable to authorized users, and *intrusion attacks*, where the goal is to gain unauthorized access to resources.

For intrusion attacks, give unauthorized access to resources, there are two basic ways to protect the network:

- Harden protocols that could be abused (for example, Telnet to a router)
- Make the network as inaccessible as possible. This is achieved by a combination of packet filtering or firewalling and hiding the IP addresses in the MPLS core.

Denial-of service attacks are easier to execute, since in the simplest case, a known IP address might be enough to attack a machine. The only way to be certain that you are not be vulnerable to this kind of attack is to make sure that machines are not reachable, again by packet filtering and pinging IP addresses.

MPLS networks must provide at least the same level of protection against both forms of attack as current Layer 2 networks provide.

To attack an element of an MPLS network it is first necessary to know this element, that is, its IP address. It is possible to hide the addressing structure of the MPLS core to the outside world, as discussed in the previous section. Thus, an attacker does not know the IP address of any router in the core that he wants to attack. The attacker could guess addresses and send packets to these addresses. However, due to the address separation of MPLS, each incoming packet is treated as belonging to the address space of the customer. It is therefore impossible to reach an internal router, even through guessing the IP addresses. There is only one exception to this rule—the peer interface of the PE router.

Securing the Routing Protocol

The routing between the VPN and the MPLS core can be configured two ways:

1. **Static**. In this case, the PE routers are configured with static routes to the networks behind each CE, and the CEs are configured to statically point to the PE router for any network in other parts of the VPN (usually a default route).

   The static route can point to the IP address of the PE router, or to an interface of the CE router (for example, serial0).

   Although in the static case the CE router does not know any IP addresses of the PE router, it is still attached to the PE router via some method, and could guess the address of the PE router and try to attack it with this address.

   In the case of a static route from the CE router to the PE router, which points to an interface, the CE router does not need to know any IP address of the core network, not even of the PE router. This has the disadvantage of a more extensive (static) configuration, but from a security point of view, it is preferable to the other cases.

2. **Dynamic**. A routing protocol (for example, RIP, OSPF, or BGP) is used to exchange the routing information between the CE and the PE at each peering point.

In all other cases, each CE router needs to know at least the router ID (RID; peer IP address) of the PE router in the MPLS core, and thus has a potential destination for an attack.

In practice, access to the PE router over the CE-PE interface can be limited to the required routing protocol by using access control lists (ACLs). This limits the point of attack to one routing protocol, for example BGP. A potential attack could send an extensive number of routes, or flood the PE router with routing updates. Both of these attacks could lead to a denial-of-service attack, however, not to an intrusion attack.

To restrict this risk it is necessary to configure the routing protocol on the PE router as securely as possible. This can be done in various ways:

- Use ACLs. Allow the routing protocol only from the CE router, not from anywhere else. Furthermore, no access other than that should be allowed to the PE router in the inbound ACL on each PE interface.

  ACLs must be configured to limit access only to the port(s) of the routing protocol, and only from the CE router.

- Where available, configure MD-5 authentication for routing protocols.

  This is available for BGP, OSPF, and RIP2. It avoids the possibility that packets could be spoofed from other parts of the customer network than the CE router. This requires that the service provider and customer agree on a shared secret between all CE and PE routers. The problem here is that it is necessary to do this for all VPN customers; it is not sufficient to do this only for the customer with the highest security requirements.

  MD5 authentication in routing protocols should be used on all PE-CE peers. It is easy to track the source of such a potential denial-of-service attack.

- Configure, where available, the parameters of the routing protocol to further secure this communication.

  In BGP, for example, it is possible to configure *dampening*, which limits the number of routing interactions. Also, a maximum number of routes accepted per VRF should be configured where possible.

In summary, it is not possible to intrude from one VPN into other VPNs or the core. However, it is theoretically possible to exploit the routing protocol to execute a denial-of-service attack against the PE router. This in turn might have negative impact on other VPNs. For this reason, PE routers must be extremely well secured, especially on their interfaces to the CE routers.

Label Spoofing

Assuming the address and routing separation as discussed above, a potential attacker might try to gain access to other VPNs by inserting packets with a label that he does not own. This is called *label spoofing*. This kind of attack can be done from the outside, that is, another CE router or from the Internet, or from within the MPLS core. The latter case (from within the core) is not discussed since the assumption is that the core network is provided in a secure manner. Should protection against an insecure core be required, it is necessary to run IPsec on top of the MPLS infrastructure.

Within the MPLS network, packets are not forwarded based on the IP destination address, but based on the labels that are prepended by the PE routers. Similar to IP spoofing attacks, where an attacker replaces the source or destination IP address of a packet, it is also possible to spoof the label of an MPLS packet.

The interface between any CE router and its peering PE router is an IP interface, that is, without labels. The CE router is unaware of the MPLS core, and is only aware of the destination router. The intelligence exits in the PE device, where based on the configuration, the PE chooses a label and prepends it to the packet. This is the case for all PE routers, toward CE routers, as well as to the upstream service provider. All interfaces into the MPLS cloud require IP packets without labels.

For security reasons, a PE router should never accept a packet with a label from a CE router. Cisco routers implementation is such that packets that arrive on a CE interface with a label are dropped. Thus, it is not possible to insert fake labels because no labels are accepted.

There remains the possibility to spoof the IP address of a packet that is being sent to the MPLS core. However, since there is strict addressing separation within the PE router, and each VPN has its own VRF, this can only do harm to the VPN the spoofed packet originated from, in other words, a VPN customer can attack himself. MPLS does not add any security risk here.

Securing the MPLS Core

The following is a list of recommendations and considerations on configuring an MPLS network securely.

Note The security of the overall solution depends on the security of its weakest link. This could be the weakest single interconnection between a PE and a CE, an insecure access server, or an insecure TFTP server.

Trusted Devices

The PE and P devices, as well as remote access servers and AAA servers must be treated as trusted systems. This requires strong security management, starting with physical building security and including issues such as access control, secure configuration management, and storage. There is ample literature available on how to secure network elements, so these topics are not discussed here in more detail.

CE routers are typically not under full control of the service provider and must be treated as “untrusted.”

PE-CE Interface

The interface between PE and CE routers is crucial for a secure MPLS network. The PE router should be configured as close as possible. From a security point of view, the best option is to configure the interface to the CE router unnumbered and route statically.

Packet filters (Access Control Lists) should be configured to permit only one specific routing protocol to the peering interface of the PE router, and only from the CE router. All other traffic to the router and the internal service provider network should be denied. This avoids the possibility that the PE and P routers can be attacked, since all packets to the corresponding address range are dropped by the PE router. The only exception is the peer interface on the PE router for routing purposes. This PE peer interface must be secured separately.

If private address space is used for the PE and P routers, the same rules with regard to packet filtering apply—it is required to filter all packets to this range. However, since addresses of this range should not be routed over the Internet, it limits attacks to adjacent networks.

Routing Authentication

All routing protocols should be configured with the corresponding authentication option toward the CEs and toward any Internet connection. Specifically: BGP, OSPF, and RIP2. All peering relationships in the network need to be secured this way:

- CE-PE link: use BGP MD-5 authentication
- PE-P link: use LDP MD5 authentication

- P-P

This prevents attackers from spoofing a peer router and introducing bogus routing information. Secure management is particularly important regarding configuration files, which often contain shared secrets in clear text (for example for routing protocol authentication).

Separation of CE-PE Links

If several CEs share a common Layer 2 infrastructure to access the same PE router (for example, an ethernet VLAN), a CE router can spoof packets as belonging to another VPN that also has a connection to this PE router. Securing the routing protocol is not sufficient, since this does not affect normal packets.

To avoid this problem, Cisco recommends that you implement separate physical connections between CEs and PEs. The use of a switch between various CE routers and a PE router is also possible, but it is strongly recommended to put each CE-PE pair into a separate VLAN to provide traffic separation. Although switches with VLANs increase security, they are not unbreakable. A switch in this environment must thus be treated as a trusted device and configured with maximum security.

LDP Authentication

The Label Distribution Protocol (LDP) can also be secured with MD-5 authentication across the MPLS cloud. This prevents hackers from introducing bogus routers, which would participate in the LDP.

Connectivity Between VPNs

MPLS provides VPN services with address and routing separation between VPNs. In many environments, however, the devices in the VPN must be able to reach destinations outside the VPN. This could be for Internet access or for merging two VPNs, for example, in the case of two companies merging. MPLS not only provides full VPN separation, but also allows merging VPNs and accessing the Internet.

To achieve this, the PE routers maintain various tables: A *routing context table* is specific to a CE router, and contains only routes from this particular VPN. From there, routes are propagated into the *VRF* (virtual routing and forwarding instance) *routing table*, from which a *VRF forwarding table* is calculated.

For separated VPNs, the VRF routing table contains only routes from one routing context. To merge VPNs, different routing contexts (from different VPNs) are put into one single VRF routing table. In this way, two or several VPNs can be merged to a single VPN. In this case, it is necessary that all merged VPNs have mutually exclusive addressing spaces; in other words, the overall address space must be unique for all included VPNs.

For a VPN to have Internet connectivity, the same procedure is used: Routes from the Internet VRF routing table (the default routing table) are propagated into the VRF routing table of the VPN that requires Internet access. Alternatively to propagating all Internet routes, a default route can be propagated. In this case, the address space between the VPN and the Internet must be distinct. The VPN must use private address space since all other addresses can occur in the Internet.

From a security point of view, the merged VPNs behave like one logical VPN, and the security mechanisms described above apply now between the merged VPN and other VPNs. The merged VPN must have unique address space internally, but further VPNs can use the same address space without interference. Packets from and to the merged VPNs cannot be routed to other VPNs. All the separation functions of MPLS apply also for merged VPNs with respect to other VPNs.

If two VPNs are merged in this way, hosts from either part can reach the other part as if the two VPNs were a common VPN. With the standard MPLS features, there is no separation or firewalling or packet filtering between the merged VPNs. Also, if a VPN receives Internet routes through MPLS/BGP VPN mechanisms, firewalling or packet filtering has to be engineered in addition to the MPLS features.

MP-BGP Security Features

Security in VPN Solutions Center MPLS-based networks is delivered through a combination of MP-BGP and IP address resolution. In addition, service providers can ensure that VPNs are isolated from each other.

Multiprotocol BGP is a routing information distribution protocol that, through employing multiprotocol extensions and community attributes, defines who can talk to whom. VPN membership depends upon logical ports entering the VPN, where MP-BGP assigns a unique Route Distinguisher (RD) value (see "Route Distinguishers and Route Targets" below).

RDs are unknown to end users, making it impossible to enter the network on another access port and spoof a flow. Only preassigned ports are allowed to participate in the VPN. In an MPLS VPN, MP-BGP distributes forwarding information base (FIB) tables about VPNs to members of the same VPN only, providing native security via logical VPN traffic separation. Furthermore, IBGP PE routing peers can perform TCP segment protection using the MD5 Signature Option when establishing IBGP peering relationships, further reducing the likelihood of introducing spoofed TCP segments into the IBGP connection stream among PE routers (for information on the MD5 Signature Option, see RFC 2385).

The service provider, not the customer, associates a specific VPN with each interface when provisioning the VPN. Users can only participate in an intranet or extranet if they reside on the correct physical or logical port and have the proper RD. This setup makes a Cisco MPLS VPN virtually impossible to enter.

Within the core, a standard Interior Gateway Protocol (IGP) such as OSPF or IS-IS distributes routing information. Provider edge routers set up paths among one another using LDP to communicate label-binding information. Label binding information for external (customer) routes is distributed among PE routers using MP-BGP multiprotocol extensions instead of LDP, because they easily attach to VPN IP information already being distributed.

The MP-BGP community attribute constrains the scope of reachability information. MP-BGP maps FIB tables to provider edge routers belonging to only a particular VPN, instead of updating all edge routers in the service provider network.

Security Through IP Address Resolution

MPLS VPN networks are easier to integrate with IP-based customer networks. Subscribers can seamlessly interconnect with a provider service without changing their intranet applications because MPLS-based networks have built-in application awareness. Customers can even transparently use their existing IP address space without Network Address Translator (NAT) because each VPN has a unique identifier.

MPLS VPNs remain unaware of one another. Traffic is separated among VPNs using a logically distinct forwarding table and RD for each VPN. Based on the incoming interface, the PE selects a specific forwarding table, which lists only valid destinations in the VPN. To create extranets, a provider explicitly configures reachability among VPNs.

The forwarding table for a PE contains only address entries for members of the same VPN. The PE rejects requests for addresses not listed in its forwarding table. By implementing a logically separate forwarding table for each VPN, each VPN itself becomes a private, connectionless network built on a shared infrastructure.

IP limits the size of an address to 32 bits in the packet header. The VPN IP address adds 64 bits in front of the header, creating an extended address in routing tables that classical IP cannot forward. MPLS solves this problem by forwarding traffic based on labels, so one can use MPLS to bind VPN IP routes to label-switched paths. PEs are concerned with reading labels, not packet headers. MPLS manages forwarding through the provider's MPLS core. Since labels only exist for valid destinations, this is how MPLS delivers both security and scalability.

When a virtual circuit is provided using the overlay model, the egress interface for any particular data packet is a function solely of the packet's ingress interface; the IP destination address of the packet does not determine its path in the backbone network. Thus, unauthorized communication into or out of a VPN is prevented.

In MPLS VPNs, a packet received by the backbone is first associated with a particular VPN by stipulating that all packets received on a certain interface (or subinterface) belong to a certain VPN. Then its IP address is looked up in the forwarding table associated with that VPN. The routes in that forwarding table are specific to the VPN of the received packet.

In this way, the ingress interface determines a set of possible egress interfaces, and the packet's IP destination address is used to choose from among that set. This prevents unauthorized communication into and out of a VPN.

Ensuring VPN Isolation

To maintain proper isolation of one VPN from another, it is important that the provider routers not accept a labeled packet from any adjacent PE unless the following conditions are met:

- The label at the top of the label stack was actually distributed by the provider router to the PE device.
- The provider router can determine that use of that label will cause the packet to exit the backbone before any labels lower in the stack and the IP header will be inspected.

These restrictions are necessary to prevent packets from entering a VPN where they do not belong.

The VRF tables in a PE are used only for packets arriving from a CE that is directly attached to the PE device. They are not used for routing packets arriving from other routers that belong to the service provider backbone. As a result, there may be multiple different routes to the same system, where the route followed by a given packet is determined by the site from which the packet enters the backbone. So one may have one route to a given IP network for packets from the extranet (where the route leads to a firewall), and a different route to the same network for packets from the intranet.

VPN Routing and Forwarding Tables (VRFs)

The VPN routing and forwarding table (VRF) is a key element in the MPLS VPN technology. VRFs exist on PEs only. A VRF is a routing table instance, and more than one VRF can exist on a PE. A VPN can contain one or more VRFs on a PE.The VRF contains routes that should be available to a particular set of sites. VRFs use CEF technology, therefore the VPN must be CEF-enabled.

A VRF is associated with the following elements:

- IP routing table
- Derived forwarding table, based on the Cisco Express Forwarding (CEF) technology
- A set of interfaces that use the derived forwarding table
- A set of routing protocols and routing peers that inject information into the VRF

Each PE maintains one or more VRFs. VPNSC software looks up a particular packet's IP destination address in the appropriate VRF only if that packet arrived directly through an interface that is associated with that VRF. The so-called "color" MPLS label tells the destination PE to check the VRF for the appropriate VPN so that it can deliver the packet to the correct CE and finally to the local host machine.

A VRF is named based on the VPN or VPNs it services, and on the role of the CE in the topology. The schemes for the VRF names are as follows:

The VRF name for a hub: `ip vrf V`*x*`:[`*VPN_name*`]`

The *x* parameter is a number assigned to make the VRF name unique.

For example, if we consider a VPN called Blue, then a VRF for a hub CE would be called:

```
ip vrf V1:blue
```

A VRF for a spoke CE in the Blue VPN would be called:

```
ip vrf V1:blue-s
```

A VRF for an extranet VPN topology in the Green VPN would be called:

```
ip vrf V1:green-etc
```

Thus, you can read the VPN name and the topology type directly from the name of the VRF.

Figure 1-5 shows a network in which two of the four sites are members of two VPNs, and illustrates which routes are included in the VRFs for each site.

Figure 1-5 VRFs for Sites in Multiple VPNs

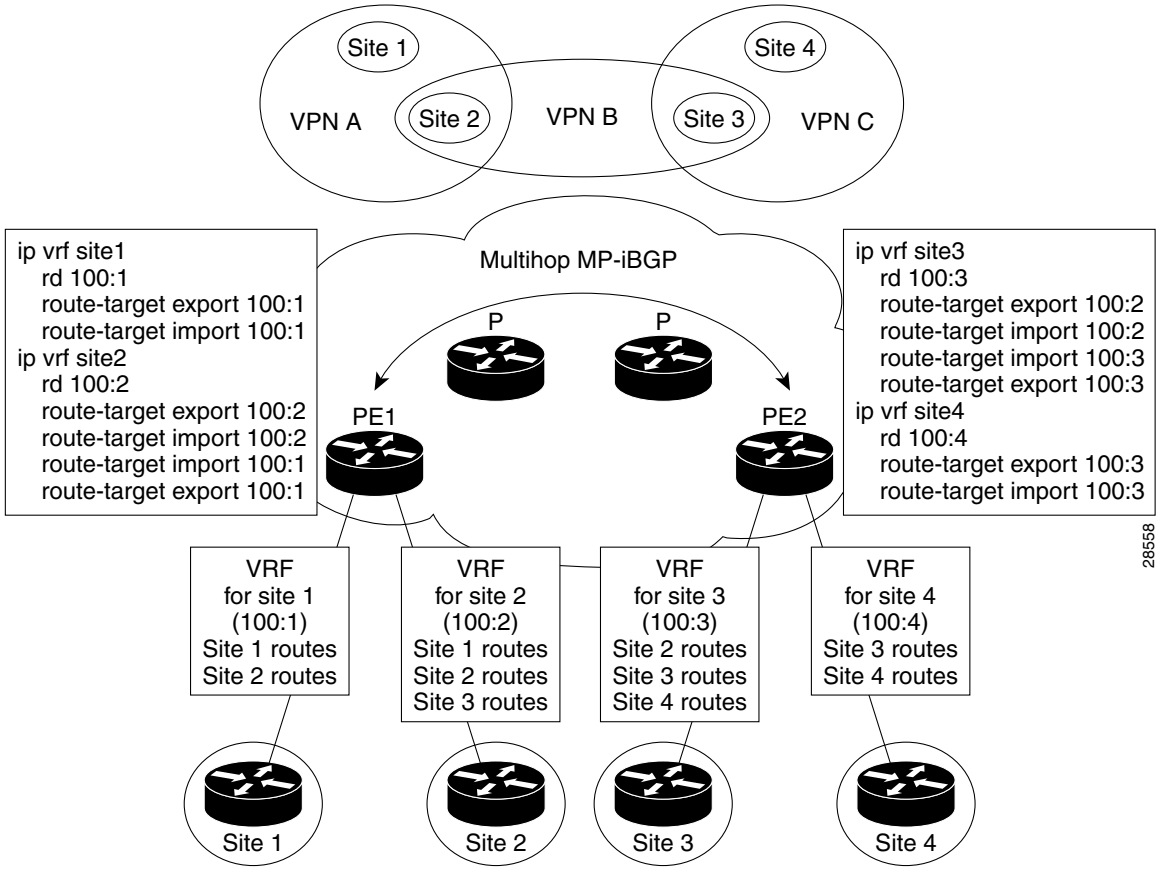

VRF Implementation Considerations

When implementing VPNs and VRFs, Cisco recommends you keep the following considerations in mind:

- A local VRF interface on a PE is not considered a directly-connected interface in a traditional sense. When you configure, for example, a Fast Ethernet interface on a PE to participate in a particular VRF/VPN, the interface no longer shows up as a directly-connected interface when you issue a **show ip route** command. To see that interface in a routing table, you must issue a **show ip route vrf** *vrf_name* command.
- The global routing table and the per-VRF routing table are independent entities. Cisco IOS commands apply to IP routing in a global routing table context. For example, `show ip route`, and other EXEC-level show commands—as well as utilities such as **ping**, **traceroute**, and **telnet**—all invoke the services of the Cisco IOS routines that deal with the global IP routing table.
- You can issue a standard Telnet command from a CE router to connect to a PE router. However, from that PE, you must issue the following command to connect from the PE to the CE:

  telnet *CE_RouterName* **/vrf** *vrf_name*

  Similarly, you can utilize the **Traceroute** and **Ping** commands in a VRF context.

- The MPLS VPN backbone relies on the appropriate Interior Gateway Protocol (IGP) that is configured for MPLS, for example, EIGRP, or OSPF. When you issue a **show ip route** command on a PE, you see the IGP-derived routes connecting the PEs together. Contrast that with the **show ip route vrf** *VRF_name* command, which displays routes connecting customer sites in a particular VPN.

Creating a VRF Instance

The configuration commands to create a VRF instance are as follows:

	Command	Description
Step 1	`Router# configure terminal` `Router(config)#`	Enter global configuration mode.
Step 2	`Router(config)# ip vrf` *vrf_name*	For example, `ip vrf` *CustomerA* initiates a VPN routing table and an associated CEF table named CustomerA. The command enters VRF configuration submode to configure the variables associated with the VRF.
Step 3	`Router(config-vrf)# rd` *RD_value*	Enter the eight-byte route descriptor (RD) or IP address. The PE prepends the RD to the IPv4 routes prior to redistributing the route into the MPLS VPN backbone.
Step 4	`Router(config-vrf)# route-target import | export | both` *community*	Enter the route-target information for the VRF.

For detailed information about these configuration commands, refer to Appendix B, "Cisco VPNSC: MPLS Solution Command Reference."

Route Distinguishers and Route Targets

MPLS-based VPNs employ BGP to communicate between PEs to facilitate customer routes. This is made possible through extensions to BGP that carry addresses other than IPv4 addresses. A notable extension is called the *route distinguisher* (RD).

The purpose of the route distinguisher (RD) is to make the prefix value unique across the backbone. Prefixes should use the same RD if they are associated with the same set of route targets (RTs) and anything else that is used to select routing policy. The community of interest association is based on the route target (RT) extended community attributes distributed with the Network Layer Reachability Information (NLRI).The RD value must be a globally unique value to avoid conflict with other prefixes.

The MPLS label is part of a BGP routing update. The routing update also carries the addressing and reachability information. When the RD is unique across the MPLS VPN network, proper connectivity is established even if different customers use non-unique IP addresses.

For the RD, every CE that has the same overall role should use a VRF with the same name, same RD, and same RT values. The RDs and RTs are *only* for route exchange between the PEs running BGP. That is, for the PEs to do MPLS VPN work, they have to exchange routing information with more fields than usual for IPv4 routes; that extra information includes (but is not limited to) the RDs and RTs.

The route distinguisher values are chosen by the VPN Solutions Center software.

- CEs with hub connectivity use `bgp_AS:`*value*.
- CEs with spoke connectivity use `bgp_AS:`*value* `+ 1`

Each spoke uses its own RD value for proper hub and spoke connectivity between CEs; therefore, the VPN Solutions Center software implements a new RD for each spoke that is provisioned.

VPN Solutions Center chooses route target values by default, but you can override the automatically assigned RT values if necessary when you first define a CERC in the VPN Solutions Center software (see the "Defining CE Routing Communities" section on page 5-3).

Route Target Communities

The mechanism by which MPLS VPN controls distribution of VPN routing information is through the VPN route-target extended MP-BGP communities. An extended MP-BGP community is an eight octet structure value. MPLS VPN uses route-target communities as follows:

- When a VPN route is injected into MP-BGP, the route is associated with a list of VPN route-target communities. Typically, this is set through an export list of community values associated with the VRF from which the route was learned.
- An import list of route-target communities is associated with each VRF. This list defines the values that should be matched against to decide whether a route is eligible to be imported into this VRF.

  For example, if the import list for a particular VRF is {A, B, C}, then any VPN route that carries community value A, B, or C is imported into the VRF.

CE Routing Communities

A VPN can be organized into subsets called *CE routing communities*, or CERCs. A CERC describes how the CEs in a VPN communicate with each other. Thus, CERCs describe the logical topology of the VPN. VPN Solutions Center can be employed to form a variety of VPN topologies between CEs by building hub and spoke or full mesh CE routing communities. CERCs are building blocks that allow you to form complex VPN topologies and CE connectivity.

The most common types of VPNs are *hub-and-spoke* and *full mesh*.

- A hub-and-spoke CERC is one in which one or a few CEs act as hubs, and all spoke CEs talk only to or through the hubs, never directly to each other.
- A full mesh CERC is one in which every CE connects to every other CE.

These two basic types of VPNs—full mesh and hub and spoke—can be represented with a single CERC.

Whenever you create a VPN, the VPN Solutions Center software creates one default CERC for you. This means that until you need advanced customer layout methods, you will not need to define new CERCs. Up to that point, you can think of a CERC as standing for the VPN itself—they are one and the same. If, for any reason, you need to override the software's choice of route target values, you can do so only at the time you create a CERC in the VPN Solutions Center software (see the "Defining CE Routing Communities" section on page 5-3).

To build very 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. (Note that a CE can be in more than one group at a time, so long as each group has one of the two basic 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, the provisioning software does the rest, assigning route target values and VRF tables to arrange exactly the connectivity the customer requires. You can use the Topology tool to double-check the CERC memberships and resultant VPN connectedness.

VPN Solutions Center supports multiple CEs per site and multiple sites connected to the same PE. Each CERC has unique route targets (RT), route distinguisher (RD) and VRF naming. After provisioning a CERC, it is a good idea to run the audit reports to verify the CERC deployment and view the topologies created by the service requests. The product supports linking two or more CE routing communities in the same VPN.

Figure 1-6 shows several examples of the topologies that VPN Solutions Center CERCs can employ.

Figure 1-6 Examples of CERC Topologies

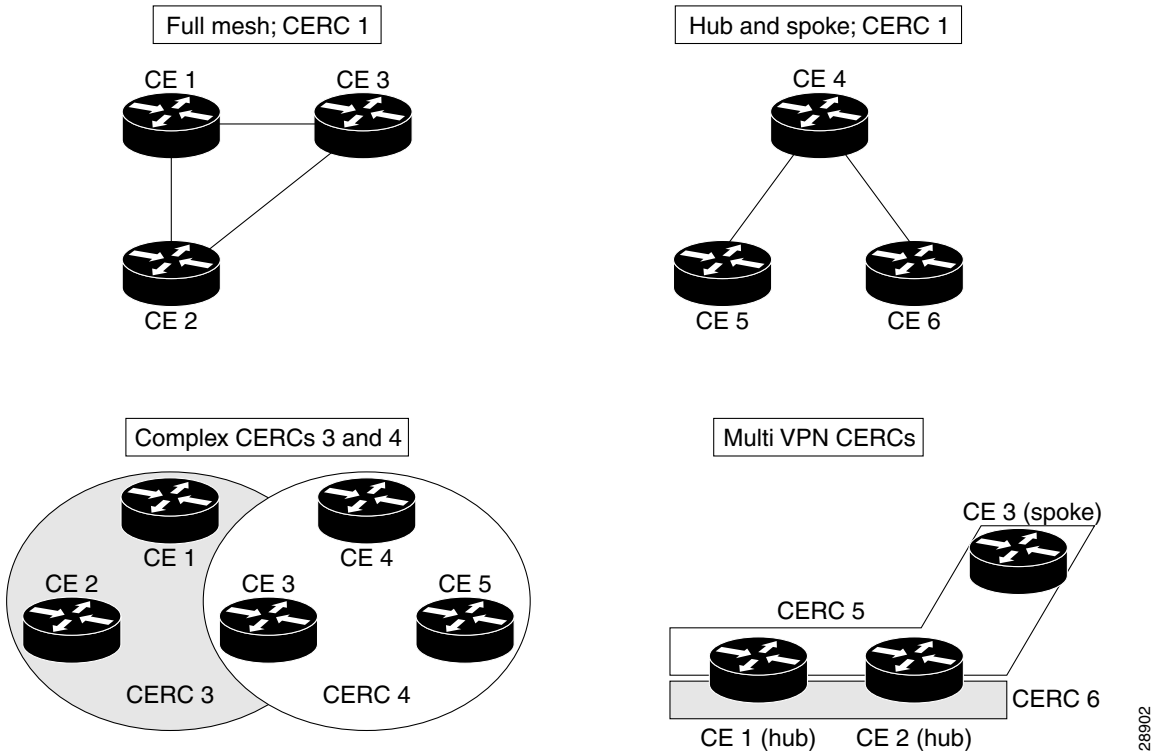

Hub and Spoke Considerations

In hub-and-spoke MPLS VPN environments, the spoke routers have to have unique Route Distinguishers (RDs). In order to use the hub site as a transit point for connectivity in such an environment, the spoke sites export their routes to the hub. Spokes can talk to hubs, but spokes never have routes to other spokes.

Due to the current MPLS VPN implementation, you must apply a different RD for each spoke VRF. The MP-BGP selection process applies to all the routes that have to be imported into the same VRF plus all routes that have the same RD of such a VRF. Once the selection process is done, only the best routes are imported. In this case this can result in a best route which is not imported. Thus, customers must have different RDs per spoke-VRF.

Full Mesh Considerations

Each CE Routing Community (CERC) has two distinct RTs: a hub RT and a spoke RT. When building a full mesh topology, always use the hub RT. Thus, when a need arises to add a spoke site for the current full mesh topology, you can easily add the spoke site without reconfiguring any of the hub sites. The existing spoke RT can be used for this purpose. This is a strategy to prevent having to do significant reprovisioning of a full mesh topology to a hub-and-spoke topology.

About the Multi-VRF CE

The Multi-VRF CE is a feature that was introduced in Cisco IOS release 12.2(4)T. The Multi-VRF CE functionality extends some of the functionality formerly reserved to the PE to a CE router in an MPLS VPN—the only PE-like functionality that this feature provides is the ability to have multiple VRFs on the CE router so that different routing decisions can be made. The packets are sent toward the PE as IP packets.

With this feature, a CE can maintain separate VRF tables to extend the privacy and security of an MPLS VPN down to a branch office, rather than just at the PE router node.

A Multi-VRF CE is unlike a CE in that there is no label exchange, no LDP adjacency, and no labeled packet flow between the PE and the CE.

CE routers use VRF interfaces to form a VLAN-like configuration on the customer side. Each VRF on the CE router is mapped to a VRF on the PE router. With Multi-VRF CE functionality, the CE router can only configure VRF interfaces and support VRF routing tables.

Figure 1-7 illustrates one method in which a Multi-VRF CE can be used. The Multi-VRF CE router associates a specific VRF by the clients connected to its interfaces and exchanges that information with the PE. Routes are installed in the VRF on the Multi-VRF CE. There also needs to be a routing protocol or a static route that propagates routes from a specific VRF on the Multi-VRF CE to the corresponding VRF on the PE.

Figure 1-7 A Multi-VRF CE in an MPLS VPN Environment

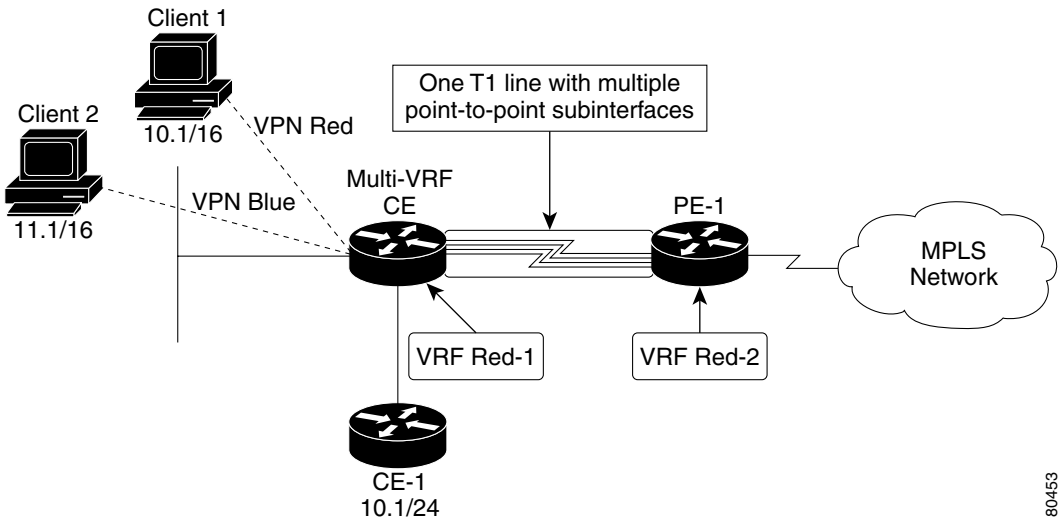

The Multi-VRF CE feature can segment its LAN traffic by placing each client or organization with its own IP address space, either on separate Ethernet interfaces such as CE-1, or through one Fast Ethernet interface segmented into multiple subinterfaces (for Client-1 and Client-2). To differentiate each client, each subinterface contains its own IP address space.

When receiving an outbound customer data packet from a directly attached interface, the Multi-VRF CE router performs a route lookup in the VRF that is associated with that site. The specific VRF is determined by the interface or subinterface over which the data packet is received. Support for multiple forwarding tables makes it easy for the CE router to provide segregation of routing information on a

per-VPN basis before the routing information is sent to the PE. The use of a T1 line with multiple point-to-point subinterfaces allows traffic from the Multi-VRF CE router to the PE router to be segmented into each individual VRF.

With Multi-VRF CE configured on the CE router, the data path is as follows from the clients to PE-1 (as shown in Figure 1-7):

1. The Multi-VRF CE learns the VPN Red routes to Client 1 from a subinterface of the Fast Ethernet interface directly attached to Multi-VRF CE.
2. The Multi-VRF CE then installs these routes into VRF Red-1 (the VRF on the Multi-VRF CE).
3. PE-1 learns the VPN Red routes to Client 1 from VRF Red-1 on the Multi-VRF CE and installs the routes into VRF Red-2 (on PE-1).
4. The local VPN Blue routes from Client 2 are not associated with VPN Red and are not imported into VRF Red-1 or VRF Red-2.

Overview of the MPLS VPN Cable Feature

Using MPLS VPN technology, service providers can create scalable and efficient private networks using a shared Hybrid Fiber Coaxial (HFC) network and Internet Protocol (IP) infrastructure. The cable MPLS VPN network consists of the following two major elements:

- The Multiple Service Operator (MSO) or cable company that owns the physical infrastructure and builds VPNs for the Internet Service Providers (ISPs) to move traffic over the cable and IP backbone.
- ISPs that use the HFC network and IP infrastructure to supply Internet service to cable customers.

You can find the complete description on how to use VPN Solutions Center software to provision cable MPLS VPNs in Chapter 9, “Provisioning MPLS VPN Cable Services.”

Benefits of Cable MPLS VPNs

Provisioning cable services with MPLS VPNs provides the following benefits:

- MPLS VPNs give cable MSOs and ISPs a manageable way of supporting multiple access to a cable plant.

  Service providers can create scalable and efficient VPNs across the core of their networks MPLS VPNs provide systems support scalability in cable transport infrastructure and management.
- Each ISP can support Internet access services from a subscriber’s PC through an MSO’s physical cable plant to their networks.
- MPLS VPNs allow MSOs to deliver value-added services through an ISP, and thus, deliver connectivity to a wider set of potential customers.

  MSOs can partner with ISPs to deliver multiple services from multiple ISPs and add value within the MSO’s own network using VPN technology.
- Subscribers can select combinations of services from various service providers.
- The Cisco IOS MPLS VPN cable feature sets build on Cable Modem Termination Server (CMTS) and DOCSIS 1.0 extensions to ensure services are reliably and optimally delivered over the cable plant.

MPLS VPN provides systems support domain selection, authentication per subscriber, selection of QoS, policy-based routing, and ability to reach behind the cable modem to subscriber end-devices for QoS and billing, while preventing session-spoofing.

- MPLS VPN technology ensures both secure access across the shared cable infrastructure and service integrity.

The Cable MPLS VPN Network

As shown in Figure 1-8, each ISP moves traffic to and from a subscriber's PC, through the MSO's physical network infrastructure, to the ISP's network. MPLS VPNs, created in Layer 3, provide privacy and security by constraining the distribution of a VPN's routes only to the routers that belong to its network. Thus, each ISP's VPN is insulated from other ISPs that use the same MSO infrastructure.

In the MPLS-based cable scheme, a VPN is a private network built over a shared cable plant and MPLS-core backbone. The public network is the shared cable plant or backbone connection points. A cable plant can support Internet access services and carry traffic for an MSO and its subscribers, as well as for multiple Internet Service Providers (ISPs) and their subscribers.

An MPLS VPN assigns a unique VPN Routing/Forwarding (VRF) instance to each VPN. A VRF instance consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine the contents of the forwarding table.

Each PE router maintains one or more VRF tables. If a packet arrives directly through an interface associated with a particular VRF, the PE looks up a packet's IP destination address in the appropriate VRF table. MPLS VPNs use a combination of BGP and IP address resolution to ensure security.

Figure 1-8 Example of an MPLS VPN Cable Network

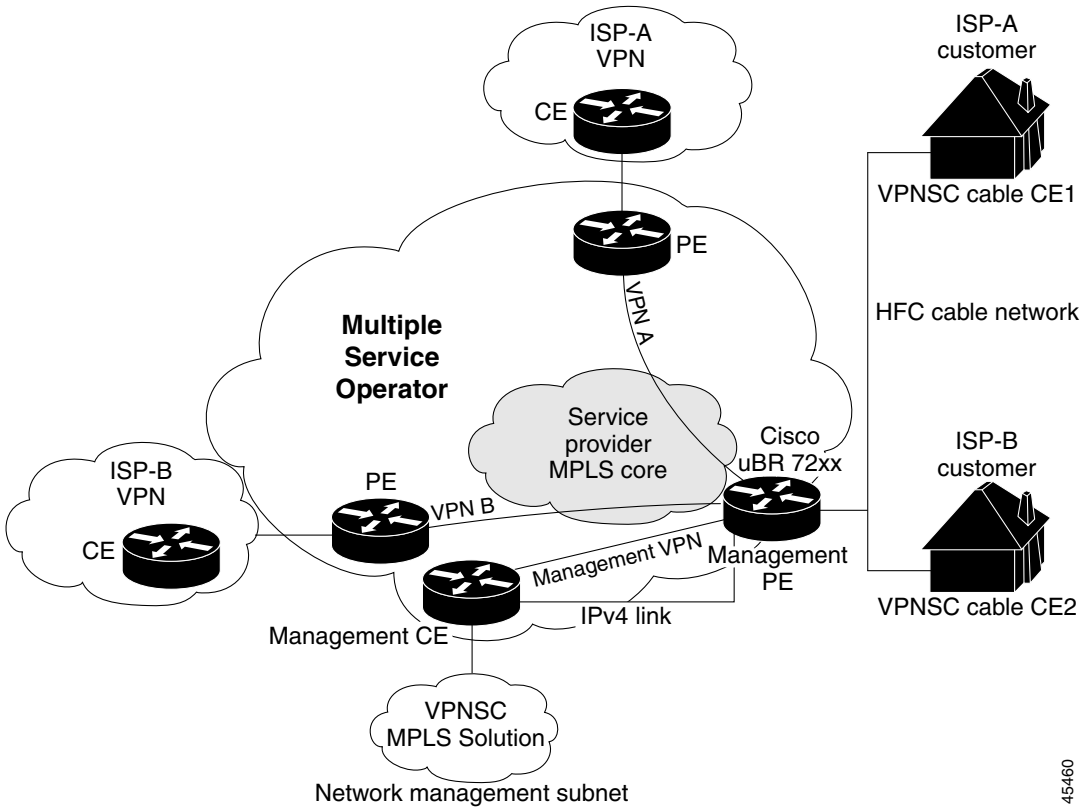

The routers in the cable network are as follows:

- Provider (P) router—Routers in the MPLS core of the service provider network. P routers run MPLS switching, and do not attach VPN labels (MPLS labels in each route assigned by the PE router) to routed packets. VPN labels direct data packets to the correct egress router.
- Provider Edge (PE) router—A router that attaches the VPN label to incoming packets based on the interface or subinterface on which they are received. A PE router attaches directly to a CE router. In the MPLS-VPN approach, each Cisco uBR72xx series router acts as a PE router.
- Customer (C) router—A router in the ISP or enterprise network.
- Customer Edge (CE) router—Edge router on the ISP's network that connects to the PE router on the MSO's network. A CE router must interface with a PE router.
- Cable CE—The cable CE is an object with the VPN Solutions Center only. In the VPN Solutions Center software, the cable CE represents the cable modem and its associated hosts for a particular site.
- Management CE (MCE) router—The MCE *emulates* the role of a customer edge router (CE), but the MCE is in provider space and serves as a network operations center gateway router. The network management subnet is connected to the Management CE (MCE). The MCE is part of a management site as defined in the VPN Solutions Center software.
- Management PE (MPE) router—The MPE *emulates* the role of a PE in the provider core network. The MPE connects the MCE to the provider core network. An MPE can have a dual role as both a PE and the MPE.

The shared cable plant supports Internet connectivity from ISP A to its subscribers and from ISP B to its subscribers.

The Management VPN in the Cable Network

The MPLS network has a unique VPN that exclusively manages the MSOs devices called the *management VPN*. It contains servers and devices that other VPNs can access. The management VPN connects the Management CE (MCE) router and the management subnet to the MSO PE router (a uBr72xx router or equivalent). VPN Solutions Center and the management servers, such as Dynamic Host Configuration Protocol (DHCP), Cisco Network Registrar (CNR) Time of Day (ToD) are part of the management subnet and are within the management VPN for ISP connectivity.

As shown in Figure 1-8, the management VPN is comprised of the network management subnet (where the VPN Solutions Center workstation resides), which is directly connected to the Management CE (MCE). The management VPN is a special VPN between the MCE and the cable VPN gateway. The cable VPN gateway is usually a Cisco uBR 72xx router that functions as both a regular PE and a Management PE. Notice that there is also a parallel IPv4 link between the MCE and the MPE.

Using Templates to Customize Configuration Files

The Template Manager in VPN Solutions Center is a provisioning system that provides fast, flexible, and extensible Cisco IOS command generation capability. The Template Manager defines standard templates to generate Cisco IOS configurations for common provisioning tasks, such as common IPv4, QoS, and VPN provisioning. For details, see Chapter 10, "Provisioning with the VPN Solutions Center Template Manager."

- A *template file* is a file created by the Template Manager that stores a VPN Solutions Center template definition.
- A *template data file* is a text file that stores variable values to generate the template file. A valid data file contains name-value pairs for all the variables defined in a template. Each template file can be associated with multiple data files; however, note that each data file can only be associated with a single template. You can select which data file to use to generate a template. The filename suffix for data files is *.dat*.
- A *template configuration file* is an IOS configuration file that stores the Cisco IOS commands created by the Template Manager. A template configuration file can be either a partial or complete configuration file. When you generate a template configuration file using a particular data file, the template configuration filename is the same as the data file's name.

The template data files are tightly linked with its corresponding template. You can use a data file and its associated template to create a template configuration file. The template configuration file is merged with (either appended to or prepended to) the VPNSC configlet. VPN Solutions Center downloads the combined configlet to the edge device router.

You can apply the same template to multiple edge devices, assigning the appropriate template data file for each device. Each template data file includes the specific data for a particular device (for example, the management IP address or host name of each device).

The template files and data files are in XML format. The template file, its data files, and all template configuration file files are mapped to a single directory.

- VPN Solutions Center creates the initial VPNSC configlet. Through the Template Manager, you can create a template configuration file. You can then associate a template configuration file with a service request, which effectively merges the VPNSC configlet and the template configuration file. For details on this process, see the "Templates" section on page 5-46. You can then download this merged VPNSC configlet to the target router (or routers).
- You can also create a template configuration file and download it directly to a router as described in the "Provisioning a Template Configuration File Directly to a Router" section on page 10-24.

Uses for the Templating Function

Service providers can use the Template Manager to enhance VPN Solutions Center functionality. Because the Telnet Gateway Server (TGS) supports console access to VPN Solutions Center targets, you can use the Template Manager to provide initial configuration for any service provider core device or edge device.

The Template Manager can be used as a stand-alone tool to generate complete configuration files that you can download to any VPN Solutions Center target.

Some of the additional uses for templating are as follows:

- IOS firewall provisioning
- Add a set of commands that VPN Solutions Center does not include to a service request; for example, provisioning ATM Class of Service.
- Use the templating feature to apply Class of Service using IP connectivity.
- Download a VPN Solutions Center service request and an Cisco IOS configuration file in one download operation through the console. This edge device staging method would create a template and apply the service request in one step.

Event Subscription Service

The Cisco VPN Solutions Center Event Subscription Service (ESS) is an event-notification service (for client-application developers) that allows you to track specific events that may be of interest to your application and your customers. Using the Event Subscription Service, client-application developers can support the following:

- Real-time response to system events
- Local caching of system data
- Synchronization of one or more tasks to create a process flow

While executing, the Cisco VPN Solutions Center software publishes events at the following times:

- Each time an element is created, modified, or destroyed in any of the four VPN Solutions Center repositories
- Each time a scheduled task begins or ends its execution
- When a Watch Dog event signals a change in execution status for any VPN Solutions Center server

Each event contains identifying information that appropriately corresponds with the event type. The ESS is supported by the following:

- Event Gateway server
- Cisco Event Gateway Callback interface (in the CiscoEventGateway.idl file)

- TIB®/Rendezvous™ software

For details on this feature, refer to "Part 5, Using the Event Subscription Service" in the *Cisco VPN Solutions Center: MPLS Solution API Programmer Reference*.

The Event Gateway Server

The Event Gateway server is a CORBA wrapper for the TIB/Rendezvous software that is used by VPN Solutions Center. The Event Gateway Server uses the client-oriented SDKs that are supplied by TIBCO to define its interaction with the TIB/Rendezvous software. The Cisco Event Gateway Callback interface, which is defined in the *CiscoEventGateway.idl* file, provides the client-development facility with which you can interact with the Event Gateway server.

There is no practical limit to the number of clients the Event Gateway server can support. Within the scope of each client, the Event Gateway server can support the Event Gateway Callback objects that subscribe to subjects of interest, which are events generated during the execution of the VPN Solutions Center software.

The Event Gateway Server complements the APIs that VPN Solutions Center provides to enable third party software access to VPNSC data. The TIB/Rendezvous software is the means by which VPNSC signals the occurrences of significant events within VPNSC (for example, the beginning and ending of tasks, the completion of data processing, and so on). If your third party software has special requirements, such as real-time notification of events within VPNSC software, you can use the Event Subscription Service to subscribe to those events.

CHAPTER 2

Setting Up Devices in the VPN Solutions Center MPLS Environment

Cisco VPN Solutions Center: MPLS Solution is an MPLS VPN provisioning and auditing tool. The software focuses on the provider edge routers (PEs), customer edge routers (CEs), and the link between them. VPN Solutions Center software uses the Telnet Gateway Server (TGS) software to transport configuration file information to and from target routers. Additional features include Class of Service (CoS) provisioning, VPN-aware NetFlow traffic profiling, and Service Level Agreement (SLA) monitoring.

The Cisco VPN Solutions Center (VPNSC) also provides external access to its provisioning, traffic profiling, and SLA monitoring features through CORBA APIs.

In an MPLS network, a customer edge router (CE) is connected to a provider edge router (PE) in such a way that the customer's traffic is encapsulated and transparently sent to other CEs, thus creating a virtual private network. The VPN Solutions Center provisioning engine for MPLS accesses the configuration files on both the CE and PE to compute the necessary changes to the configuration files to support the service on the PE-CE link.

Figure 2-1 VPN Solutions Center: MPLS Solution in the Service Provider Network

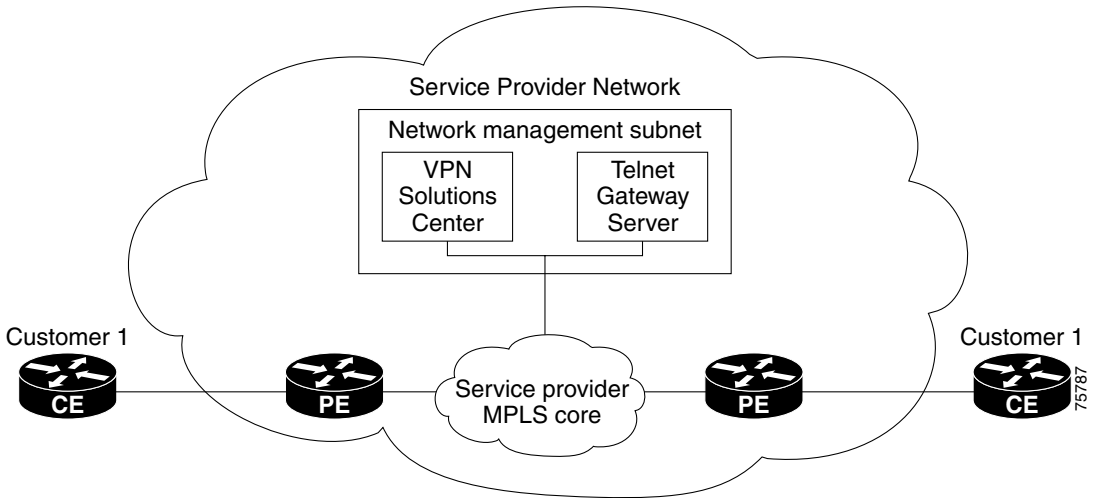

As illustrated in Figure 2-1, Cisco requires that the VPN Solutions Center software is installed on its own dedicated system. The VPN Solutions Center workstation is connected on a LAN to one or more Telnet Gateway servers.

Tasks to Be Completed Before Using VPNSC Software

Before you use VPN Solutions Center: MPLS Solution software to provision an MPLS network, the Service Provider must complete the following tasks:

1. IPv4 connectivity must be operational among all the routers in the MPLS VPN network before provisioning can take place.
2. The Service Provider or Customer must create a loopback interface on each router.
3. Each router must have a routable IP address.
4. Optionally, you can set up the Secure Shell (SSH) on the CE routers (see the next section for details).
5. Set up SNMP on all the edge routers in the network—see the "Setting Up SNMPv1 and SNMPv2 on the Routers in the Service Provider Network" section on page 2-4 and the "Setting the SNMPv3 Parameters on the Routers in the Service Provider Network" section on page 2-5.
6. Enable SA Agent on all edge devices that you want to collect SLA data from—see the "Enabling SA Agent on Edge Routers for SLA Jitter Probes" section on page 2-7.
7. If you choose to use TFTP (Trivial File Transfer Protocol) as the default configuration transport method, you must enable TFTP on the VPN Solutions Center workstation—see the "Enabling TFTP in VPN Solutions Center" section on page 2-7.
8. If you are installing and using Telnet Gateway Servers on remote networks, complete the procedures described in the "Setting Up Connectivity to a Remote Telnet Gateway Server" section on page 2-9.
9. If you are using terminal servers to access routers in the network, you must enable at least as many Telnet sessions on the terminal server as there are terminal server ports. For details, see the "Enabling Telnet Sessions for Terminal Server Ports" section on page 2-27.

Caution Make sure that the file descriptor limit is *not* set in the VPN Solutions Center workstation login shell file (which can be the *.login* file, the *.cshrc* file, or the *.kshrc* file). If the login shell file contains a line with the **ulimit -n** command (for example, "`ulimit -n <number>`"), comment out this command line in the file.

VPN Solutions Center cannot override the file descriptor limitation setting in the login shell file. If the value is set incorrectly, VPN Solutions Center may experience operational problems.

Setting Up Devices in the VPN Solutions Center MPLS Environment

This section describes the tasks the Service Provider should complete to set up devices in the VPN Solutions Center MPLS environment.

Setting Up the Secure Shell (SSH) on Edge Routers

Service providers need a mechanism to deploy VPN configuration files to remote edge routers in a secure manner. The basic requirements for secured management are as follows:

- The edge device routers and VPN Solutions Center must be able to authenticate each other.
- An encrypted channel for uploading and downloading router configuration information must be in place.

VPN Solutions Center 2.1 uses TGS as the configuration file download mechanism. One of the modes that TGS can operate in is *Secure Shell (SSH) mode*. The Telnet Gateway Server uses SSH for both authentication and encryption. In this scheme, the edge device router functions as an SSH server, while VPN Solutions Center functions as the SSH client.

Note This configuration procedure assumes that the router's authentication database is stored locally on the router and not on a TACACS (Terminal Access Controller Access Control System) server.

The procedure for configuring SSH on edge device routers is as follows:

	Command	Description
Step 1	`Router# configure terminal`	Enter global configuration mode.
Step 2	`Router(config)# ip domain-name` *domain_name*	Specify the IP domain name.

Step 3	`Router(config)# crypto key generate rsa`	Generate keys for the SSH session.
		The **crypto key generate rsa** command is interactive. You will see the following prompt:
		`Choose the size of the key modulus in the range of 360 to 2048 for your general purpose keys.`
		`How many bits in the modulus (nnn):`

Step 4		Press **Enter** to accept the default number of bits.
Step 5	`Router(config)# username` *username* `password` *password*	Configure the user ID and password. Enter the VPNSC workstation username and password you are logged in as. For example, **username admin password vpnsc**.
Step 6	`Router(config)# line vty 0 4`	Enable SSH as part of the vty login transport.
Step 7	`Router(config-line)# login local`	The **login** command can take either **local** or **tacacs** as its value. This command indicates that the router stores the authentication information locally.

	Command	Description
Step 8	`Router(config-line)# transport input telnet ssh`	
Step 9	`Router(config-line)# Ctrl+Z`	Return to Privileged Exec mode.
Step 10	`Router# copy running startup`	Save the configuration changes to NVRAM.

Setting Up SNMPv1 and SNMPv2 on the Routers in the Service Provider Network

The Simple Network Management Protocol (SNMP) must be configured on each router and edge device in the service provider network. To determine whether SNMP is enabled and to set the SNMP community strings on a router, execute the following steps for each router.

	Command	Description
Step 1	`> telnet router_name`	Telnet to the router you want to configure.
Step 2	`Router> enable` `Router> enable_password`	Enter enable mode, then enter the enable password.
Step 3	`Router# show snmp`	Check the output of the **show snmp** command to see whether the following statement is present: "*SNMP agent not enabled.*" If SNMP is not enabled, complete the steps in this procedure.
Step 4	`Router# configure terminal`	Enter global configuration mode.
Step 5	`Router(config)# snmp-server community userstring RO`	Set the community read-only string.
Step 6	`Router(config)# snmp-server community userstring RW`	Set the community read-write string.
Step 7	`Router(config)# Ctrl+Z`	Return to Privileged Exec mode.
Step 8	`Router# copy running startup`	Save the configuration changes to NVRAM.

Tip The SNMP strings defined in the VPN Solutions Center for each target device must be identical with those configured for the corresponding edge devices in the service provider network. The procedure for setting the SNMP parameters in the VPN Solutions Center software is described in the "Specifying the Default SNMPv3 Attributes for PEs" section on page 4-22 and in the "Specifying the SNMPv3 Attributes for CEs" section on page 4-58.

Setting the SNMPv3 Parameters on the Routers in the Service Provider Network

Simple Network Management Protocol Version 3 (SNMPv3) is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting packets over the network.

This section describes how to set the SNMPv3 parameters on the routers in the service provider network. To complete the task regarding SNMPv3 parameters, you also must set a selected set of parameters in the VPN Solutions Center software. The SNMPv3 parameters you set on the routers must match the SNMPv3 parameters you specify in the VPN Solutions Center software (see the "Specifying the Default SNMPv3 Attributes for PEs" section on page 4-22 and in the "Specifying the SNMPv3 Attributes for CEs" section on page 4-58.).

The security features provided in SNMPv3 are as follows:

- Message integrity—Ensuring that a packet has not been tampered with in-transit.
- Authentication—Determining the message is from a valid source.
- Encryption—Scrambling the contents of a packet prevent it from being seen by an unauthorized source.

Using SNMPv3, data can be collected securely from SNMP devices without fear of the data being tampered with or corrupted. Also, using the **SNMP Set** command, packets that change a router's configuration can be encrypted to prevent its contents from being exposed on the network.

SNMPv3 provides for both security models and security levels. A *security model* is an authentication strategy that is set up for a user and the group in which the user resides. A *security level* is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

Three security models are available: SNMPv1, SNMPv2c, and SNMPv3. Table 2-1 identifies the combinations of security models and levels.

Table 2-1 SNMP Security Models and Levels

Model	Level	Authentication	Encryption	What Happens
v1	noAuthNoPriv	Community String	No	Uses a community string match for authentication
v2c	noAuthNoPriv	Community String	No	Uses a community string match for authentication.
v3	noAuthNoPriv	Username	No	Uses a username match for authentication.
v3	authNoPriv	MD5 or SHA	No	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms.
v3	authPriv	MD5 or SHA	DES	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Provides DES 56-bit encryption in addition to authentication based on the CBC-DES (DES-56) standard.

SNMPv3 objects have the following characteristics:

- Each user belongs to a group.
- A group defines the access policy for a set of users.
- An access policy is what SNMP objects can be accessed for reading, writing, and creating.
- A group determines the list of notifications its users can receive.
- A group also defines the security model and security level for its users.

To check the existing SNMP configuration, use these commands:

- **show snmp group**
- **show snmp user**

To set the SNMPv3 *server group* and *server users* parameters on a router, execute the following steps:

	Command	Description
Step 1	`> telnet router_name`	Telnet to the router you want to configure.
Step 2	`Router> enable` `Router> enable_password`	Enter enable mode, then enter the enable password.
Step 3	`Router# configure terminal`	Enter global configuration mode.
Step 4	`Router(config)# snmp-server group` [*groupname* {**v1** | **v2c** | **v3** {**auth** | **noauth** | **priv**}}] [**read** *readview*] [**write** *writeview*] [**notify** *notifyview*] [**access** *access-list*]	The **snmp-server group** command configures a new SNMP group or a table that maps SNMP users to SNMP views. Each group belongs to a specific security level. Example: **snmp-server group v3auth v3 auth read v1default write v1default**
Step 5	`Router(config)# snmp-server user` *username* [*groupname* **remote** *ip-address* [**udp-port** *port*] {**v1** | **v2c** | **v3** [**encrypted**] [**auth** {md5 | **sha**} *auth-password* [**priv des56** *priv-password*]] [**access** *access-list*]	The **snmp-server user** command configures a new user to an SNMP group. Example: **snmp-server user user1 v3auth v3 auth md5 user1Pass**
Step 6	`Router(config)# Ctrl+Z`	Return to Privileged Exec mode.
Step 7	`Router# copy running startup`	Save the configuration changes to NVRAM.

Enabling SA Agent on Edge Routers for SLA Jitter Probes

If you want to use the (voice) jitter protocol to collect SLA data from the edge devices in your network, you must enable SA Agent on each device from which you want to collect this data.

This procedure assumes that you have already enabled SNMP and set the SNMP parameters on the edge device router, as described in the previous sections of this chapter.

To enable SA Agent on an edge router for jitter probes, execute the following steps:

	Command	Description
Step 1	`> telnet router_name`	Telnet to the router you want to configure.
Step 2	`Router> enable` `Router> enable_password`	Enter enable mode, then enter the enable password.
Step 3	`Router# configure terminal`	Enter global configuration mode.
Step 4	`Router(config)# rtr responder`	Enable SA Agent on the SLA probe's target router.
Step 5	`Router(config)# Ctrl+Z`	Return to Privileged Exec mode.
Step 6	`Router# copy running startup`	Save the configuration changes to NVRAM.

Setting Up Various Elements in VPN Solutions Center

This section describes the elements or components you can set up on the VPN Solutions Center workstation.

Enabling TFTP in VPN Solutions Center

The VPN Solutions Center software in MPLS mode is set by default to use **TGS_TELNET** to transport configuration files to and from routers. To set VPN Solutions Center software to use TGS in TFTP mode instead, edit the *csm.properties* file as described below. Changing this value in the *csm.properties* file sets the default download mechanism for all the targets defined in the VPN Solutions Center Repository.

Note Setting the transport method in the VPN Console overrides the transport method setting in the *csm.properties* file. For instructions, see the "Default Transport Mechanism" section on page 4-14.

Step 1 On the VPN Solutions Center workstation, log in as the *vpnadm* administrative user.

Step 2 Go to the */<installation_directory>/vpnadm/vpn/etc* directory.

Step 3 Open the *csm.properties* file with a text editor.

Step 4 Find the following line in the *csm.properties* file:

```
netsys.gtl.transportMechanism = TGS_TELNET
```

Step 5 Change the `TGS_TELNET` value as follows:

```
netsys.gtl.transportMechanism = TGS_TFTP
```

Step 6 To specify the TFTP server IP address:

a. Find the **netsys.tgs.myTftpServer** property.

b. Specify the IP address of the VPN Solutions Center workstation.

Step 7 Save your changes and exit from the file.

Step 8 If the Watch Dog is running, be sure to stop the Watch Dog, then start it again to enable this change.

Step 9 Restart VPN Solutions Center.

When Using Multiple TGS Servers as TFTP Hosts

When using multiple Telnet Gateway servers as TFTP hosts, each Telnet Gateway server in each host system must have the **netsys.tgs.myTftpServer** property set to the IP address of the local host in which TGS is running.

This must be done whenever:

- The device's transport mechanism is set to **TGS_TFTP**, or
- An operator selects the **Download to Startup** option in either the Download Console (see the "Using the Task Manager" section on page 6-34) or the Template Manager.

In addition, if all the remote Telnet Gateway servers and local Telnet Gateway servers are symbolically linked, make sure they are not pointing directly to the same NFS mounted directory. Cisco recommends that */tftpboot* on each system that is running TGS should point to local directories on each TFTP host.

Setting a Local Solaris Host as a TFTP Server

This section describes how to set up a local Solaris host as a TFTP server. If the VPNSC Network Management Subnet includes one or more Telnet Gateway servers, you must set up the VPN Solutions Center workstation and the Telnet Gateway servers as TFTP hosts.

To set up a local Solaris host as a TFTP server:

Step 1 Open a new terminal window on the local host machine.

Step 2 Log in as the **root** user.

Step 3 With a text editor, open the *etc/inetd.conf* file.

Step 4 Find the following line in the file:

```
#tftp dgram udp wait root /usr/sbin/in.tftpd in.tftpd -s /tftpboot
```

Step 5 To activate the statement, delete the pound symbol (#) at the beginning of the line.

Step 6 Check to make sure that the command line indicated in Step 4 specifies the */tftpboot* directory. If any other directory is specified here, change it to **/tftpboot**.

Step 7 Save the change and exit from the *inetd.conf* file.

Step 8 Make sure that the */tftpboot* directory has read and write permissions for *user*, *group*, and *world*.

Step 9 At the terminal window command line, issue the following command to see if *inetd* processes are running:

```
ps -ef | grep inetd
```

The output of this command is as follows:

```
root <pid> 1 0 <date> <time> /usr/sbin/inetd -s
```

where *pid* is the process ID for the inetd process.

Step 10 If the process is running, send the SIGHUP (hang-up) signal to the inetd daemon with this command:

```
kill -1 <pid>
```

If the inetd daemon process is not running, start the inetd daemon.

Step 11 Exit from the terminal window.

Setting Up Connectivity to a Remote Telnet Gateway Server

When you install the VPN Solutions Center software on the VPNSC workstation, the installation includes a Telnet Gateway server (TGS). The VPN Solutions Center uses TGS for all communication with routers, including downloading and uploading configuration files (with the exception of SNMP communications, which are handled through the poller server).

Service providers can install multiple Telnet Gateway servers, either in the same network that VPN Solutions Center resides in, or on a remote network. However, installing the TGS servers on a remote network requires that TIBCO event connectivity between the VPNSC network and the remote network must be in place.

If you install multiple Telnet Gateway servers on the LAN connected to the VPN Solutions Center workstation (which is called the *VPNSC Network Management Subnet*), no special setup is required. However, if you want to install and use TGS on remote networks, the *TIBCO rvrd* software must be properly configured on both the VPN Solutions Center workstation and on one TGS machine in each remote network.

Note Even if a remote network contains multiple Telnet Gateway servers running on multiple machines, only one instance of TIBCO *rvrd* needs to run on that network.

Before You Begin the Setup Process

If VPN Solutions Center is currently running, you must bring it down before proceeding with the remote TGS setup procedure.

Step 1 Bring down VPN Solutions Center as described in the "Shutting Down the VPN Solutions Center Software" section on page 3-8.

Step 2 Check to see if the TIBCO software is already running:

```
ps -A | grep rv
```

Step 3 If TIBCO rvd or rvrd processes are running, kill them.

Step 4 Complete the TIBCO connectivity setup procedure on the VPN Solutions Center workstation and on the remote TGS machine as described in the following sections.

Setting Up the VPNSC Workstation for Connectivity to the Remote TGS Host

To set up the VPN Solutions Center workstation to allow TIBCO event connectivity to a TGS host in a remote network, follow these steps.

You must complete this procedure before you start the Watch Dog, bring up the VPN Solutions Center software, and start TGS on the VPN Solutions Center workstation.

Note On the VPN Solutions Center workstation, this is a one-time procedure. If you need to add additional remote TGS server machines, you do not need to repeat this procedure on the VPN Solutions Center workstation.

Step 1 On the VPN Solutions Center workstation, change directories to the */<installation_directory>/vpnadm/vpn* directory.

Step 2 Issue the following command to source the environment:

```
source vpnenv.csh
```

Step 3 Check to see if the TIBCO software is already running:

```
ps -A | grep rv
```

Step 4 If any **rvd** or *rvrd* processes are running, kill them.

Step 5 Issue the following command:

```
rvrd -store rvrd.store
```

Step 6 Start Netscape and go to the following URL, where "*VPNSC_hostname*" is the hostname of the VPN Solutions Center workstation:

http:*VPNSC_hostname***:7580**

The TIB/Rendezvous home page appears (see Figure 2-2).

Figure 2-2 TIBCO/Rendezvous Home Page

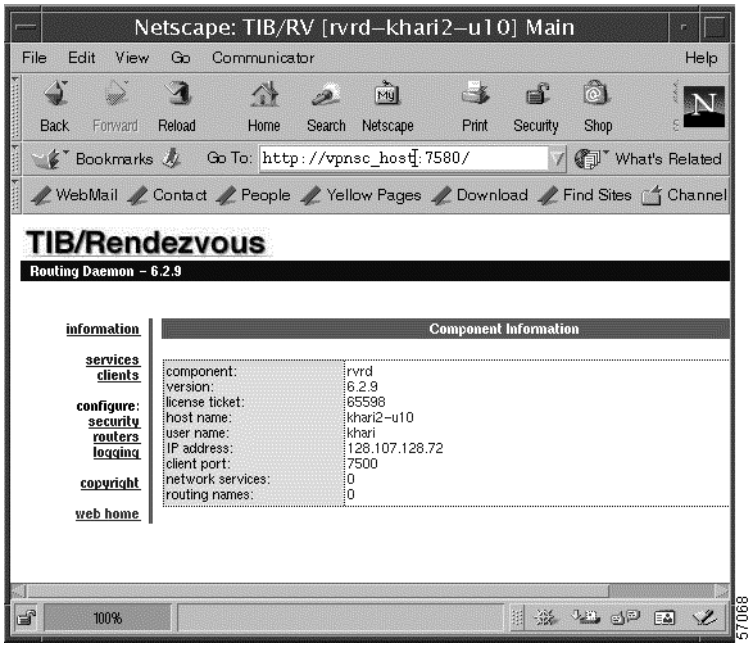

Step 7 From this page, choose the **routers** link.

The dialog box shown in Figure 2-3 appears.

Figure 2-3 Entering the VPNSC Host Name

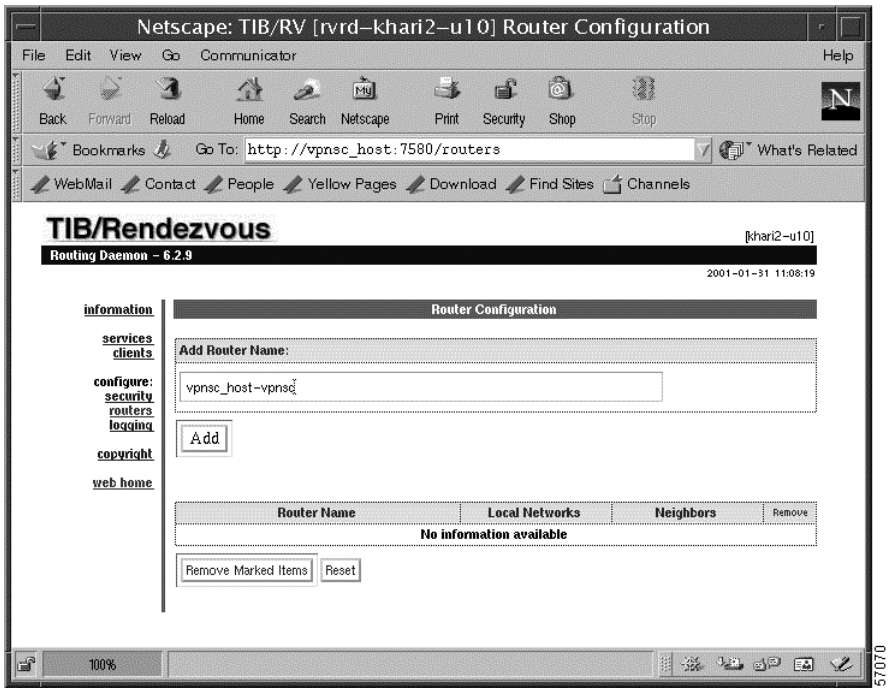

Step 8 In the *Add Router Name* field, enter the name of the VPN Solutions Center workstation followed by "**-vpnsc**," as follows: *VPNSC_host*-**vpnsc**.

Step 9 Click **Add**.

The value you entered is now displayed in the Router Name column.

Step 10 In the Local Networks column, select the current entry in the field.

The dialog box shown in Figure 2-4 appears.

Figure 2-4 Entering the VPNSC Local Network Information

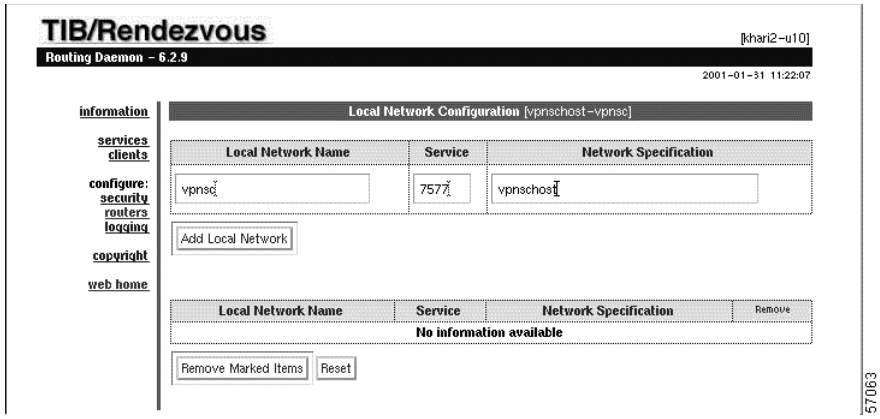

Step 11 Specify the local VPNSC network with the following values:

a. In the *Local Network Name* field, enter this value:

`vpnsc`

b. In the *Service* field, enter the TIBCO port number used for this VPN Solutions Center installation.

c. In the *Network Specification* field, enter the name of the VPNSC workstation.

Step 12 When the VPNSC network fields are specified, click **Add Local Network**.

On the lower section of the page, the values you entered are now displayed in the corresponding cells.

Step 13 From the current dialog box, choose the **routers** link.

Step 14 Click the current entry in the Neighbors column.

The dialog box shown in Figure 2-5 appears.

Figure 2-5 Entering the VPNSC Neighbor Information

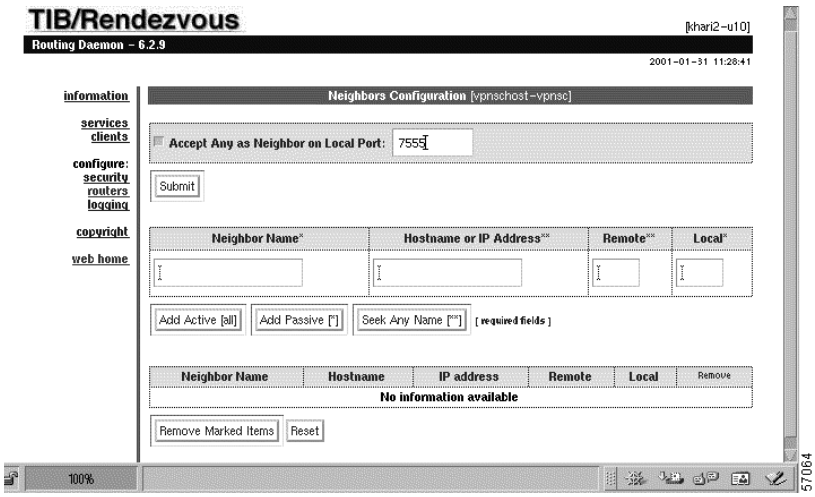

Step 15 Click the **Accept Any Neighbor on Local Port** option.

Step 16 In the *Local Port* option field, enter the following value:

```
7555
```

Step 17 Click **Submit**.

Step 18 From the current dialog box, choose the **routers** link.

Step 19 Click the current entry in Local Networks column.

The dialog box updates to the screen shown in Figure 2-4 on page 2-12. Notice that "**vpnsc**" is now displayed in the Local Network Name column.

Step 20 In the Local Network Name column, click the **vpnsc** entry.

The dialog box shown in Figure 2-6 appears.

Figure 2-6 Entering the VPNSC Neighbor Information

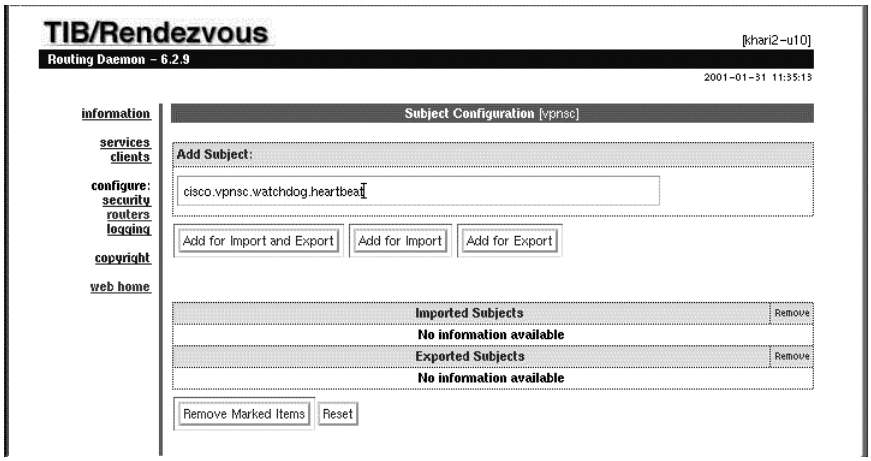

Step 21 In the *Add Subject* field, enter the following subject for import:

```
cisco.vpnsc.watchdog.heartbeat
```

Step 22 Click **Add for Import**.

The import subject you entered is now displayed in the *Imported Subjects* field.

This completes the procedure for setting up the for connectivity to the remote TGS host from the VPN Solutions Center workstation.

Enabling TIBCO Event Connectivity on the Remote TGS Host

To enable TIBCO event connectivity between a Telnet Gateway Server host on a remote network and the VPN Solutions Center workstation, follow these steps. This procedure assumes that TGS is installed on the Telnet Gateway Server host.

Note You must complete this procedure before you start TGS and before you start the VPN Solutions Center software.

In the following procedure, "*TGS_host*" refers to the hostname of the machine on which you are configuring the TIBCO rvrd software.

Step 1 On the remote Telnet Gateway Server host, change directories to the */<installation_directory>/vpnadm/vpn* directory.

Step 2 Issue the following command to source the environment:

```
source vpnenv.csh
```

Step 3 Check to see if the TIBCO software is already running:

```
ps -A | grep rv
```

Step 4 If any TIBCO *rvd* or *rvrd* processes are running, kill them.

Step 5 Issue the following command:

```
rvrd -store rvrd.store
```

Step 6 Start Netscape and go to the following URL, where "*TGS_hostname*" is the hostname of the Telnet Gateway Server installation:

http:*TGS_hostname***:7580**

The TIB/Rendezvous home page appears (see Figure 2-7).

Figure 2-7 TIBCO/Rendezvous Home Page

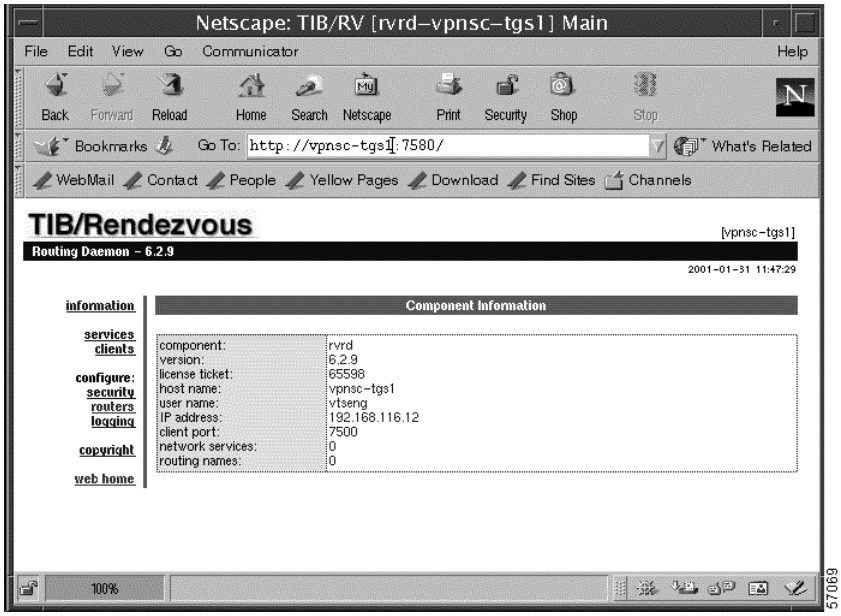

Step 7 From this page, choose the **routers** link.

The dialog box shown in Figure 2-8 appears.

Figure 2-8 Entering the TGS Host Name

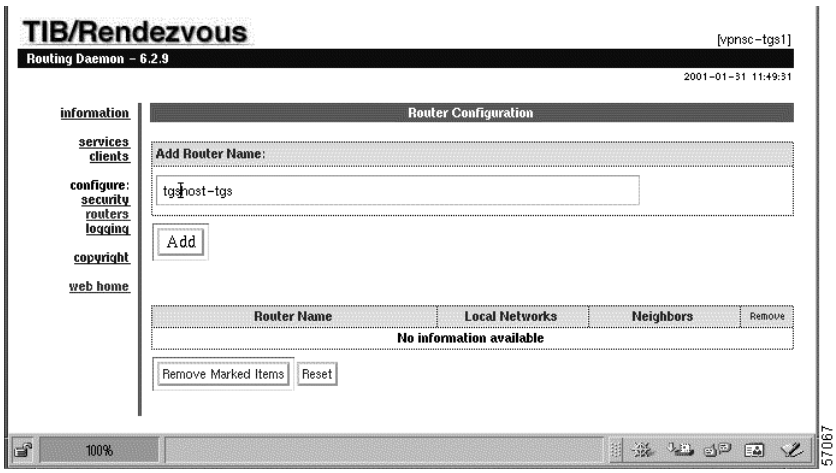

Step 8 In the *Add Router Name* field, enter the name of the TGS host followed by "**-tgs**," as follows: *TGS_host***-tgs**.

Step 9 Click **Add**.

The TGS host name is displayed in the Router Name column.

Step 10 In the *Local Networks* column, select the current entry in the field.

The dialog box shown in Figure 2-9 appears.

Figure 2-9 Entering the TGS Local Network Information

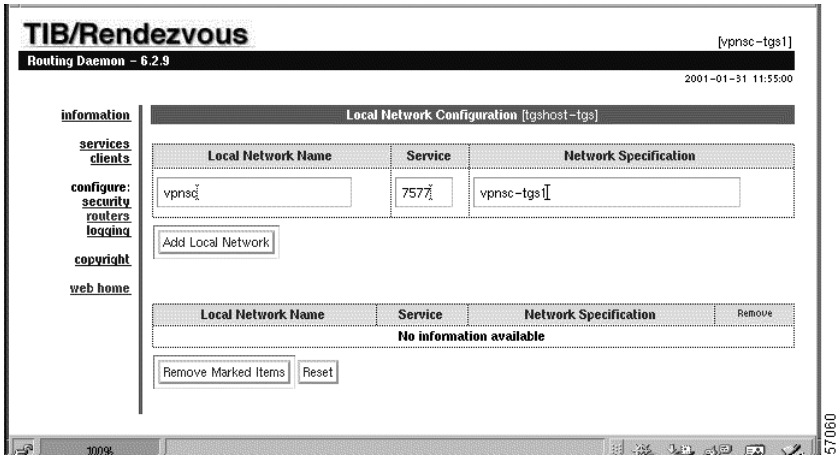

Step 11 Specify the local TGS network with the following values:

a. In the *Local Network Name* field, enter this value:

`vpnsc`

b. In the *Service* field, enter the TIBCO port number used for this installation.

The port number entered here should be same TIBCO port number entered in Step 11-b in the previous procedure to set up the VPNSC workstation for connectivity to the remote TGS host (see Figure 2-4 on page 2-12).

c. In the *Network Specification* field, enter the name of the TGS host.

Step 12 When the VPNSC network fields are specified, click **Add Local Network**.

On the lower section of the page, the values you entered are now displayed in the corresponding cells.

Step 13 From the current dialog box, choose the **routers** link.

Step 14 Click the currently displayed entry in the Neighbors column.

The dialog box shown in Figure 2-10 appears.

Figure 2-10 Entering the Telnet Gateway Server Neighbor Information

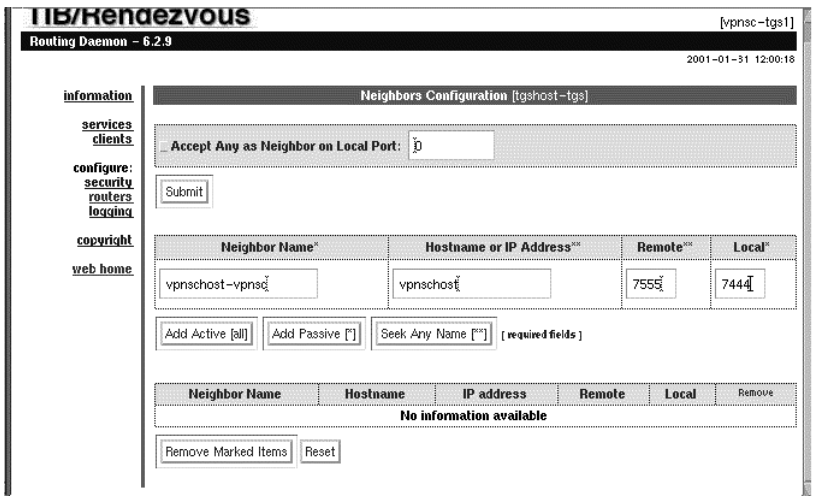

Step 15 Enter the TGS Neighbor information with these values:

a. In the *Neighbor Name* field, enter the name of the VPNSC workstation, followed by **-vpnsc**:

VPNSC_host-**vpnsc**

b. In the *Hostname or IP Address* field, enter the name of the VPNSC workstation.

c. In the *Remote* field, enter the following value:

7555

d. In the *Local* field, enter the following value:

7444

Step 16 Click **Add Active [all]**.

Step 17 From the current dialog box, choose the **routers** link.

Step 18 Click the currently displayed entry in Local Networks column.

Step 19 In the Local Network Name column, click the **vpnsc** entry.

The dialog box shown in Figure 2-11 appears.

Figure 2-11 Entering the Export Object Information

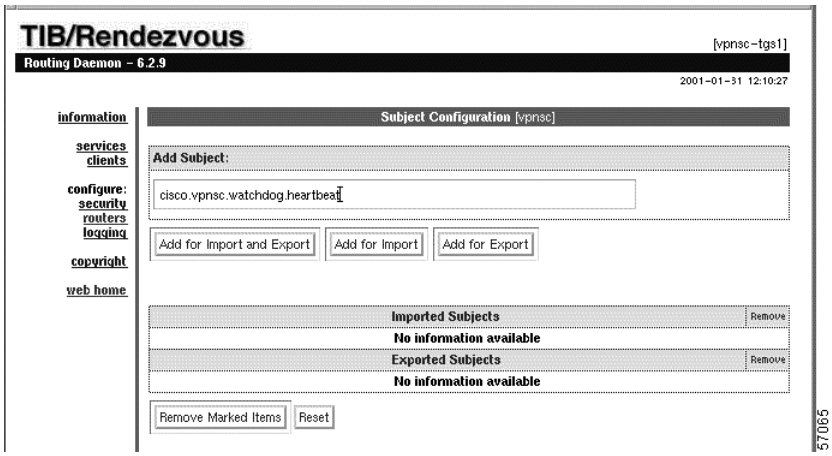

Step 20 In the *Add Subject* field, enter the following subject for export:

```
cisco.vpnsc.watchdog.heartbeat
```

Step 21 Click **Add for Export**.

The export subject you entered is now displayed in the *Exported Subjects* field.

This completes the procedure for setting up the connectivity to the VPNSC workstation on the remote TGS host.

Step 22 Start the VPN Solutions Center software as described in the "Starting the VPN Solutions Center Software" section on page 3-1.

Step 23 Start the TGS software as described in the "Starting the Telnet Gateway Server Software and the Watch Dog" section in Chapter 4 of the *Cisco VPN Solutions Center Installation Guide* (Release 2.2).

Specifying the TFTP Server Address for the TGS Host

If the TGS hosts in the VPNSC Network Management Subnet are also TFTP servers, you must set the appropriate property in the *csm.properties* file on those TGS hosts.

Step 1 On the local TGS host, log in as the *vpnadm* administrative user.

Step 2 Go to the */<installation_directory>/vpnadm/vpn/etc* directory.

Step 3 Open the *csm.properties* file with a text editor.

Step 4 Find the following line in the *csm.properties* file:

```
netsys.tgs.myTftpServer=
```

Step 5 Enter the IP address of the local TGS host.

Step 6 Save your changes and exit from the file.

Step 7 If the Watch Dog is running, be sure to stop the Watch Dog, then start it again to enable this change.

Step 8 Restart VPN Solutions Center.

Using the Cisco IE2100 with VPN Solutions Center

There are four major steps necessary to configure the VPNSC software and the IE2100 device so that they can communicate correctly:

- Modify Properties
- Configure the rvrd Daemon on the VPNSC Machine
- Configure the rvrd Daemon on the Cisco IE2100 Device
- Additional Setup Steps for Cisco IE2100 Devices

Modify Properties

There are two properties in the **csm.properties** file that must be modified to allow communication with Cisco IE2100 device. Set both properties as describes in the following steps. The **csm.properties** file is located in the **/<***install_dir***>/vpnadm/vpn/etc** directory, where <*install_dir*> is the directory in which the VPNSC software was installed.

Step 1 Set the value of the property **netsys.Cdm.IE2100.address** to the IP address of the Cisco IE2100 device.

For example:

```
netsys.Cdm.IE2100.address=12.34.56.78
```

Step 2 Set the value of the property **netsys.watchdog.startCdmServer** to be true (the default value of this property is false).

For example:

```
netsys.watchdog.startCdmServer=true
```

Configure the rvrd Daemon on the VPNSC Machine

VPN Solutions Center uses **rvd** by default. To use the IE2100 functionality, the VPNSC system needs to start the **rvrd** daemon. Do the following steps to set up and start the **rvrd** daemon on the VPNSC host machine.

Step 1 Log on as vpnadm.

```
su - vpnadm
```

Or if you are logging in remotely, enter this command:

```
rlogin <VPNSC_hostname> -l vpnadm
```

Step 2 Go to the directory where VPNSC is installed. For example:

```
cd /<installation_directory>/vpnadm/vpn
```

Step 3 Issue the following command to source the environment as required for your shell:

C-shell: `% source vpnenv.csh`

K-shell: `% . vpnenv.sh`

Step 4 Check to see if the TIBCO®/Rendezvous software is already running

```
ps -A | grep rv -
```

Step 5 If any **rvd** or **rvrd** processes are running, kill them.

Step 6 Enter the following command:

```
rvrd -store rvrd.store
```

Step 7 Start Netscape and go to the following URL, where <*VPNSC_hostname*> is the hostname of the VPN Solutions Center workstation:

```
% http://<VPNSC_hostname>:7580
```

The TIB/Rendezvous home page appears (see Figure 2-12).

Figure 2-12 TIBCO/Rendezvous Home Page

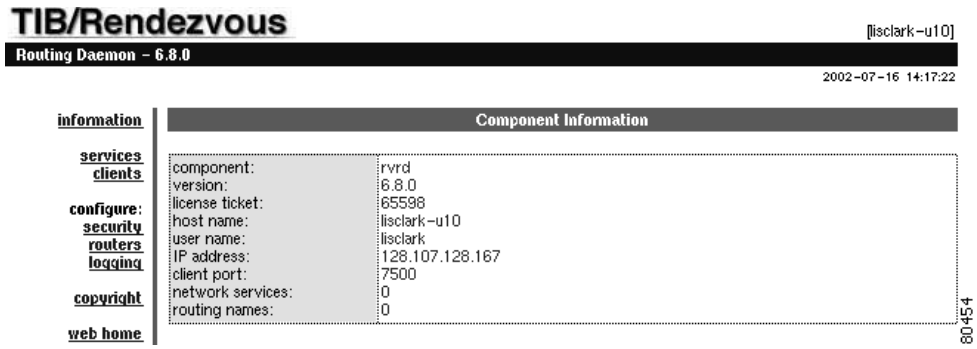

Step 8 From this page, choose the **routers** link.

The window shown in Figure 2-13 appears.

Figure 2-13 Entering the VPNSC Host Name

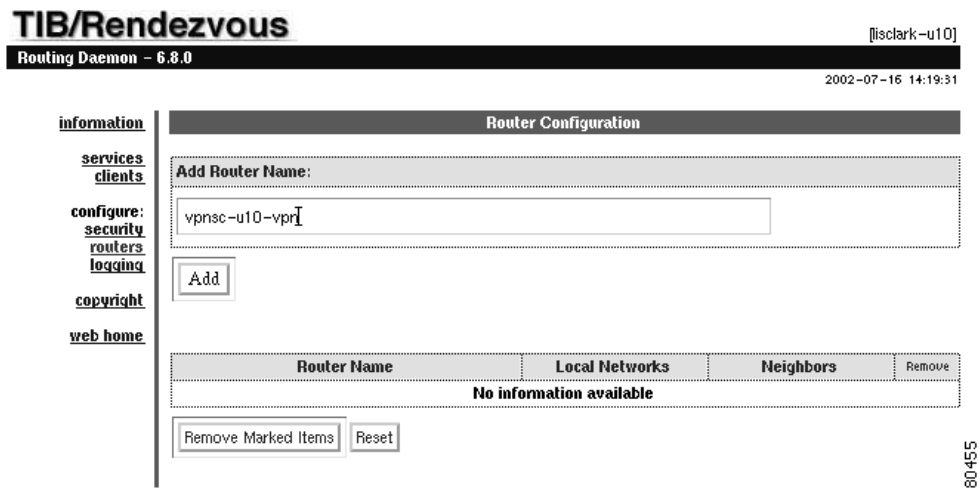

Step 9 In the *Add Router Name* field, enter the name of the VPNSC workstation followed by "**-vpn**," as follows: *<VPNSC_host>***-vpn**.

You do not have use the suffix "vpn". This is an example. Whatever you use, make sure that you enter a unique name.

Step 10 Click **Add**.

The value you entered is now displayed in the *Router Name* column, as shown in Figure 2-14.

Figure 2-14 After Adding Router Name

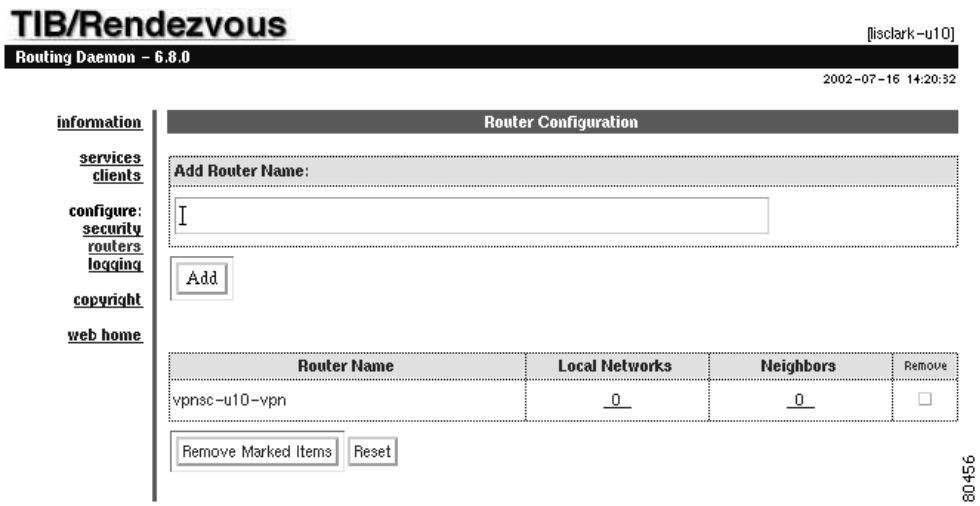

Step 11 In the *Local Networks* column, click the underlined entry in the column.

The window shown in Figure 2-15 appears.

Figure 2-15 Before Entering the VPNSC Local Network Information

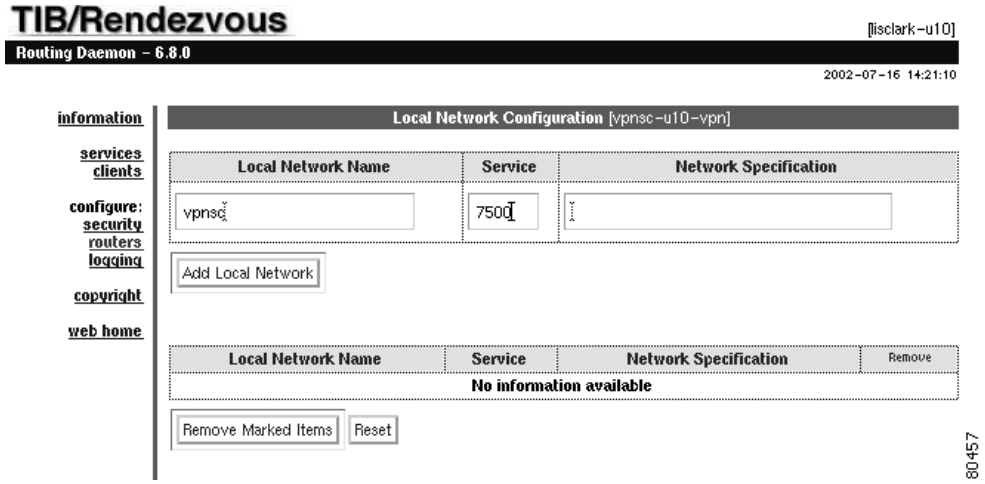

Step 12 Specify the local VPNSC network with the following values:

a. In the *Local Network Name* field, enter this value:

vpnsc

b. In the *Service* field, enter the TIBCO port number used for this VPNSC installation.

c. In the *Network Specification* field, enter the name of the VPNSC workstation.

Step 13 When the VPNSC network fields are specified, click **Add Local Network**.

On the lower section of the page, the values you entered are now displayed in the corresponding cells, as in Figure 2-16.

Figure 2-16 After Entering the VPNSC Local Network Information

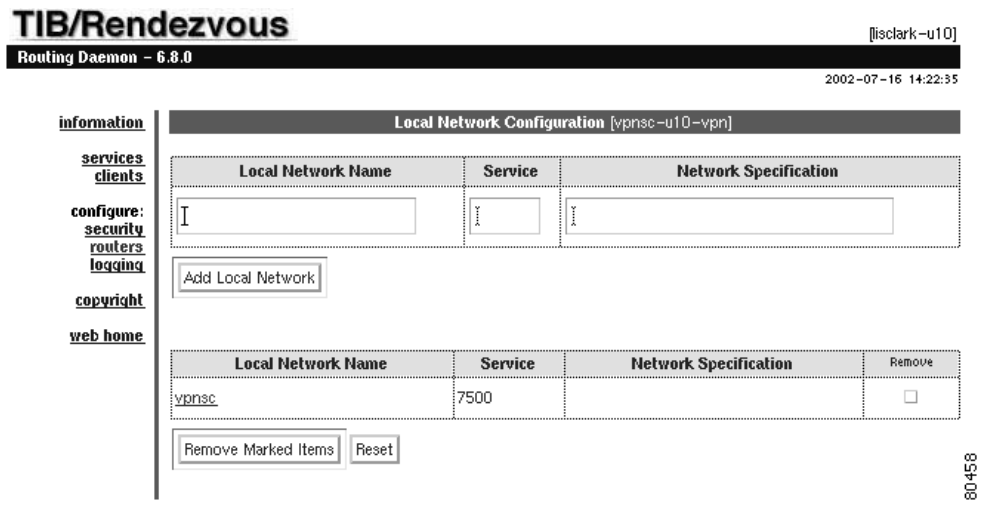

Step 14 From the current window, click the **vpnsc** entry in the *Local Network Name* column.

Step 15 In the *Add Subject* field, type the following (see Figure 2-17):

```
cisco.cns.>
```

Figure 2-17 Adding the Subject

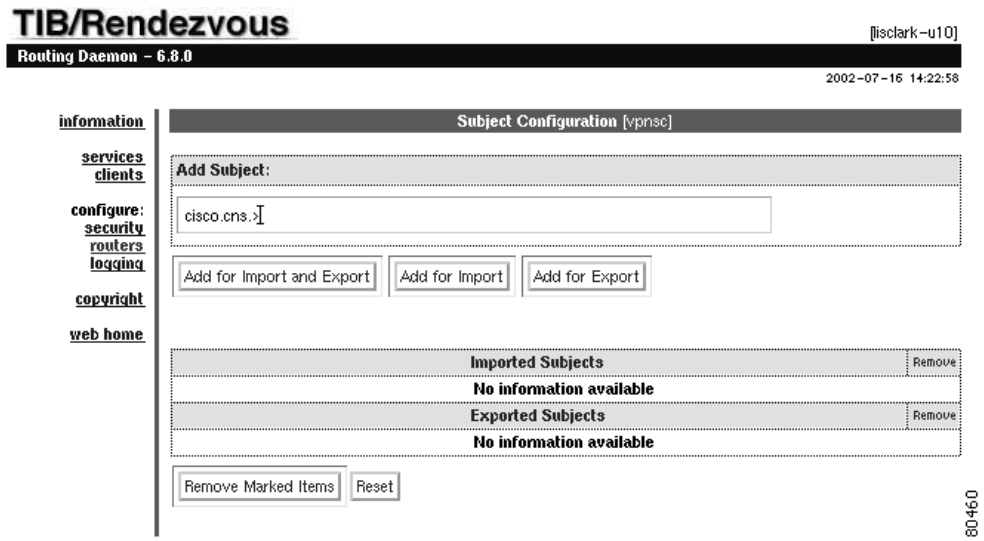

Step 16 Click **Add for Import and Export**.

Step 17 Choose the **routers** link again.

Step 18 In the Neighbors column, click the underlined entry in the column.

The window shown in Figure 2-18 appears.

Figure 2-18 Before Entering the VPNSC Neighbor Information

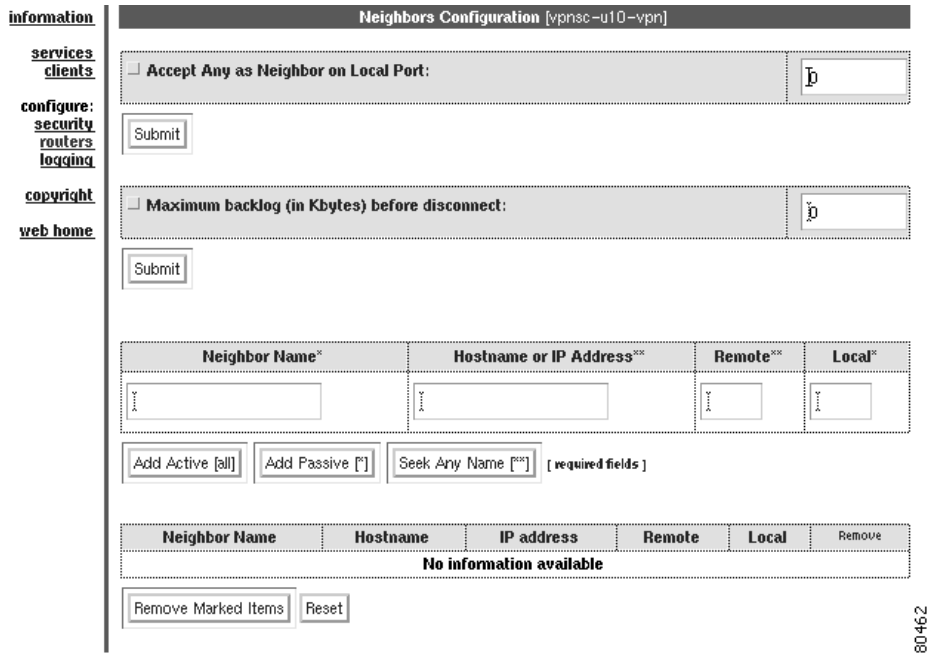

Step 19 In the *Neighbor Name* column, enter the name of the router as it is configured on the Cisco IE2100 device.

Step 20 In the *Hostname or IP Address* column, enter the IP address of the Cisco IE2100 device.

Step 21 In the *Remote* field, enter the following value:

7555

Step 22 In the *Local* field, enter the following value:

7555

Step 23 Click **Add Active**.

The dialog box shown in Figure 2-19 appears.

Figure 2-19 After Entering the VPNSC Neighbor Information

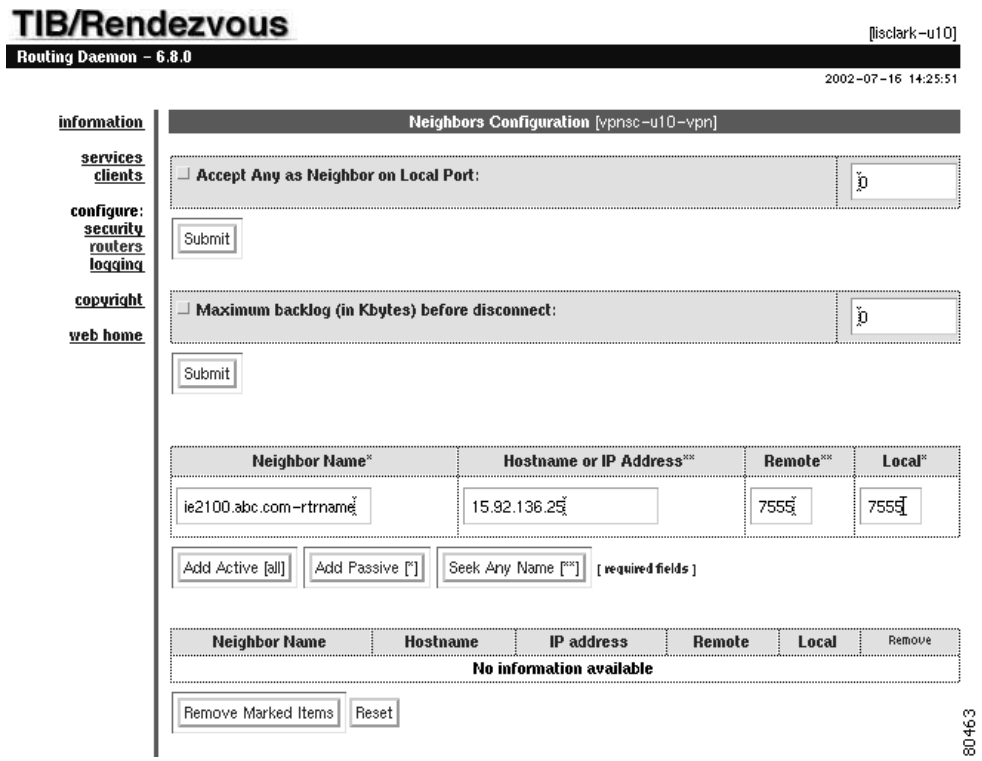

Configure the rvrd Daemon on the Cisco IE2100 Device

Configuring the rvrd daemon on the Cisco IE2100 device is similar to the steps in the previous section.

Step 1 If you closed Netscape already, start Netscape and go to the following URL, where <*IE2100_hostname*> is the hostname of the Cisco IE2100 device:

`http://<IE2100_hostname>:7580`

The TIB/Rendezvous home page appears (see Figure 2-12 on page 2-19).

Step 2 From this page, choose the **routers** link.

Step 3 In the *Add Router Name* field, enter the name of the IE2100 device followed by some suffix to make the name unique.

For example: <*IE2100_host*>**-acme_inc**

This is an example only. Whatever you use, make sure that you enter a unique name.

Step 4 Click **Add**.

The value you entered is now displayed in the *Router Name* column.

Step 5 In the *Local Networks* column, click the underlined entry in the column.

Step 6 Specify the local VPNSC network with the following values:

a. In the *Local Network Name* field, enter this value:

vpnsc

b. In the *Service* field, enter the TIBCO port number used for this VPNSC installation.

c. In the *Network Specification* field, enter the name of the IE2100 device or leave it blank.

Step 7 When the VPNSC network fields are specified, click **Add Local Network**.

On the lower section of the page, the values you entered are now displayed in the corresponding cells.

Step 8 From the current window, click on the **vpnsc** entry in the *Local Network Name* column.

Step 9 In the *Add Subject* field, type the following:

```
cisco.cns.>
```

Step 10 Click **Add for Import and Export**.

Step 11 Choose the **routers** link again.

Step 12 In the *Neighbors* column, click the underlined entry in the column.

Step 13 In the *Neighbor Name* column, enter the name of the VPNSC host machine.

Step 14 In the *Hostname or IP Address* column, enter the IP address of the VPNSC host machine.

Step 15 In the *Remote* field, enter the following value:

7555

Step 16 In the *Local* field, enter the following value:

7555

Step 17 Click **Add Active**.

A good indication that the connection is established is when the device in *Neighbor Name* appears as a hyperlink.

Additional Setup Steps for Cisco IE2100 Devices

The routers used with the IE2100 functionality (that is, using the CNS transport mechanism and/or set up with the Cisco Router Inactive type) need to run the IOS image 12.2(8)T or later.

Running Configuration

n addition, the router's *running configuration* must contain the following two CNS commands:

- **cns config partial <*IE2100 IP address*> 80**
- **cns event <*IE2100 IP address*> 11011**
  - Instead of the above command, you could use the following command instead:

    cns event <*IE2100 address*> **11011 keepalive** <*seconds*> <*retries*>
    where the **keepalive** option makes sure the TCP connection between the IE2100 device and the router is alive at all times by sending keepalive messages at intervals specified by **<***seconds***>**, with the number of retries specified by <*retries*>.

Startup Configuration

The router's *startup configuration* must contain the following two CNS commands:

- **cns config initial <***IE2100 address***> 80**

  The **cns config initial** command is only present the first time the router comes up on a network. It triggers the router to pick up and apply any initial configuration that might be waiting for it on the IE2100 device; then the command changes to **cns config partial**. There is a **no persist** option for this command that keeps the **cns config initial** in the startup configuration, which forces the router to contact the Cisco IE2100 for a new configuration every time it comes up.
- **cns event <*IE2100 address*> 11011** or **cns event <*IE2100 address*> 11011 keepalive <*seconds*> <*retries*>**

Please refer to the IE2100/CNS documentation for more details on other CNS commands and their options.

Modifying Frame Relay LMI Types

Local Management Interface (LMI) is a signalling standard between the router and the Frame Relay switch that provides a Frame Relay management mechanism. The LMI type must match the type used by the network. Changing the LMI type is a global change that affects all service requests (for related information, see the next section, "Applying a Mixed Set of LMI Types").

If a service provider or Customer needs to modify the Frame Relay Local LMI types, they can do so by modifying the appropriate property in the *csm.properties* file. Changing the LMI type in this way applies the Frame Relay modification to the Customer Edge router (CE) only.

You can set the LMI type to any one of four values:

LMI Value	Description
ansi	Annex D defined by ANSI standard T1.617
cisco	LMI type defined jointly by Cisco and other companies
none	This is the default.
q933a	ITU-T Q.933 Annex A

To modify the LMI type in the *csm.properties* file:

Step 1 On the VPN Solutions Center workstation, log in as the **vpnadm** user.

Step 2 Go to the */<installation_directory>/vpnadm/vpn/etc* directory.

Step 3 Open the *csm.properties* file with a text editor.

Step 4 Find the following line in the *csm.properties* file:

```
netsys.watchdog.server.CVPIMServer.frameRelayLmiType = none
```

Step 5 Change the **none** value to the appropriate LMI type value. For example, to change the LMI type to **cisco**, you would edit the line as follows:

```
netsys.watchdog.server.CVPIMServer.frameRelayLmiType = cisco
```

Step 6 Save your changes and exit the file.

Step 7 Log out (exit) from the **vpnadm** user.

Step 8 If the Watch Dog is running, be sure to stop the Watch Dog, then start it again to enable this change.

Step 9 Restart VPN Solutions Center.

Applying a Mixed Set of LMI Types

Changing the LMI type is a global change that affects all active service requests. To apply a mixed set of LMI types, do the following:

Step 1 Modify the *csm.properties* file to set the desired LMI type as described in the previous section.

Step 2 In the VPN Console, deploy the service requests that are associated with the LMI value set in Step 1.

Step 3 Modify the *csm.properties* file again to set the desired LMI type for the next set of service requests.

Step 4 In the VPN Console, deploy the service requests that are associated with the LMI value set in Step 3.

Enabling Telnet Sessions for Terminal Server Ports

You must enable at least as many Telnet sessions on the terminal server as there are terminal server ports. Otherwise, concurrent access to all the routers via the terminal server may fail.

To enable the appropriate number of Telnet sessions for terminal server access, follow these steps:

	Command	Description
Step 1	`> telnet terminal_server_name`	Telnet to the terminal server.
Step 2	`Terminalserver> enable` `Terminalserver> enable_password`	Enter enable mode, then enter the enable password.
Step 3	`Terminalserver# configure terminal`	Enter global configuration mode.
Step 4	`Terminalserver(config)# line vty 0 31`	Set the number of Telnet sessions to the number of available ports on the terminal server. This example sets 32 Telnet sessions.
Step 5	`Terminalserver(config)# Ctrl+Z`	Return to Privileged Exec mode.
Step 6	`Terminalserver# copy running startup`	Save the configuration changes to NVRAM.

Time Zone Support in VPNSC

VPN Solutions Center supports only the time zones that are in the */usr/share/lib/zoneinfo* directory of the Solaris workstation on which the VPN Solutions Center software is installed. The contents of this directory may change with each version of Solaris.

VPN Solutions Center cannot change the manner in which these time zones are configured, most notably the variations in Daylight Savings Time.

Note VPN Solutions Center does not support custom time zones.

CHAPTER 3

Starting and Stopping the VPN Solutions Center Software

Cisco VPN Solutions Center is a network service and management system that defines and monitors both MPLS-based and IPsec-based virtual private network (VPN) services for service providers. VPNSC allows service providers to seamlessly provision and manage intranet and extranet VPNs. The product provides the aspect of operations management that addresses flow-through provisioning, service auditing, and Service Level Agreement (SLA) measurement of IP-based VPN environments.

VPN Solutions Center focuses on provisioning, auditing, and monitoring the links between the customer's edge routers through the service provider's network.

Starting the VPN Solutions Center Software

Before you can start the VPN Solutions Center software, you must install the license key for the number of edge devices in your network (see the "Installing the VPN Solutions Center Licenses for Edge Devices" section on page 3-4). The license key is supplied in the *Right to Use Notification* document that is provided when you purchase VPN Solutions Center software.

Before you start the VPN Solutions Center software, complete these tasks:

Step 1 If you are bringing up VPN Solutions Center 2.2 for the first time:

- If your previous installation was VPNSC v1.x, you must convert your VPNSC v1.x Repository before proceeding. For instructions, see the "Converting a VPN Solutions Center 1.x Repository to 2.x Format" section on page 14-1.
- If your previous installation was VPNSC v2.0, you must convert your VPNSC v2.0 Repository before proceeding. See the "Converting a 2.0 Repository to 2.x Format" section on page 14-3.

Step 2 Log into the VPN Solutions Center (VPNSC) workstation under your own login name.

Step 3 *To keep the startup operations conveniently organized, open three terminal windows*—one window for the xhost process, one for the VPN Console and Watch Dog, and a third window for Orbix.

Step 4 In the first terminal window, enter the following command:

```
xhost + VPNSC_hostname
```

The *VPNSC_hostname* parameter is the name of the VPN Solutions Center workstation. This command configures your system so that the Orbix administrative user and the MPLS VPN administrative user (*vpnadm*) can communicate with your client system.

Starting Orbix

Starting the VPN Solutions Center software requires that you first start the Orbix process and then start the Watch Dog process and the VPN Console as described below. To start the Orbix software, follow these steps:

Step 1 Go to the terminal window for the Orbix software.

Step 2 Log in as administrative user of the Orbix process (*vpnadm*).

```
su - vpnadm
```

Or if you are logging in remotely, enter this command:

```
rlogin VPNSC_hostname -l vpnadm
```

Step 3 Go to the VPN Solutions Center installation directory.

```
cd /<install_dir>/vpnadm/vpn
```

Step 4 Issue the following command to source the environment as required for your shell:

C-shell: `source vpnenv.csh`

K-shell: `. vpnenv.sh`

Step 5 Start the Orbix process in the background:

```
startorbixd &
```

Starting the Watch Dog and the VPN Console

Step 1 Go to the terminal window for the Watch Dog and the VPN Console.

Step 2 Log in as the administrative user of the VPN Solutions Center software (*vpnadm).*

```
su - vpnadm
```

Or if you are logging in remotely, enter this command:

```
rlogin VPNSC_hostname -l vpnadm
```

Step 3 Go to the VPN Solutions Center installation directory.

```
cd /<install_dir>/vpnadm/vpn
```

Step 4 Issue the following command to source the environment as required for your shell:

C-shell: `source vpnenv.csh`

K-shell: `. vpnenv.sh`

Step 5 Set the display variable for the VPN Solutions Center workstation.

C-shell: **setenv DISPLAY** *VPNSC_hostname***:0.0**

K-shell: **export DISPLAY=***VPNSC_hostname***:0.0**

Step 6 Start the application's Watch Dog processes:

```
startwd
```

To stop the Watch Dog process, issue the **stopwd -y** command.

The Watch Dog log file is located at */opt/vpnadm/vpn/tmp/wdlog*.

Step 7 If you would like to confirm that the servers are running, issue the following command:

```
wdclient status
```

If you would prefer to bring up the Watch Dog graphical user interface, issue this command:

```
wdgui &
```

The Watch Dog interface appears (see Figure 3-1):

Figure 3-1 The VPN Solutions Center Watch Dog Interface

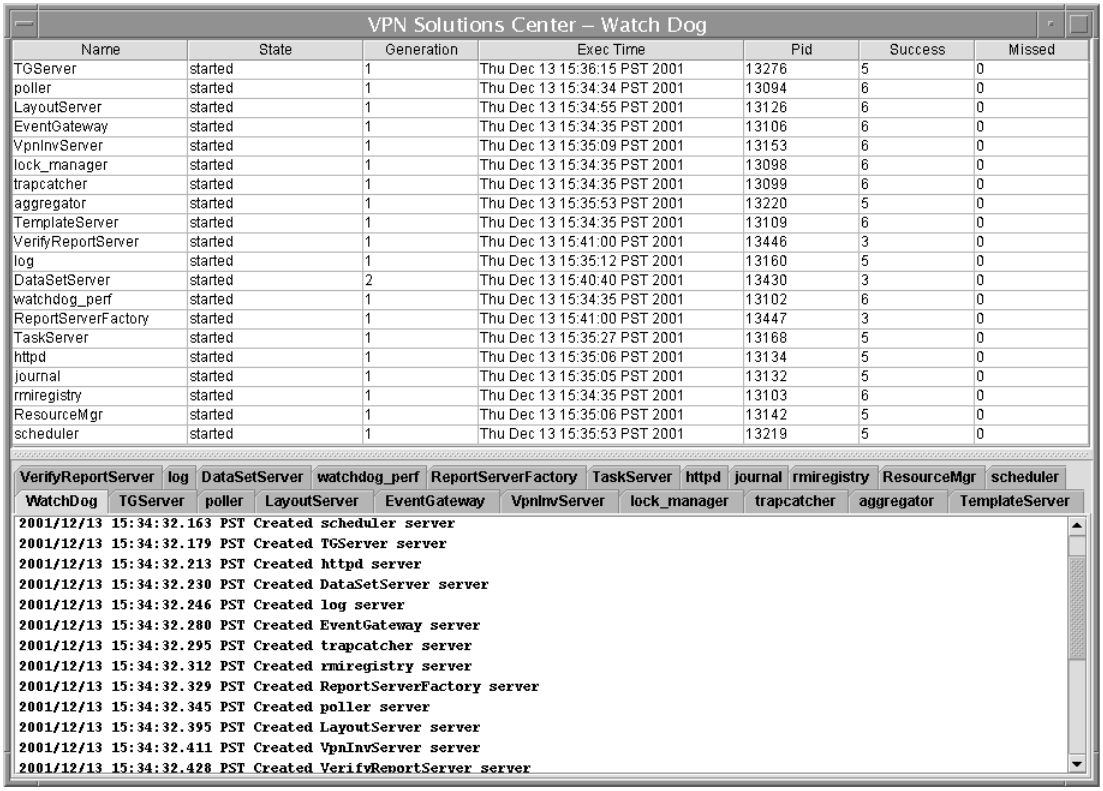

For a detailed description of the Watch Dog graphical user interface, refer to the "wdgui Command" section in Chapter 2, "Watch Dog Commands," of the *Cisco VPN Solutions Center: MPLS Solution User Reference*.

Step 8 Issue the following command to start the VPN Console:

```
vpnconsole -mode mpls &
```

The VPN Solutions Center Security dialog box appears (see Figure 3-2).

Figure 3-2 The VPN Solutions Center Security Dialog Box

Step 9 Enter a valid user name and password, then click **OK**.

The default username is **admin**. The default password is **admin**.

Note You cannot change the *admin* username, but you can change the password for the *admin* user. Additionally, you can add new usernames and assign their associated passwords.
The password for the VPN Console must be at least six (6) characters in length and no more than eight (8) characters in length. For details, see "User Administration" in Chapter 9, "VPN Console: Tools Menu," of the *Cisco VPN Solutions Center: MPLS Solution User Reference, Software Release 2.2*.

Installing the VPN Solutions Center Licenses for Edge Devices

An edge device is counted only once when it is added to a VPN as a CE or PE (in an MPLS network) or as an edge device router (in an IPsec network). If another VPN uses an edge device already counted in another VPN, it is not added again to the license count. Your license keys are supplied in the *Right to Use Notification* document.

The Install License dialog box (see Figure 3-7 on page 3-7) displays the number of edge devices you have a license for and the number of edge devices currently created. As you add edge devices to the VPNs defined in the VPN Solutions Center software, the number of edge devices created is incremented.

Checking the Number of Edge Devices Created

You can check the number anytime by choosing **Tools > License Administration** from the VPN Console menu bar and viewing the updated number of edge devices created.

When You Approach Your License Limit

When you reach 90 percent of the number of devices you are licensed for, VPN Solutions Center sends you a message to alert you to that situation. Cisco Systems strongly recommends that you take measures to upgrade your license agreement at that time.

Note When the number of edge routers created equals the number of edge routers licensed, you receive an email message to inform you. You cannot add any additional edge routers until you purchase an upgraded license and install the license key as described in the "Installing Upgrade Licenses" section on page 3-6.

Cisco recommends that you install the license keys in the following order:

- Application license key for MPLS
- Application license key for IPsec (if employing both VPN technologies)
- VPN license key for the number of VPNs in your installation
- Base license key for the number of edge devices in your network
- API license key (if purchased)

After you log in, the Install MPLS License dialog box appears (see Figure 3-3).

Figure 3-3 Installing the MPLS License Key

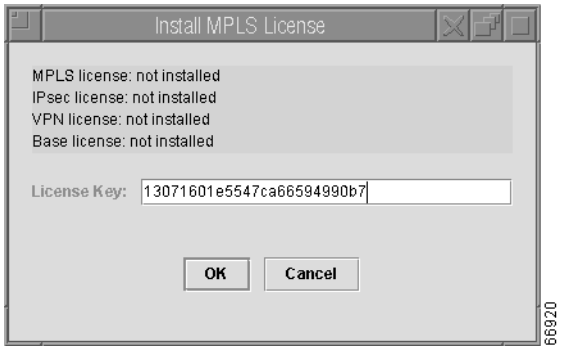

Step 1 Enter the MPLS License Key that you received with the VPN Solutions Center software, then click **OK**. You receive the message, "License key <*key number*> was installed successfully." Click **OK**.
The Install VPN License dialog box appears (see Figure 3-4).

Figure 3-4 Installing the VPN License Key

Step 2 Enter the VPN License Key that you received with the VPN Solutions Center software, then click **OK**. You receive the message, "License key <*key number*> was installed successfully." Click **OK**.
The Install Base License dialog box appears (see Figure 3-5).

Figure 3-5 Installing the Base License Key

Step 3 Enter the Base License Key that you received with the VPN Solutions Center software, then click **OK**.

You receive the message, "License key <*key number*> was installed successfully." Click **OK**.

The necessary licenses for your VPN Solutions Center installation are now successfully installed.

When all the VPNSC servers initialize, the VPN Solutions Center VPN Console appears, as illustrated in Figure 3-6.

Figure 3-6 The VPN Console: MPLS Solution

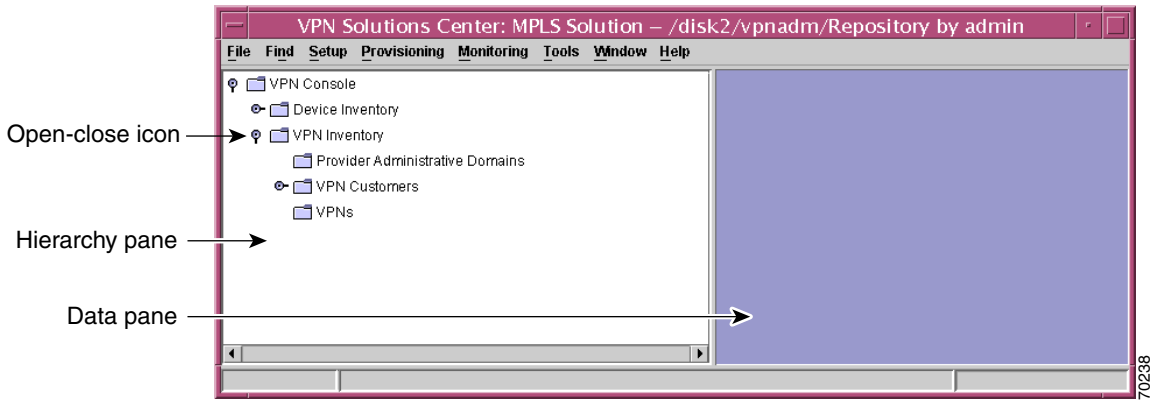

Proceed to the "Setting Up Networks in the VPN Solutions Center Software" section on page 4-3.

Installing Upgrade Licenses

When you purchase a new upgrade license—for example, to increase the number of licensed edge devices or add an API license—you can add the upgrade licenses to your VPNSC installation.

If you upgrade to more than one level of the number of edge devices licensed, you must:

- Enter the license key for each numerical increase.
- Enter the license keys in ascending order.

For example, if you were to upgrade from a license for 500 edge devices to a license for 3,000 edge devices, you would receive two Right to Use Notification documents: one document and license key for the upgrade from 500 to 1,500 edge devices, and another document and license key for the upgrade from 1,500 to 3,000 edge devices.

In this case, you would first enter the license key to upgrade from 500 to 1,500 edge devices, then repeat the procedure to upgrade from 1,500 to 3,000 edge devices.

To install VPNSC upgrade licenses, do the following:

Step 1 From the VPN Console menu bar, choose **Tools > License Administration**.

The Install License dialog box appears (see Figure 3-7).

Figure 3-7 Installing an Upgrade License Key

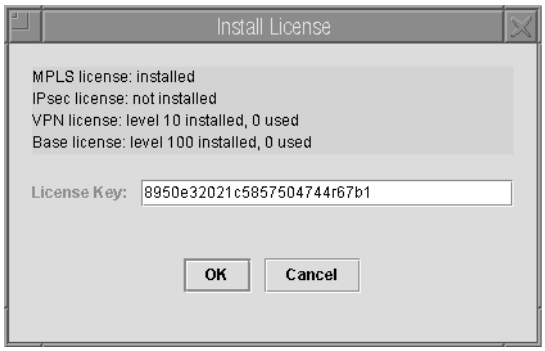

Step 2 Enter the license key provided in the Right to Use Notification document you received when you purchased the upgrade license, then click **OK**.

You receive the message, "License key <*key number*> was installed successfully." Click **OK**.

Step 3 If you need to add another level of upgrade, repeat Steps 1 and 2.

You can confirm and view the installed licenses by choosing **Tools > License Administration**. The new license level is displayed in the list of installed license keys.

Shutting Down the VPN Solutions Center Software

This section assumes that the VPN Solutions Center software is running, and that the administrative user name—*vpnadm*—is active. It also assumes that Orbix is running as a background process. To shut down the VPN Solutions Center software, execute these commands:

Step 1 If the VPN Console is running, close it by choosing **File > Exit**.

Step 2 If the Watch Dog user interface (**wdgui**) is running, close it by selecting the window, **right-click**, then choose **Close** from the menu.

Step 3 From the window where Watch Dog was launched, close the Watch Dog by issuing this command:

```
stopwd -y
```

Step 4 Log out (exit) from the *vpnadm* user.

Shutting down Orbix is optional. To shut down Orbix, follow these steps:

Step 5 From the terminal window from which you launched Orbix, shut down the Name Server:

```
killit NS
```

Step 6 Discover the process ID of orbixd:

```
ps -ef | grep orbixd
```

Step 7 Shut down the Orbix process by issuing this command:

kill *orbixd_process_ID*

Step 8 Log out (exit) from the *vpnadm* user.

Setting Up Networks and Importing Configurations Into VPN Solutions Center

The main topics presented in this chapter are as follows:

VPN Solutions Center in a Service Provider MPLS Environment

In an MPLS network, a customer edge router (CE) is connected to a provider edge router (PE) in such a way that the customer's traffic is encapsulated and transparently sent to other CEs, thus creating a virtual private network. The VPN Solutions Center provisioning engine for MPLS accesses the configuration files on both the CE and PE to compute the necessary changes to the configuration files to support the service on the PE-CE link.

Figure 4-1 VPN Solutions Center: MPLS Solution in the Service Provider Network

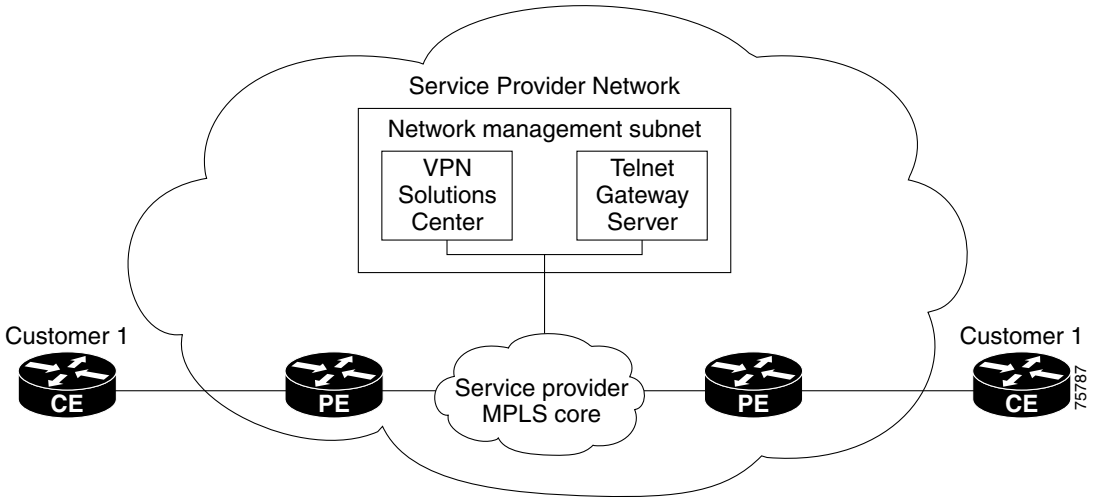

As illustrated in Figure 4-1, Cisco requires that the VPN Solutions Center software is installed on its own dedicated system. The VPN Solutions Center workstation is connected on a LAN to one or more Telnet Gateway servers.

Network Administrators and Network Operators

Cisco assumes that there are at least two categories of users in VPN Solutions Center:

- **Network Administrators**, who do all the setup tasks, including:
  - Setting up the networks as described in the next section.
  - Using the Import Manager to define the PEs, CEs, and other devices in the provider networks.
  - Using the Import Manager to import device configuration files into VPNSC, as described in the "Importing PEs and CEs into VPN Solutions Center" section on page 4-3.
  - Defining VPNs in VPNSC, which is described in Chapter 5, "Creating MPLS VPNs and Administering Service Request Profiles."
  - Using the MPLS Service Request Editor to create and organize Service Request Profiles. For details, see Chapter 5, "Creating MPLS VPNs and Administering Service Request Profiles."
- **Network Operators**, who use the MPLS Service Request Editor to provision service requests.

  For details, see Chapter 6, "Provisioning MPLS VPN Service Requests."

Setting Up Networks in the VPN Solutions Center Software

In this product, an MPLS VPN *network* is a unique group of *targets*; a target can be a member of only one network. Thus, an MPLS VPN network allows a provider to partition the working space into manageable segments that are unique and do not overlap other networks.

To use VPN Solutions Center to set up MPLS VPN networks, complete the following tasks:

1. Import the target PE and CE routers—see the next section, "Importing PEs and CEs into VPN Solutions Center."
2. Create the VPN customer definitions for each VPN customer—see the "Defining a New VPN Customer Name" section on page 4-44 and the "Specifying a Customer Site for Each CE" section on page 4-46.
3. Define the Provider Administrative Domain(s)—see the "Defining Provider Administrative Domains" section on page 4-30.
4. Define the VPNs—see the "Defining a New VPN in the VPNSC Software" section on page 5-2.
5. If you are using a management VPN to manage your customers' VPNs, define a management VPN—see the "Implementing the Management VPN Technique" section on page 8-12.

Caution Make sure that the file descriptor limit is *not* set in the VPN Solutions Center workstation login shell file (which can be the *.login* file, the *.cshrc* file, or the *.kshrc* file). If the login shell file contains a line with the **ulimit -n** command (for example, "`ulimit -n <number>`"), comment out this command line in the file.

VPN Solutions Center cannot override the file descriptor limitation setting in the login shell file. If the value is set incorrectly, VPN Solutions Center experiences operational problems.

Importing PEs and CEs into VPN Solutions Center

Every device that the VPN Solutions Center software manages must be defined as a *target*. A target is any device from which the VPN Solutions Center software can collect information (a router or NetFlow Collector). In most cases, these targets are Cisco routers that function either as a provider edge router (PE) or a customer edge router (CE).

Tip When you define target names in the VPN Solutions Center software, the target names you specify must match the actual IOS host names of the corresponding devices.

The Simple Network Management Protocol (SNMP) must be configured on each PE router and CE router in the service provider network. To determine whether SNMP is enabled and set the SNMP community strings on a router, see the "Setting Up SNMPv1 and SNMPv2 on the Routers in the Service Provider Network" section on page 2-4 and the "Setting the SNMPv3 Parameters on the Routers in the Service Provider Network" section on page 2-5.

The VPNSC Import Manager

The principal method for defining targets and organizing them into the appropriate networks (or target groups) is to use the Import Manager to import all the pertinent Provider Edge router and Customer Edge router configuration files.

Note A *valid* configuration file is one in which the *hostname statement* is present in the file. If a configuration file does not contain the hostname statement, VPN Solutions Center software regards the file as invalid and does not import the configuration file into the Repository.

Importing router configuration files into VPN Solutions Center is an efficient way to define the MPLS VPN networks and the devices in them. Using the VPNSC Import Manager, you define the attributes for multiple PEs and CEs at once prior to actually importing the configuration files into VPN Solutions Center. VPNSC creates the network and the devices in the network based on the imported configuration files. When you have completed defining the attributes for the device configuration files and then import the configuration files into VPN Solutions Center, network operators can immediately begin provisioning VPNs and service requests.

Tip The Import Manager is designed to facilitate the process of setting up many devices at once so that the Network Administrators and Network Operators can quickly provision all the PEs and CEs in the service provider administrative domain.

When you reimport a configuration file in VPN Solutions Center that was imported previously (that is, a device that already exists in VPNSC), importing that configuration file does not create a new target—it updates the target entry in the Repository and adds a new dataset for the existing target.

Note In VPN Solutions Center 2.2, you must import PEs first, then import CEs.

Importing Provider Edge Routers into VPN Solutions Center

To import PE router configuration files into VPN Solutions Center, follow these steps:

Step 1 Create a directory of configuration files for a given set of devices and copy the appropriate configuration files into the directory.

Device names within each directory must be unique.

A typical set includes Provider and Customer edge routers (PEs and CEs).

Step 2 From the VPN Console menu, choose **Setup > Create Targets From Configuration Files**.

The opening window for the Import Manager appears (see Figure 4-2).

Figure 4-2 Import Manager: Opening Screen

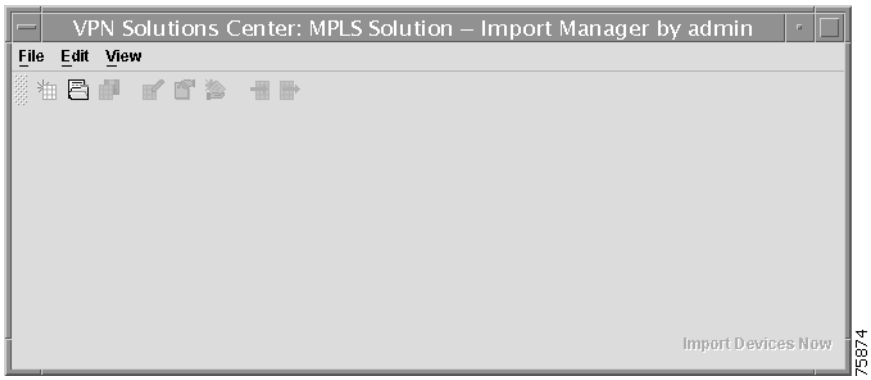

You must first import the configuration files for PEs.

Step 3 From the Import Manager menu bar, choose **File > Import > Import Provider Edge Devices** (see Figure 4-3).

Figure 4-3 Import Provider Edge Devices Menu Option

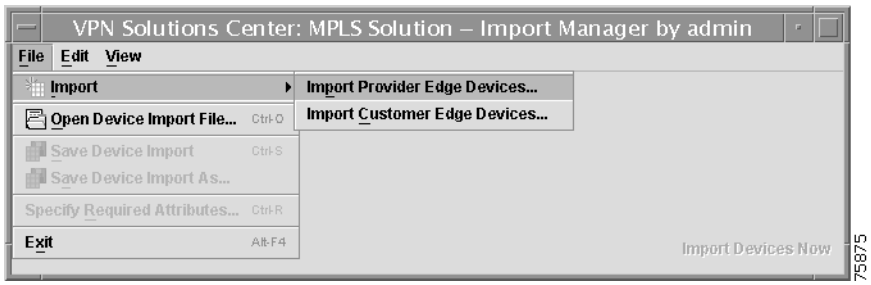

A dialog box appears that lets you change directories to where a set of Cisco IOS configuration files are located (see Figure 4-4).

Figure 4-4 Selecting Provider Edge Router Configuration Files

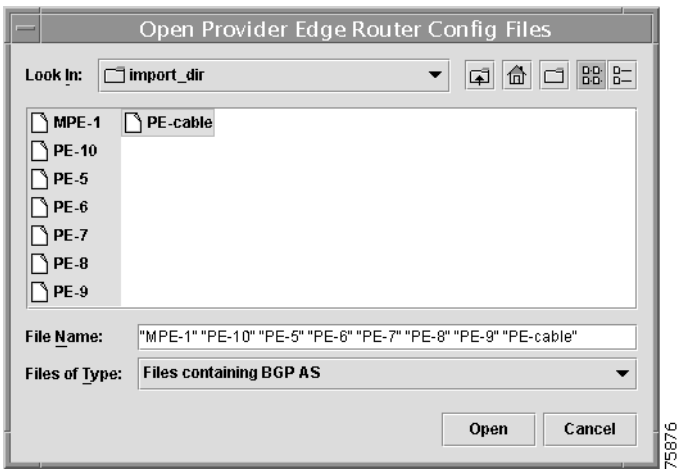

Step 4 Use the Open Provider Edge Router Config Files dialog box to locate and specify the PE router configuration files that you want to import into VPN Solutions Center.

a. *Look In*: Navigate to the directory where the PE configuration files reside.

b. *Files of Type*: When importing PE files, the Import Manager by default lists **Files containing BGP AS**.

Valid Provider Edge Router configuration files include the BGP AS (Border Gateway Protocol Autonomous System) number. With **Files containing BGP AS** set for the file type, the Import Manager displays PE configuration files only.

The BGP AS number corresponds to a Provider Administrative Domain (PAD). When VPNSC finds a BGP AS number in a configuration file, it searches the Repository for a PAD with a matching BGP AS number. If a match is found, the PAD is automatically assigned to the imported PE. If a match is not found, the Import Manager prompts you to create a PAD for the PE.

In addition, the *Fields of Type* field provides two other file type options:

- **All Files**
- **Files not containing BGP AS**

c. *File Name*: When you select PEs from the displayed list, their filenames are displayed. You can also type in the pertinent PE filenames if you wish.

d. When you have selected the PE configuration files you want to import, click **Open**.

The Import Manager displays a dialog box in spreadsheet format that shows the list of PEs you selected for import (see Figure 4-5).

Figure 4-5 PEs Displayed in Import Manager Spreadsheet

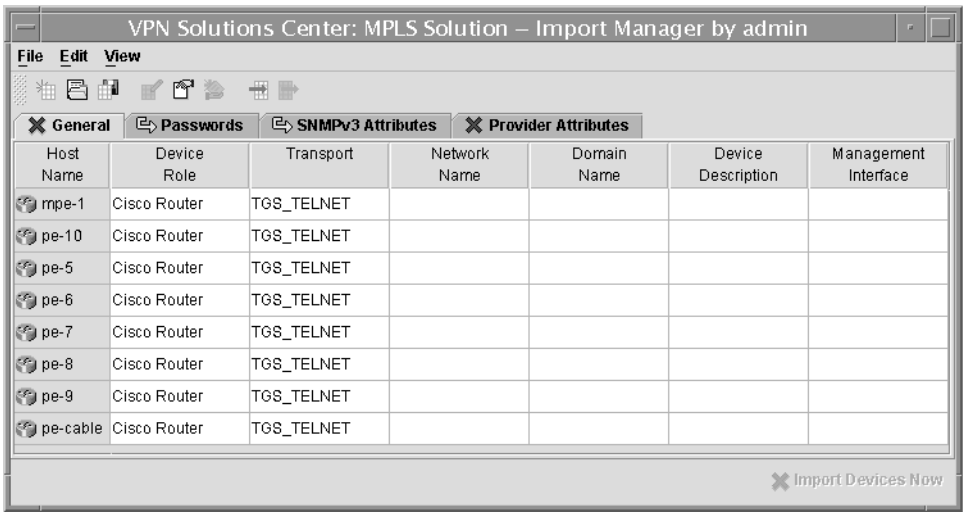

Observe the Import Manager dialog box as shown in Figure 4-5:

- Two of the tabs—**General** and **Provider Attributes**—have red X's in the tab title. The red X's indicate that some or all of the parameters in the spreadsheets for those tabs must be filled in before you can import the devices into VPN Solutions Center. When these values are filled in, the red X's change into green arrows, indicating that device description for that spreadsheet is complete and ready for import.

- The other two tabs—**Passwords** and **SNMPv3 Attributes**—display yellow arrows. This indicates that although there are some values that are not defined in those areas, they are considered ready for import.
- The **Import Devices Now** button (in the lower right corner) is temporarily disabled. When the required parameters for the PEs are specified, this button is enabled so that you can then import the PE configuration files into VPN Solutions Center.
- The *Device Role* for each device listed is already set to **Cisco Router**, and the *Transport Mechanism* for each PE router is set by default to **TGS_TELNET**. You can change the transport mechanism for all or some of the devices as necessary.

Specifying the Required Attributes for PEs

Most of the attributes for PEs are optional by default. However, you can specify which attributes are required for provisioning.

To specify the required attributes for PEs, follow these steps:

Step 1 From the Import Manager menu bar, choose **File > Specify Required Attributes**.

The Specify Required Attributes editor for MPLS PE devices appears (see Figure 4-6).

Figure 4-6 Specify Required Attributes Editor

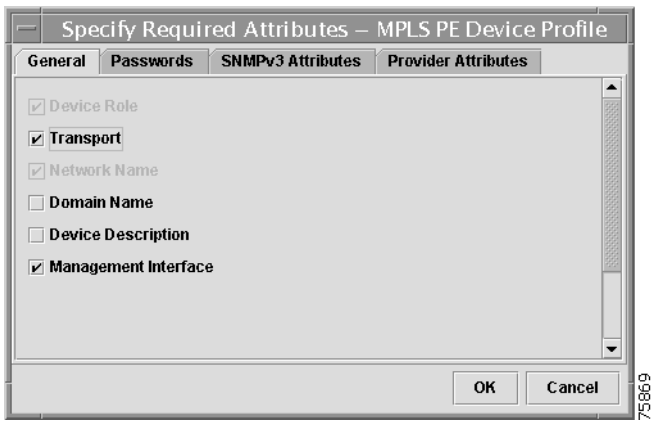

The Specify Required Attributes editor (for PEs) is organized into four tabs:

- *General*

  As shown in Figure 4-6, the General tab includes the following attributes: **Device Role**, **Transport**, **Network Name**, **Domain Name**, **Device Description**, and **Management Interface**.

  Only **Device Role** and **Network Name** are required by default. In this example, we show **Transport** and **Management Interface** selected to be required attributes.
- *Passwords*
- *SNMPv3 Attributes*
- *Provider Attributes*

The Provider Attributes—**Provider Administrative Domain** and **Region**—are both required by default.

Step 2 In the list of attributes in the General tab, select the checkboxes for those attributes that you want to be required.

Step 3 Click the **Passwords** tab.

The Specify Required Password Attributes dialog box appears (see Figure 4-7).

Figure 4-7 Specify Required Passwords Attributes

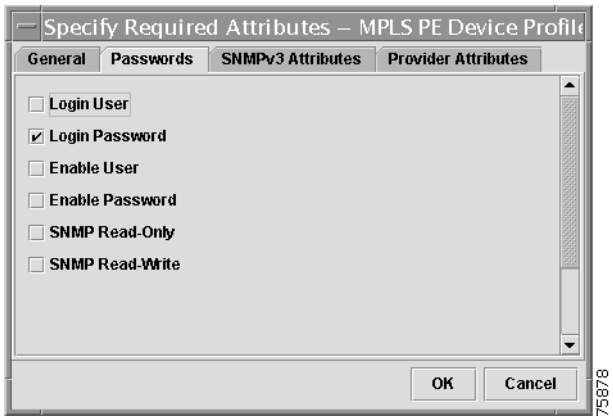

Step 4 In the list of attributes in the Passwords tab, select the checkboxes for those password attributes that you want to be required.

Step 5 Click the **SNMPv3 Attributes** tab.

The Specify Required SNMPv3 Attributes dialog box appears (see Figure 4-8).

Figure 4-8 Specify Required SNMPv3 Attributes

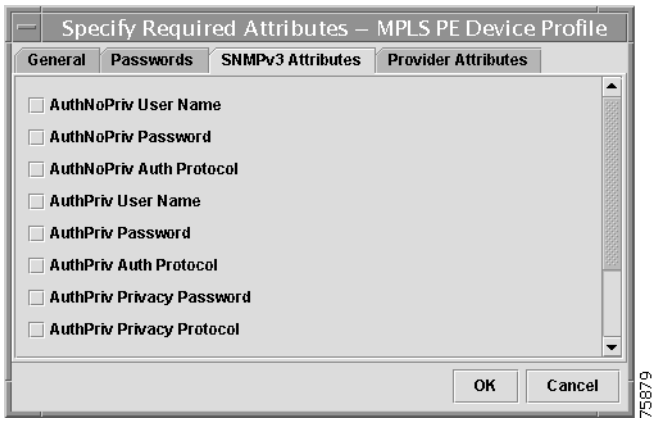

Step 6 In the list of attributes in the SNMPv3 Attributes tab, select the checkboxes for those SNMPv3 attributes that you want to be required.

Step 7 When satisfied with the settings for required PE attributes, click **OK**.

VPN Solutions Center will now require the selected required attributes to be defined in the Import Manager before the set of PE routers can be imported into VPNSC.

Two Ways to Proceed: Edit Parameters or Set Default Values

Once you bring in the desired router configuration files into the Import Manager and you define the required attributes for those devices, you have two ways you can proceed:

- You can enter the values for each device parameter displayed in the Import Manager (as described in the next section, "Setting Parameter Values for the Devices to be Imported" section on page 4-11).
- You can specify the default values for the devices' parameters and attributes and then apply those default values to the set of imported devices (as described in the "Specifying the Default Values for the Imported PE Routers" section on page 4-13).

  Specifying default values is not a requirement—it is provided as a convenience to the Network Administrator. You always have the option of defining the parameters and attributes of the selected devices without specifying default values.

Saving a Device Import Instance to a File

In the event that you are not able to complete the entire device import procedure at one session, you can always save the work you have done so far to a device import file. You can then open that device import file and proceed from the point you left off.

To save the device import to a file, follow these steps:

Step 1 From the Import Manager dialog box, choose **File > Save Device Import**.

The Save dialog box appears (see Figure 4-9).

Figure 4-9 Saving Import Work to a File

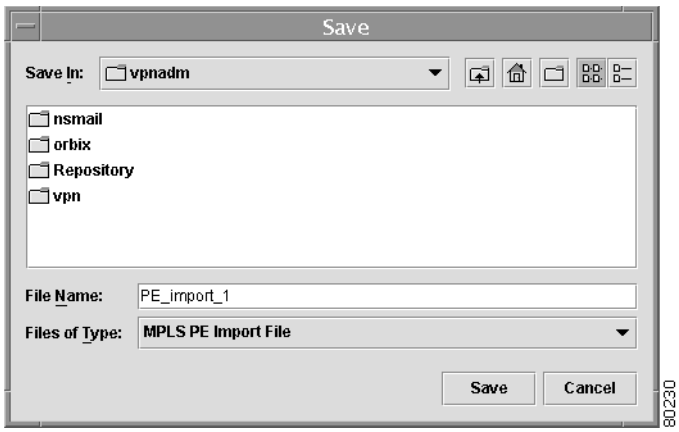

Step 2 *Save In*: Navigate to the location where you want to save the device import file.

Step 3 *File Name*: Enter the name of the device import file. Do not apply a filename extension—VPNSC automatically appends the appropriate extension.

The *Files of Type* are already set to **MPLS PE Import File**.

Step 4 When ready, click **Save**.

VPN Solutions Center saves the device import file to the location you specify.

Opening a Device Import File

To open a device import file, follow these steps:

Step 1 From the VPN Console menu, choose **Setup > Create Targets From Configuration Files**.

The opening window for the Import Manager appears (see Figure 4-10).

Figure 4-10 Import Manager: Opening Screen

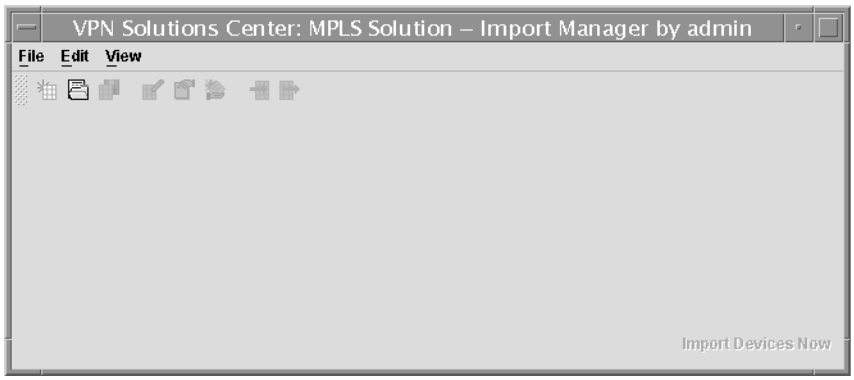

Step 2 From the Import Manager menu bar, choose **File > Open Device Import File**.

The Open dialog box appears (see Figure 4-11).

Figure 4-11 Opening a Device Import File

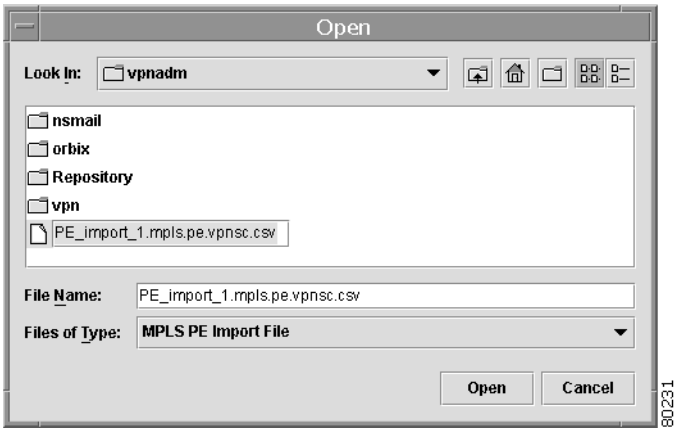

Step 3 *Look In*: Navigate to the location where you device import file resides.

Step 4 *File Name*: Select the name of the device import file. The filename is displayed in the *File Name* field.

The *Files of Type* are already set to **MPLS PE Import File**.

Step 5 When ready, click **Open**.

VPNSC opens the device import file. You can then resume work on the device import procedure.

Setting Parameter Values for the Devices to be Imported

The Import Manager provides several tools to allow you to define the parameters for multiple devices at once. In this section we describe how you can define a single parameter for all the devices you are importing into VPN Solutions Center. In this example, we specify the management interface for all the PEs that are to be imported.

To define a single parameter for all the devices you are importing, follow these steps:

Step 1 From the Import Manager, select the cells in the column for the device parameter you want to define.

Step 2 Choose **Edit > Edit Selected Devices** (see Figure 4-12).

Figure 4-12 Setting Attribute for Multiple Devices

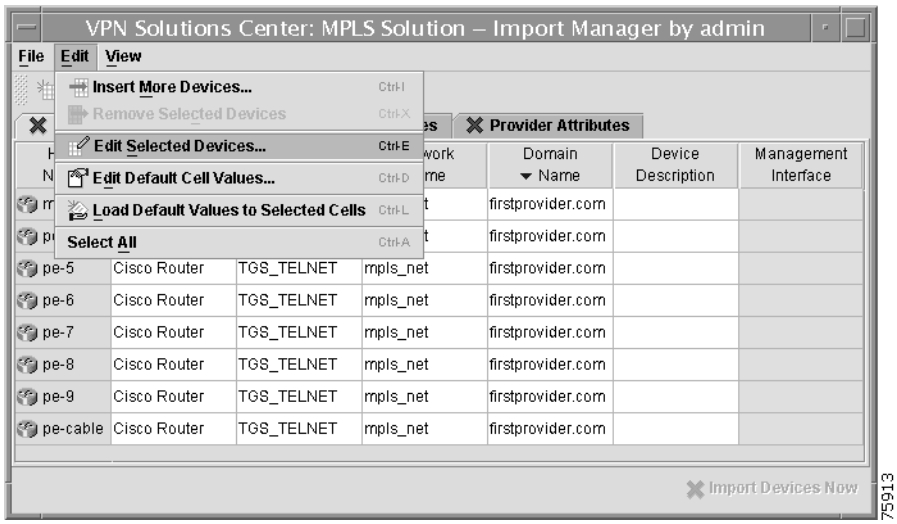

The Edit Parameters dialog box is displayed (see Figure 4-13).

Figure 4-13 Editing the Management Interface Parameter

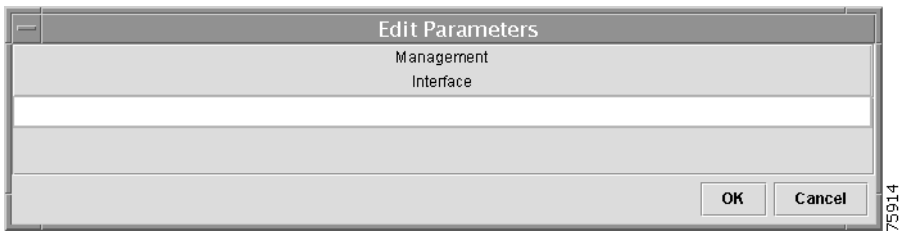

As you can see in Figure 4-13, the Edit Parameters dialog box displays the name of the parameter being edited—in this example, the management interface for the selected PEs.

Step 3 Place the cursor anywhere in the edit field and **double-click**.

A dialog box for the selected attribute appears (see Figure 4-14).

Figure 4-14 Selecting the Management Interface for the Devices

Step 4 Select the management interface, then click **OK**.

You return to the Edit Parameters dialog box, where the value you selected is displayed.

Step 5 To accept the selected value, click **OK**.

The new value is populated to all the selected devices (see Figure 4-15).

Figure 4-15 Attribute Populated to Multiple Devices

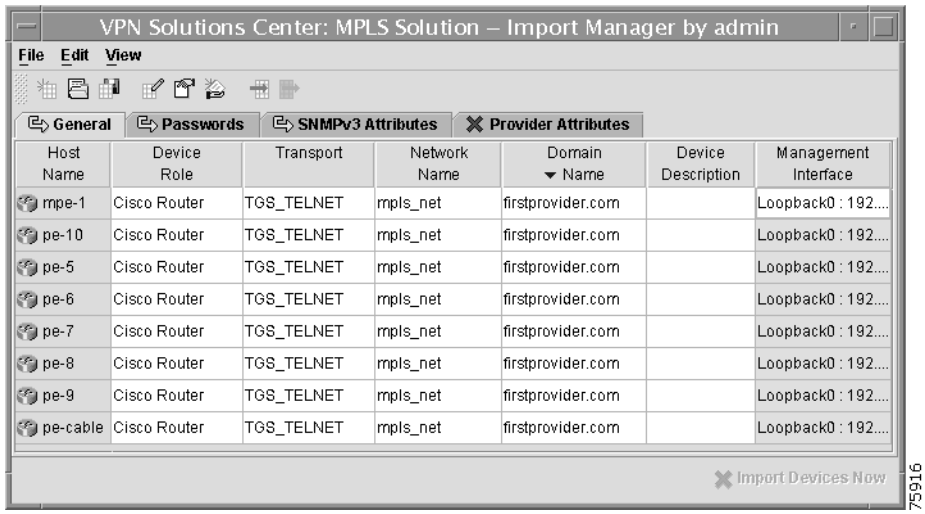

As you can see in Figure 4-15, the red X in the General tab is now a yellow arrow, indicating that all the required general attributes are defined and ready for device import.

Specifying a Parameter Value for a Single Device

Although the Import Manager allows you to define the attributes for any number of imported devices at once, there may be occasions when you need to edit an attribute for a particular device, not a group of devices.

In this example, we show the management interface being set for a particular router.

To specify a parameter value for a particular device, follow these steps:

Step 1 From the Import Manager's General attributes dialog box, select the appropriate cell for the specific device attribute.

Step 2 In the column for the device attribute of interest, **double-click** the appropriate cell.

The edit dialog box for the selected attribute is displayed. This dialog box can display either a list of items you can choose from (as shown in Figure 4-16) or a field in which you need to enter a value.

Figure 4-16 Selecting a Device's Management Interface

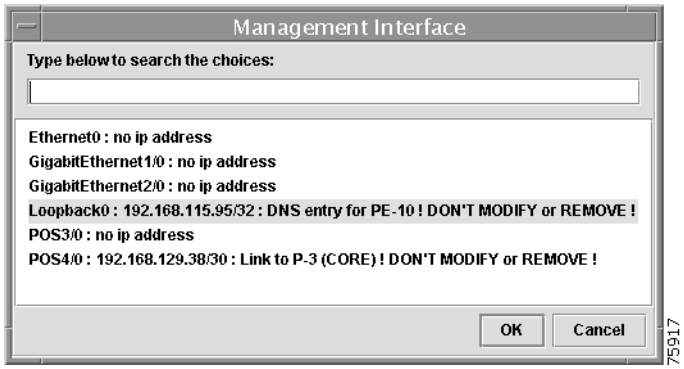

Step 3 Select the appropriate option or enter the appropriate value, then click **OK**.

Specifying the Default Values for the Imported PE Routers

You are now ready to specify any default values for the PE parameters. Once you specify the default values, you can apply them to all or any subset of the imported devices.

Specifying default values is not a requirement—it is provided as a convenience to the Network Administrator. You always have the option of defining the parameters and attributes of the selected devices without specifying default values.

Step 1 From the Import Manager, choose **Edit > Edit Default Cell Values**.

The MPLS PE Default Values Editor appears (see Figure 4-17).

Step 2 To specify the default value for any value, **double-click** the appropriate cell. A value editor is displayed.

Step 3 Select or enter the appropriate default value for each parameter.

Step 4 When you have defined each parameter to your satisfaction, click **OK**, then proceed to the next tab and define the next set of attributes as needed.

Step 5 When all the default attributes are assigned, you will then import the default values into the Import Manager.

Figure 4-17 MPLS PE Default Values Editor: General Parameters

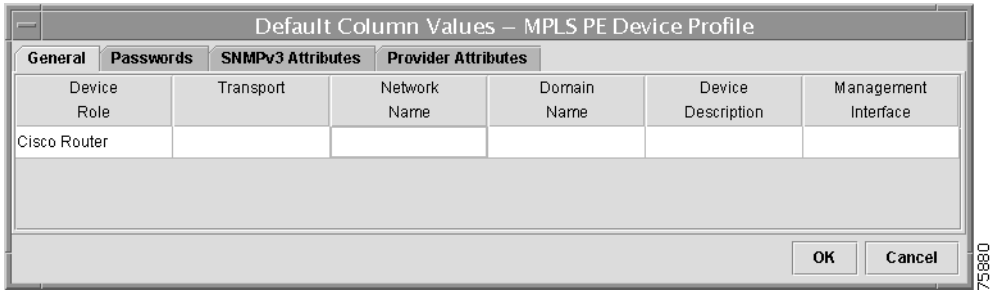

Specifying the General Parameters for PEs

As you can see in Figure 4-17, you can set the default values for all the selected devices for the following general parameters:

- Device Role (set by default to *Cisco Router*)
- Transport (set by default to *TGS_TELNET*)
- Network Name
- Domain Name
- Device Description (optional)
- Management Interface

Default Device Role

When importing PEs into VPNSC, the device role is set by default to **Cisco Router**.

Default Transport Mechanism

The **Transport** parameter configures the method of communication between the VPN Solutions Center workstation and the specified PE routers. The default transport mechanism for MPLS operations is **TGS_TELNET**.

Step 1 To change the default transport mechanism, **double-click** the *Transport* cell.

The Transport dialog box is displayed (see Figure 4-18).

Figure 4-18 Specifying the Default Transport Mechanism

Step 2 From the list of transport mechanisms, choose the configuration file transport method you are using.

- *TGS_Telnet*: The *TGS_TELNET* option is the default transport method for MPLS VPNs.
- *TGS_SSH*: The configuration file transport method for VPN Solutions Center IPsec mode is *TGS_SSH* (Telnet Gateway Server—Secure Shell).
- *TGS_TFTP*: If you choose TGS_TFTP as the default transport method, be sure to enable TFTP (Trivial File Transfer Protocol) on the VPN Solutions Center workstation and on the target routers.

  For details, see the "Enabling TFTP in VPN Solutions Center" section on page 2-7.

Step 3 Click **OK**.

You return to the MPLS PE Default Values Editor, where the selected default transport method is now displayed in the *Transport* cell.

Default Network Name

Step 1 To set the default network for all the PE routers imported into VPNSC, **double-click** the *Network Name* cell.

The Network Name dialog box is displayed (see Figure 4-19).

Figure 4-19 Specifying the Default Network

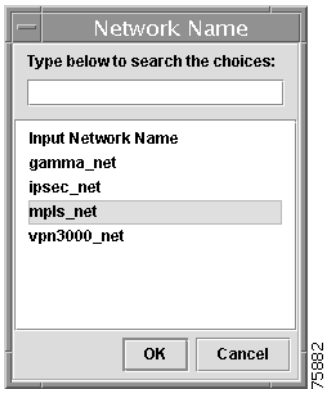

Step 2 From the list of networks, you can either choose the name of one of the networks listed or enter a new network name.

Step 3 Click **OK**.

The network name you specified is displayed in the Default Editor's *Network Name* cell.

Default Domain Name

If you do not specify the domain server name, VPN Solutions Center uses the target name; then, using, DNS (Domain Name System), it performs a domain lookup.

Step 1 To set the default domain for all the PE routers imported into VPNSC, **double-click** the *Domain Name* cell.

The Domain Name dialog box is displayed (see Figure 4-20).

Figure 4-20 Specifying the Default Domain

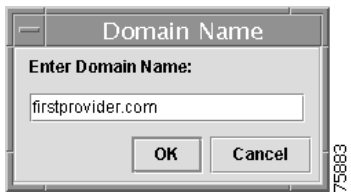

Step 2 Enter the default domain name for the selected PEs.

Step 3 Click **OK**.

The domain name you specified is displayed in the Default Editor's *Domain Name* cell.

Default Device Description

A default description of all the imported PEs may or may not be useful. For this reason, this field is optional. You can enter up to 256 characters in a device description.

If you do wish to enter a default device description for the imported PEs, **double-click** the *Device Description* cell, enter the description in the **Device Description** dialog box, then click **OK**.

Default Management Interface

The VPN Solutions Center Network Management Subnet resides inside the service provider network, and communicates with edge routers through an assigned *management interface*. Configuration changes are managed by VPN Solutions Center software and transported to the appropriate edge routers through the management interface.

This task assigns a default management interface to the set of PEs that you are importing. Doing so is possible only if the edge devices all use the same type of interface; for example, if they all use S0 for their management interface.

Tip When setting up network devices in VPN Solutions Center, be sure to have both the DNS-resolvable hostname and an IP address designated as the management interface specified for each device. If VPNSC cannot access a device via DNS, and no routable IP address is specified for the target devices in VPNSC, data collection operations will fail.

Step 1 To set the default management interface for all the PE routers imported into VPNSC, **double-click** the *Management Interface* cell.

The Management Interface dialog box is displayed (see Figure 4-21).

Figure 4-21 Specifying the Default Management Interface

Step 2 Enter one or two default management interface(s) for the set of PEs.

Note Use a *semicolon* (**;**) to separate multiple interface names.

The interface entered first is the interface that VPNSC searches for and selects first. Using the example in Figure 4-21, VPNSC will search each of the PE devices for a Loopback0 interface. If a Loopback0 interface is found on a device, VPNSC assigns the management interface for that device to Loopback0. If that interface is not found, VPNSC searches each PE in the set for the next interface in the list, in this case, S0. In cases where S0 exists, but Loopback0 does not, VPNSC will assign the S0 interface as the management interface.

Step 3 Click **OK**.

The management interfaces you specified are displayed in the Default Editor's *Management Interface* cell.

This completes the general default values for PEs (see Figure 4-22).

Figure 4-22 General Parameters Completed in the MPLS PE Default Values Editor

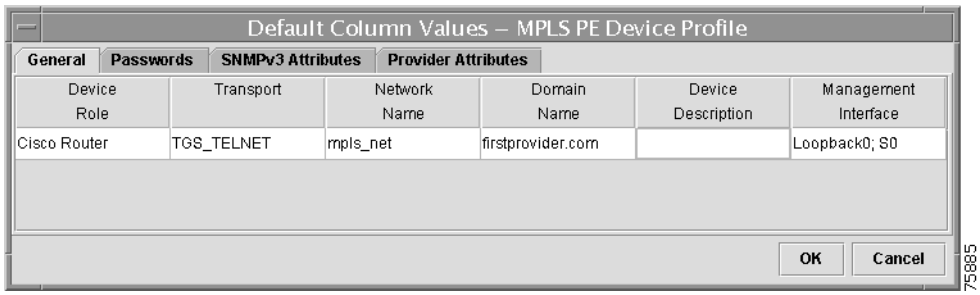

Step 4 To save the work you have done so far, click **OK**.

The default values for the general PE parameters are saved to the Repository. You return to the Import Manager dialog box.

Specifying the Default Passwords for PEs

Caution VPN Solutions Center requires that the PEs and managed CEs in the network have a login password (also called the *virtual terminal* password). Data collection operations fail if VPN Solutions Center does not find the login password set on routes it attempts to collect data from.

To specify the default passwords for the set of PEs you are importing, follow these steps:

Step 1 From the MPLS PE Import Manager, choose the **Passwords** tab.

The Passwords dialog box appears (see Figure 4-23).

Figure 4-23 Password Parameters in the MPLS PE Import Manager

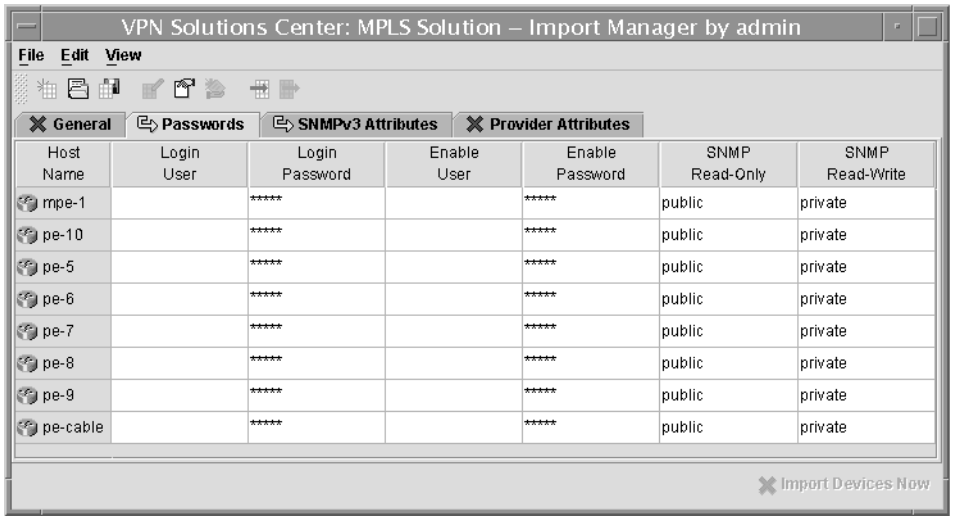

As you can see in this figure, some of the password attributes, such as the login password, enable password, and so on, were picked up from the routers' configuration files and populated into the appropriate columns in the dialog box.

The yellow arrow in the Passwords tab indicates that even though some attributes are not filled in, the passwords for the selected PEs could be imported without further additions or changes. However, if you wish to specify additional default password attributes, proceed to the next step.

Step 2 From the Import Manager menu bar, choose **Edit > Default Cell Values**.

Step 3 Choose the **Passwords** tab (see Figure 4-24).

Figure 4-24 Default PE Passwords Parameters Dialog Box

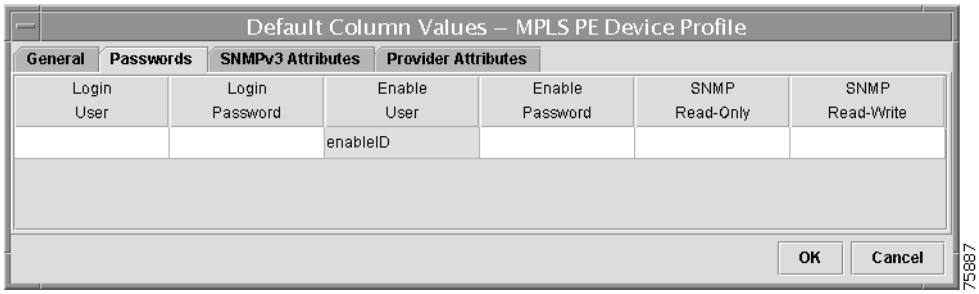

Default Login User

Step 1 To set the default login username for all the PE routers imported into VPNSC, **double-click** the *Login User* cell. The Login User dialog box is displayed (see Figure 4-25).

Figure 4-25 Specifying the Default Login User

Step 2 Enter the default login user name here, then click **OK**.

Default Login Password

The login password is the router's virtual terminal password, which establishes password protection on incoming Telnet sessions.

On a router that is to be accessed by a terminal server, the login password and the console password should be identical. A terminal server accesses a router via the router's console port.

Step 1 To set the default login password for all the PE routers imported into VPNSC, **double-click** the *Login Password* cell. The Login Password dialog box is displayed (see Figure 4-25).

Figure 4-26 Specifying the Default Login Password

Step 2 Enter the default login password.

Step 3 In the *Verify Password* field, enter the login password again, then click **OK**.

Default Enable User

Step 1 To set the default enable username for all the PE routers imported into VPNSC, **double-click** the *Enable User* cell. The Enable User dialog box is displayed (see Figure 4-27).

Figure 4-27 Specifying the Default Enable Username

Step 2 Enter the default enable username, then click **OK**.

Default Enable Password

If desired, set the default enable password here for the selected devices.

Step 1 To set the default enable password for all the PE routers imported into VPNSC, **double-click** the *Enable Password* cell. The Enable Password dialog box is displayed (see Figure 4-28).

Figure 4-28 Specifying the Default Enable Password

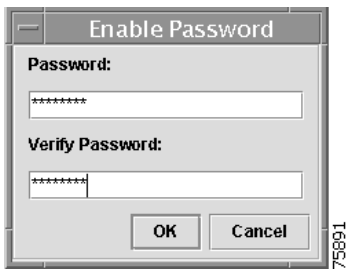

Step 2 Enter the default enable password.

Step 3 In the *Verify Password* field, enter the enable password again, then click **OK**.

Default SNMP Read-Only Community String

The SNMP Read-Only community string is used to read MIB variables.

Note The SNMP community strings must be set on all the PEs and CEs in the service provider's network; the SNMP settings specified in VPN Solutions Center must match the SNMP string values configured for the routers. For related information, see the "Setting Up SNMPv1 and SNMPv2 on the Routers in the Service Provider Network" section on page 2-4 and the "Setting the SNMPv3 Parameters on the Routers in the Service Provider Network" section on page 2-5.

Step 1 To set the SNMP Read-Only community string for all the PE routers imported into VPNSC, **double-click** the *SNMP Read-Only* cell (see Figure 4-29).

Figure 4-29 Specifying the Default SNMP Read-Only String

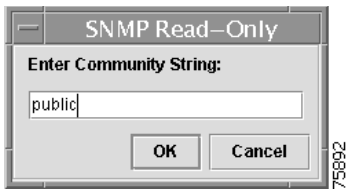

Step 2 Enter the default SNMP Read-Only community string, then click **OK**.

Default SNMP Read-Write Community String

The SNMP Read-Write community string is used to set MIB variables.

Step 1 To set the SNMP Read-Write community string for all the PE routers imported into VPNSC, **double-click** the *SNMP Read-Write* cell (see Figure 4-30).

Figure 4-30 Specifying the Default SNMP Read-Write String

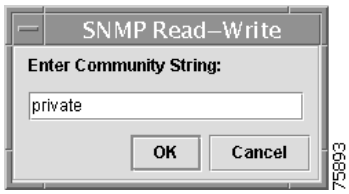

Step 2 Enter the default SNMP Read-Write community string, then click **OK**.

Step 3 If satisfied with the default password settings, click **OK**.

Specifying the Default SNMPv3 Attributes for PEs

Simple Network Management Protocol Version 3 (SNMPv3) is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting packets over the network.

The values you set here must match the actual SNMPv3 values configured on the selected device (see "Setting the SNMPv3 Parameters on the Routers in the Service Provider Network" section on page 2-5).

To specify the default SNMPv3 attributes for the set of PEs you are importing, follow these steps:

Step 1 From the MPLS PE Default Values Editor, choose the **SNMPv3** tab.

The Import Manager's SNMPv3 Attributes dialog box appears (see Figure 4-31).

Figure 4-31 SNMPv3 Attributes in the MPLS PE Import Manager

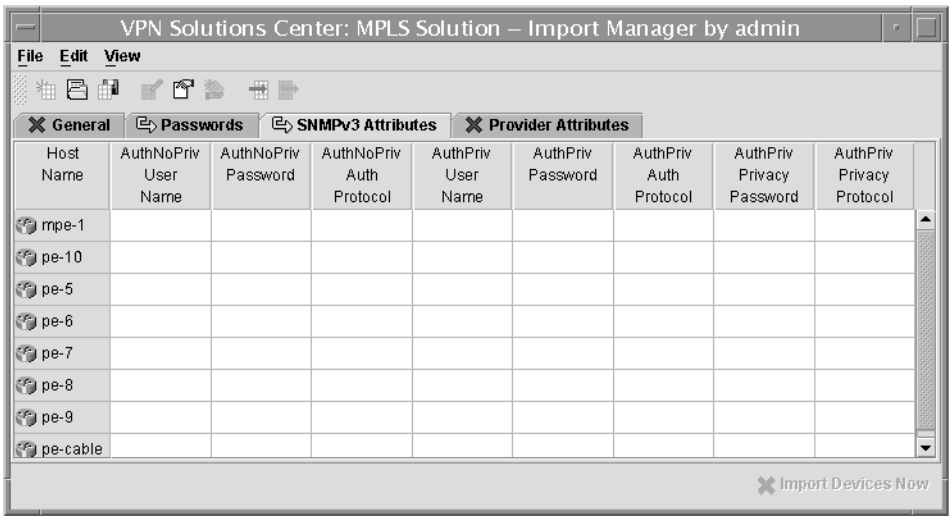

The yellow arrow in the Import Manager's SNMPv3 tab indicates that even though none of the SNMPv3 attributes are filled in, the selected PEs could be imported without further additions or changes to the SNMPv3 area. However, if you wish to specify additional default SNMPv3 attributes, proceed to the next step.

Step 2 From the Import Manager menu bar, choose **Edit > Default Cell Values**.

The MPLS PE Default Values Editor appears.

Step 3 Choose the **SNMPv3 Attributes** tab (see Figure 4-32).

Figure 4-32 Default SNMPv3 Attributes

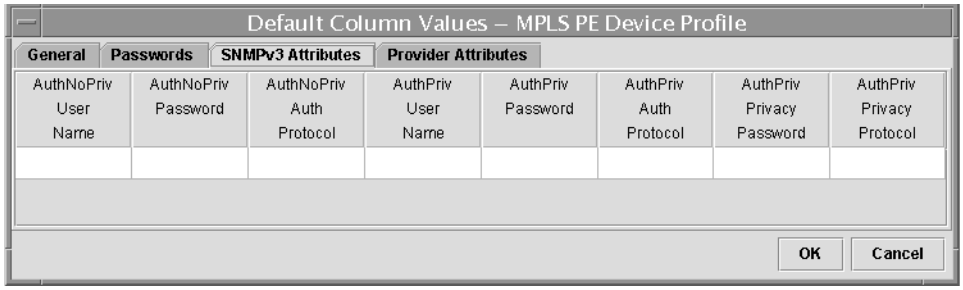

Default AuthNoPriv User Name

The AuthNoPriv user must have permission to the object identification numbers (OIDs) specified in the security request (that is, write permission for a set request, and read permission for a get request).

Step 1 To set the default AuthNoPriv username for all the PE routers imported into VPNSC, **double-click** the *AuthNoPriv User Name* cell.

The AuthNoPriv User Name dialog box is displayed.

Step 2 Enter the default AuthNoPriv user name configured on the specified edge device routers, then click **OK**.

Default AuthNoPriv Password

Step 1 To set the default AuthNoPriv password for all the PE routers imported into VPNSC, **double-click** the *AuthNoPriv Password* cell.

The AuthNoPriv Password dialog box is displayed.

Step 2 *Password*: Enter the default AuthNoPriv authentication password configured on the specified edge device routers.

Step 3 *Verify Password*: Enter the password again to verify it, then click **OK**.

Default AuthNoPriv Authentication Protocol

The available AuthNoPriv Authentication protocol options are **None**, **MD5**, or **SHA**.

- **MD5**: The MD5 algorithm is intended for digital signature applications, where a large file must be compressed in a secure manner before being encrypted with a private (secret) key under a public-key cryptosystem such as RSA. The MD5 algorithm takes as input a message of arbitrary length and produces as output a 128-bit *fingerprint* or *message digest* of the input.

- **SHA**: Secure Hash Algorithm. Computes a condensed representation of a message or a data file. When a message of any length is input, the SHA-1 produces a 160-bit output called a *message digest*. The message digest can then be input to the Digital Signature Algorithm (DSA), which generates or verifies the signature for the message. The creator of the digital signature and the verifier of the digital signature must use the same hash algorithm.

Step 1 To set the default AuthNoPriv authentication protocol for all the PE routers imported into VPNSC, **double-click** the *AuthNoPriv Auth Protocol* cell.

The AuthNoPriv Auth Protocol dialog box is displayed (see Figure 4-33).

Figure 4-33 Specifying the Default Authentication Protocol

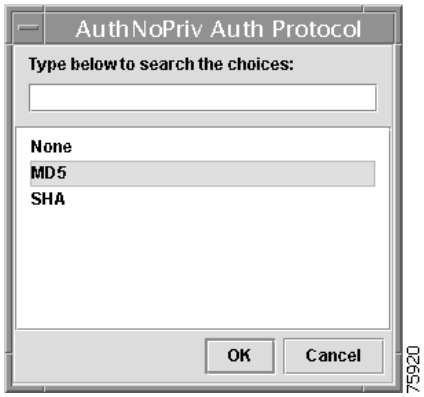

Step 2 Select the authentication protocol from the list, then click **OK**.

Default AuthPriv User Name

The AuthPriv user must have permission to the object identification numbers (OIDs) specified in the security request (that is, write permission for a set request, and read permission for a get request).

Step 1 To set the default AuthPriv username for all the PE routers imported into VPNSC, **double-click** the *AuthPriv User Name* cell.

The AuthPriv User Name dialog box is displayed.

Step 2 Enter the default AuthPriv user name configured on the specified edge device routers, then click **OK**.

Default AuthPriv Password

Step 1 To set the default AuthPriv password for all the PE routers imported into VPNSC, **double-click** the *AuthPriv Password* cell.

The AuthPriv Password dialog box is displayed.

Step 2 *Password*: Enter the default authentication password configured on the specified edge device routers.

Step 3 *Verify Password*: Enter the password again to verify it, then click **OK**.

Default AuthPriv Authentication Protocol

The available AuthPriv Authentication protocol options are **None**, **MD5**, or **SHA**.

- **MD5**: The MD5 algorithm is intended for digital signature applications, where a large file must be compressed in a secure manner before being encrypted with a private (secret) key under a public-key cryptosystem such as RSA. The MD5 algorithm takes as input a message of arbitrary length and produces as output a 128-bit *fingerprint* or *message digest* of the input.
- **SHA**: Secure Hash Algorithm. Computes a condensed representation of a message or a data file. When a message of any length is input, the SHA-1 produces a 160-bit output called a *message digest*. The message digest can then be input to the Digital Signature Algorithm (DSA), which generates or verifies the signature for the message. The creator of the digital signature and the verifier of the digital signature must use the same hash algorithm.

Step 1 To set the default AuthPriv authentication protocol for all the PE routers imported into VPNSC, **double-click** the *AuthPriv Auth Protocol* cell.

The AuthPriv Auth Protocol dialog box is displayed.

Step 2 Select the appropriate authentication protocol from the list, then click **OK**.

Default AuthPriv Privacy Password

The privacy password is the encryption password.

Step 1 To set the default AuthPriv privacy password for all the PE routers imported into VPNSC, **double-click** the *AuthPriv Password* cell.

The AuthPriv Privacy Password dialog box is displayed.

Step 2 *Password*: Enter the default AuthPriv privacy (encryption) password configured on the specified edge device routers.

Step 3 *Verify Password*: Enter the encryption password again to verify it, then click **OK**.

Default AuthPriv Privacy Protocol

Currently, the only AuthPriv privacy protocol supported is **DES-56**.

Data Encryption Standard (DES) encrypts packet data. The Cisco IOS implements the mandatory 56-bit DES-CBC (Cipher Block Chaining) with the explicit initialization vector. Cipher Block Chaining requires an initialization vector to start encryption. The initialization vector is given in the packet. Triple DES (3DES) adds security by performing the operation three times with different subkeys.

Step 1 To set the default AuthPriv privacy protocol for all the PE routers imported into VPNSC, **double-click** the *AuthPriv Protocol* cell.

The AuthPriv Privacy Protocol dialog box is displayed.

Step 2 Specify the privacy protocol, then click **OK**.

Step 3 When finished defining the default AuthPriv privacy protocol for the selected devices, click **OK**.

You return to the Import Manager dialog box.

Defining the Default Provider Attributes for PEs

A Provider Administrative Domain (PAD) is an administrative domain defined by an Internet Service Provider. A PAD is also the set of all PE devices in one BGP autonomous system. For more details about Provider Administrative Domains and instructions on how to create a new PAD, see the "Defining Provider Administrative Domains" section on page 4-30.

A Region is a group of provider edge routers (PEs) within a single BGP autonomous system. Regions must have a name, assigned PEs, and their corresponding IP address pools. For more details about Regions and instructions on how to create a new Region, see the "Defining Provider Administrative Domains" section on page 4-30.

Specifying a default Region may or not be useful, depending on whether the service provider has one Region or multiple Regions. If the provider has multiple Regions, you may choose to specify the appropriate Regions directly in the Import Manager dialog box.

To specify the default Provider attributes for the set of PEs you are importing, follow these steps:

Step 1 From the MPLS PE Import Manager, choose the **Provider Attributes** tab.

The Provider Attributes dialog box appears (see Figure 4-34).

Figure 4-34 Provider Attributes in the MPLS PE Import Manager

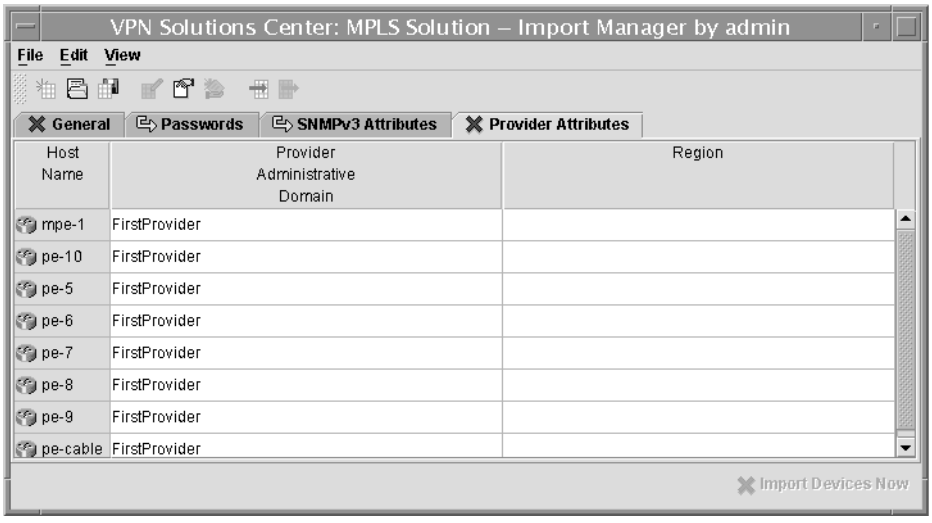

As you can see in Figure 4-34, VPNSC populated the Provider Administrative Domain column with the name of the service provider. But the Region column is not filled in.

The red X in the Provider Attributes tab indicates that the Region attribute is a required attribute and it must be defined before you can import the PE configuration files.

Step 2 From the Import Manager menu bar, choose **Edit > Default Cell Values**.

The MPLS PE Default Values Editor appears.

Step 3 Choose the **Provider Attributes** tab (see Figure 4-35).

Figure 4-35 Default Provider Attributes

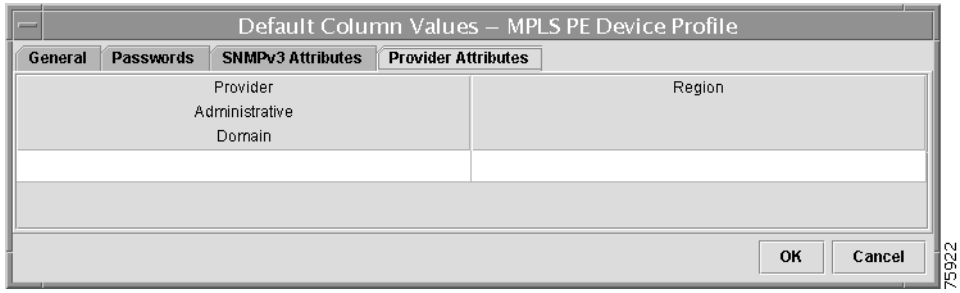

Default Provider Administrative Domain

You can specify the default Provider Administrator Domain or create a new PAD. For details on the procedure for creating a new PAD, see the "Defining Provider Administrative Domains" section on page 4-30.

Step 1 To set the default Provider Administrative Domain for all the PE routers imported into VPNSC, **double-click** the *Provider Administrator Domain* cell.

The Provider Administrator Domain dialog box is displayed (see Figure 4-36).

Figure 4-36 Specifying the Default Provider Administrative Domain

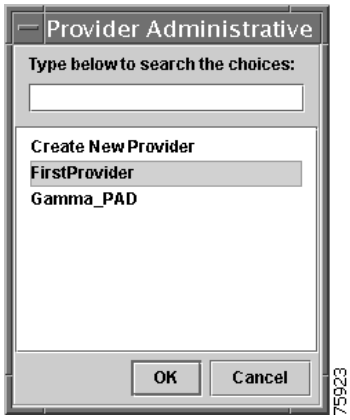

Step 2 Select the default Provider Administrator Domain from the list, then click **OK**.

Default Region

You can specify the default Region or create a new Region. For details on creating a new Region, see the "Defining Provider Administrative Domains" section on page 4-30.

Step 1 To set the default Region for all the PE routers imported into VPNSC, **double-click** the *Region* cell.

The Region dialog box is displayed (see Figure 4-37).

Figure 4-37 Specifying the Default Region

Step 2 Select the default Region from the list, then click **OK**.

You return to the Default Provider Attributes dialog box.

Step 3 If you are finished defining the default Provider attributes for the selected devices, click **OK**.

You return to the Import Manager dialog box.

You have now completed the default attributes for the PEs. You may notice that the red X's are still displayed in some of the Import Manager tabs. This is because the default values you have specified need to be imported into the Import Manager.

Importing the PE Configuration Files Into the Repository

When you have completed defining the default attributes for a set of devices, you must first import those default values into the Import Manager. The default values for each tab are loaded separately. For example, if you defined default values for the General attributes and the Provider attributes, you need to load the default values into both areas.

When all the necessary values are specified in the Import Manager, you can then import the configuration files into the VPN Solutions Center Repository.

To import the default values and import the PE configuration files, follow these steps:

Step 1 From the MPLS PE Import Manager, choose a tab category for which you created default values.

Step 2 Choose **Edit > Select All**.

This operation selects all the devices listed in the Import Manager dialog box.

Tip If you wish to load the default values to specific devices displayed in the Import Manager dialog box (not all the listed devices), place the cursor on the appropriate host name icon, then press **Ctrl+Click**. Repeat this as necessary to select the desired devices.

Step 3 With the appropriate devices selected, choose **Edit > Load Default Values to Selected Cells**.

The Import Manager loads the default values that you specified into the appropriate cells (see Figure 4-38).

Figure 4-38 Default PE Values Loaded into the Import Manager

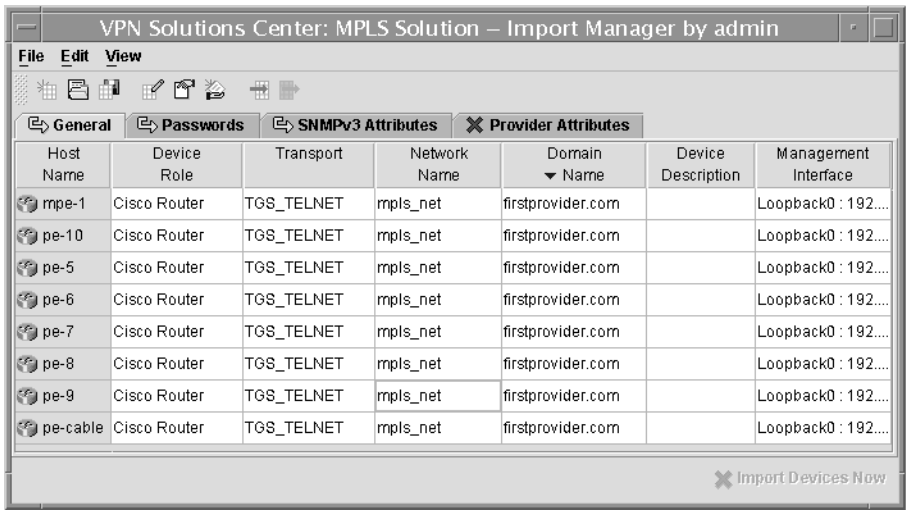

Step 4 Repeat this procedure for each tab category for which you created default values.

When you are finished, the Import Devices Now button (in the bottom right corner of the Import Manager dialog box) is enabled (see Figure 4-39).

Figure 4-39 Import Now Button Enabled

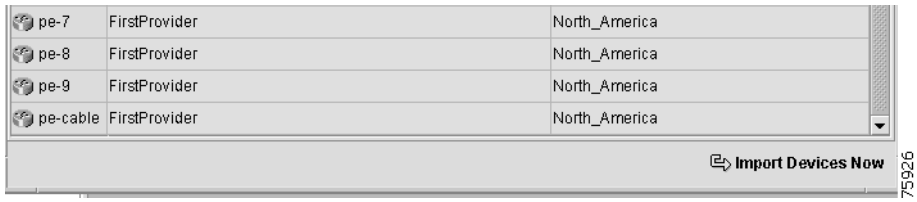

You are now ready to import the PE devices into the VPN Solutions Center Repository.

Step 5 To import the PEs into VPNSC, click **Import Devices Now**.

The selected configuration files and the additional information specified in this import procedure are imported into VPN Solutions Center.

With this task is completed, you can proceed to importing Customer Edge routers in VPN Solutions Center.

Defining Provider Administrative Domains

A Provider Administrative Domain (PAD) is an administrative domain defined by an Internet Service Provider. In practical terms, a PAD is the set of all PE devices in one BGP autonomous system (AS). The network owned by the PAD is called a *backbone network*. Each PAD includes a route distinguisher and route target and IP address pools. Each Provider Administrative Domain can have many *Regions* within it. If an ISP requires two AS numbers, it must consist of two provider administrative domains. Each provider administrative domain has Regions that have a route distinguisher (*RD*), a route target (*RT*), and an IP address pool from which to automatically generate IP address values during provisioning.

The VPN Solutions Center software allows you to define as many *Regions* within a Provider Administrative Domain (PAD) as you need. PADs are divided into Regions in much the same way that customers are divided into sites. A Region is considered to be a group of provider edge routers (PEs) within a single BGP autonomous system. The primary objective for defining Regions is to allow a provider to employ unique IP address pools in large Regions, such as Europe, Asia Pacific, and so forth.

Note that a provider can also assign PEs to these Regions, thereby simplifying the PE selection process (for example, only presenting PEs in the European Region when adding services to a European customer edge router).

Tip Cisco recommends that providers create one Provider Administrative Domain and then define the Regions within the PAD.

Before You Begin

Before you begin this procedure, have the following information at hand:

- The BGP autonomous system (AS) number

  There is generally one BGP AS number per Provider Administrative Domain.
- The names of the PE routers within the Region
- The IP address pools for point-to-point links (that is, the IP numbered links)
- The IP address pools for loopback links (that is, the IP unnumbered links)

To define a new Provider Administrative Domain, follow these steps:

Step 1 From the MPLS PE Import Manager, choose the **Provider Attributes** tab.

The Provider Attributes dialog box appears (see Figure 4-40).

Figure 4-40 Provider Attributes in the MPLS PE Import Manager

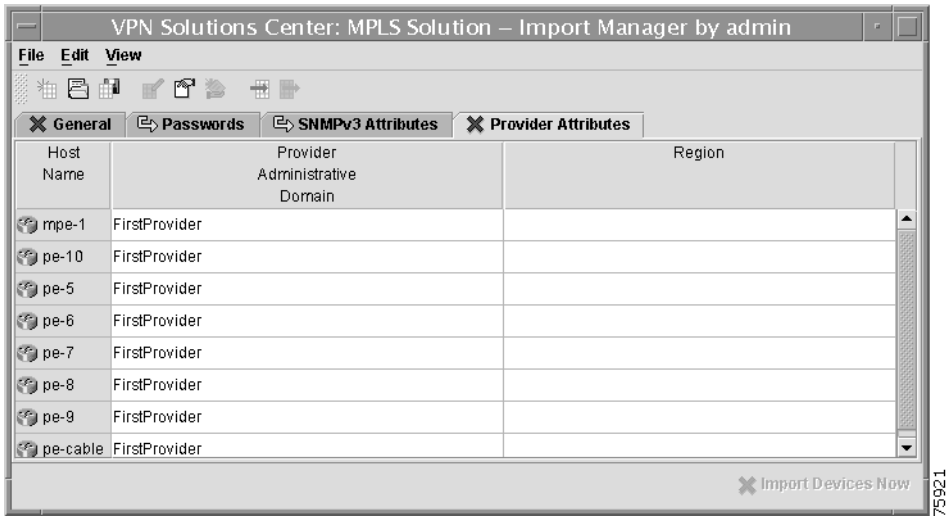

Step 2 Select the *Provider Administrative Domain* cells of interest, then choose **Edit > Edit Selected Devices**.

The Edit Parameters: Provider Administrative Domain dialog box appears (see Figure 4-41).

Figure 4-41 Editing the Provider Administrative Domain

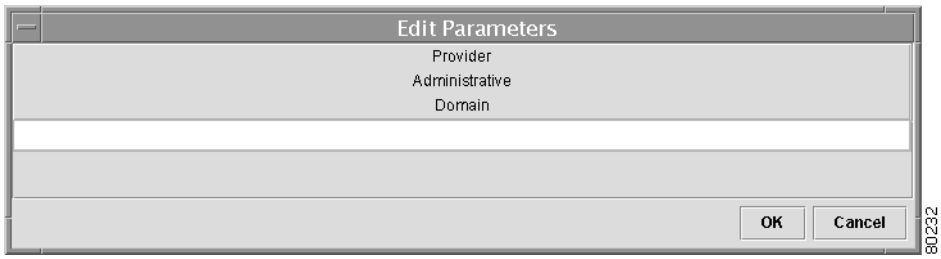

Step 3 Place your cursor anywhere in the field, then **double-click**.

The Select Provider Administrative Domain dialog box appears (see Figure 4-42).

Figure 4-42 Creating a New Provider Domain

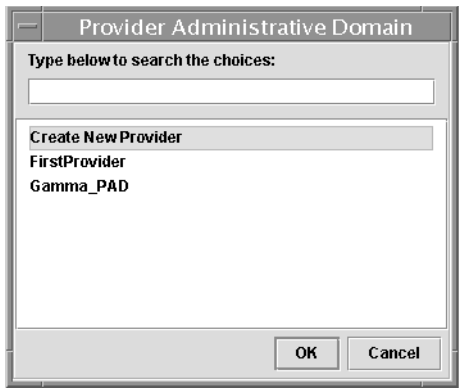

Step 4 Choose **Create New Provider**, then click **OK**.

The New Provider Administrative Domain dialog box appears (see Figure 4-43).

Figure 4-43 The New Provider Administrative Domain Dialog Box

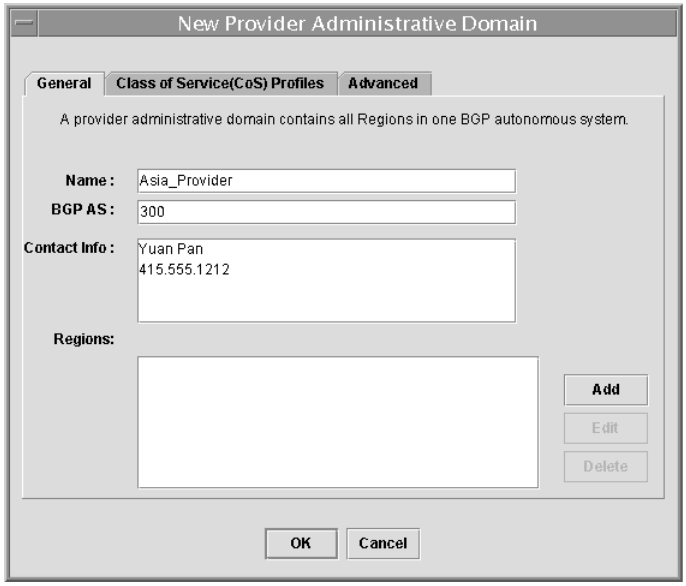

Step 5 Enter the general information for the Provider Administrative Domain.

a. *Name*: Enter the name of the Provider Administrative Domain.

b. *BGP AS*: Enter the BGP autonomous system number.

Each autonomous system is assigned a unique 16-bit number by the same central authority that assigns IP network numbers.

c. *Contact Info*: The contact information is optional, but it is a good idea to provide it.

Step 6 To add a Region to the new Provider Domain, click **Add**.

The New Region dialog box appears (see Figure 4-44).

A Region is a group of provider edge routers (PEs) within a single BGP autonomous system. Regions must have a name, assigned PEs, and their corresponding IP address pools.

Figure 4-44 Specifying a New Region

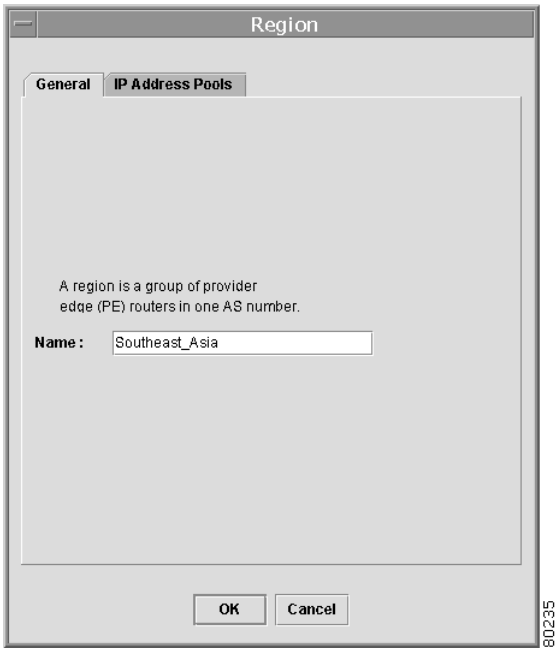

Step 7 *Name*: Enter the name of the Region.

Now you must define the IP address pools for the Region.

The VPN Solutions Center software uses IP address pools to automatically assign IP addresses to PEs and CEs. Each Region has an IP address pool to use for IP numbered addresses (point-to-point address pool) and a separate IP address pool for IP unnumbered address (loopback address pool).

Within a VPN or extranet, all IP addresses must be unique. Customer IP addresses must not overlap with the provider's IP addresses. Overlapping IP addresses are only possible when two devices cannot see each other—that is, when they are in isolated VPNs.

Caution Due to security and maintenance issues, Cisco does not recommend using customer IP addresses on the PE-CE link.

The VPN Solutions Center 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.

Step 8 Choose the **IP Address Pools** tab (see Figure 4-45).

Figure 4-45 Defining a Region's IP Address Pool

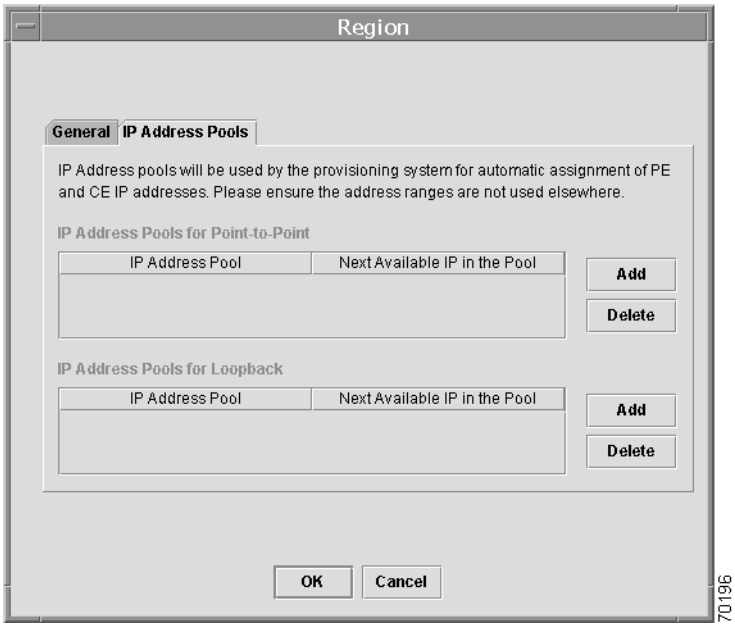

From this dialog box, you can specify IP address pool information for *point-to-point (IP numbered)* links or *loopback (IP unnumbered)* links.

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 9 Choose which type of address pool you are defining and click **Add**.

The New IP Address Pool dialog box appears (see Figure 4-46).

Figure 4-46 Specifying a New Region's IP Address Pool

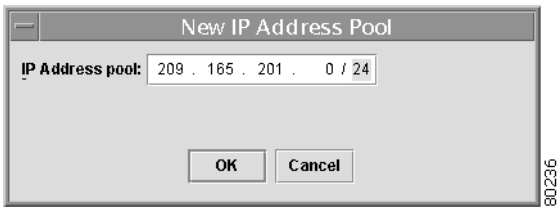

Step 10 Enter the address and subnet mask for the IP address pool, then click **OK**.

You return to the Region: IP Address Pools dialog box, where the new IP address pool information is displayed.

Step 11 Click **OK**.

You return to the New Provider Administrative Domain dialog box, where the new Region name is displayed in the *Regions* field.

Step 12 To complete the Provider Domain and Regions definition, click **OK**.

Defining a Class of Service Profile

A Class of Service (CoS) profile represents a set of CoS configurations offered by a provider to its customer. Each CoS profile consists of a set of CoS classes that record information on how traffic shaping and policing are configured.

The VPN Solutions Center software requires that you create a Class of Service (CoS) profile only if you want the product to provision CoS on the PE-CE link. You can add additional CoS profiles at any time. This procedure only defines the CoS profile—until you invoke it when you activate a service request, the CoS profile has no effect.

Class of Service profiles are applied to the Provider Edge Router (PE), but the CoS definition is enforced across the PE-CE link on both the PE and CE.

To define a Class of Service profile, follow these steps:

Step 1 From the New Provider Administrative Domain dialog box (see Figure 4-43 on page 4-32), choose the **Class of Service (CoS) Profiles** tab.

This dialog box is initially empty.

Step 2 To add a new Class of Service profile, click **Add**.

The New Class of Service Profile dialog box appears (see Figure 4-47).

Figure 4-47 Creating a New Class of Service Profile

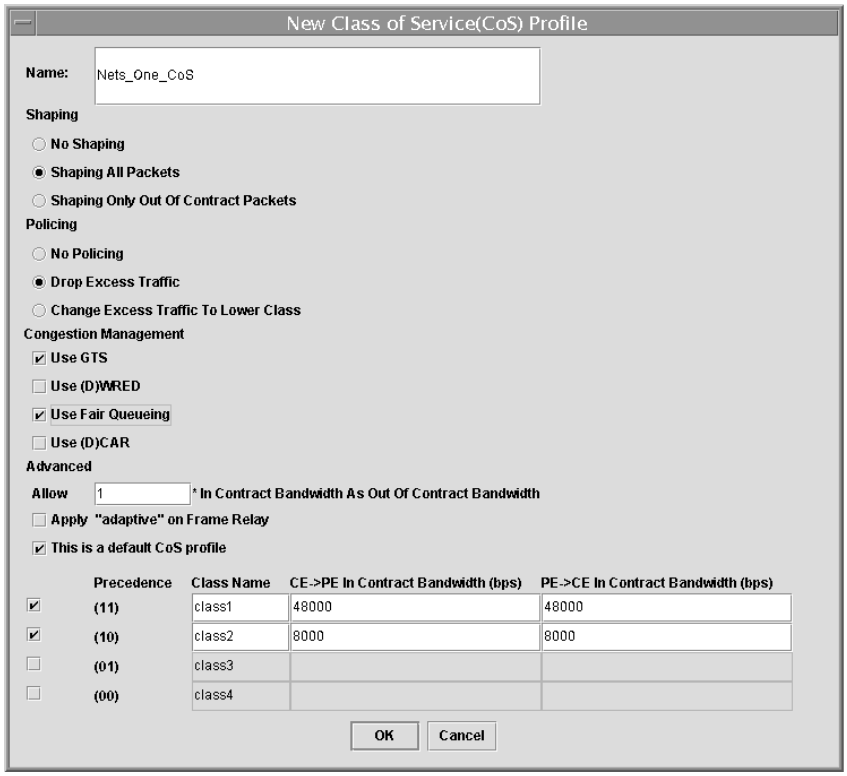

Step 3 Complete the fields in the Class of Service profile that suits your network, then click **OK**.

The CE-to-PE rate must be entered as a multiple of 8,000 and in the range between 8,000 to 2,000,000,000 (in bits per second).

Do not include commas in the numbers you enter.

About In-Contract and Out-of-Contract Bandwidth

The PE can rate limit traffic to the subscribed bandwidth and mark the traffic that is within the specified bandwidth as *in-contract*, and mark traffic above the specified bandwidth as *out-of-contract*.

Marking a packet as in-contract or out-of-contract is done by setting the first bit of the precedence bits in the IP header. The appropriate class is indicated by the remaining two precedence bits (see Table 4-1). Traffic that exceeds any class is marked as out-of-contract, and this traffic can be dropped or mapped to a lower class of service. The out-of-contract bandwidth is initially set to the in-contract bandwidth, but you can set this to the values appropriate for the customer.

Table 4-1 Mapping IP Precedence to Class of Service

IP Precedence	Contract Status	Class of Service
111	In-contract	Class 1
110	In-contract	Class 2
101	In-contract	Class 3
100	In-contract	Class 4
011	Out-of-contract	Class 1
010	Out-of-contract	Class 2
001	Out-of-contract	Class 3
000	Out-of-contract	Class 4

The customer can initially “paint” the packets that leave the customer edge router (the PE is the destination router), and VPN Solutions Center allows policing or repainting of packets that enter the provider edge router.

Customizing the Route Distinguisher and Route Target Values

MPLS-based VPNs employ BGP to communicate between PEs to facilitate customer routes. This is made possible through extensions to BGP that carry addresses other than IPv4 addresses. A notable extension is called the *route distinguisher* (RD).

The purpose of the route distinguisher (RD) is to make the prefix value unique across the network backbone. Prefixes should use the same RD if they are associated with the same set of route targets (RTs) and anything else that is used to select routing policy. The community of interest association is based on the route target (RT) extended community attributes distributed with the Network Layer Reachability Information (NLRI). The RD value must be a globally unique value to avoid conflict with other prefixes.

The MPLS label is part of a BGP routing update. The routing update also carries the addressing and reachability information. When the RD is unique across the MPLS VPN network, proper connectivity is established even if different customers use non-unique IP addresses.

For the RD, every CE that has the same overall role should use a VRF with the same name, same RD, and same RT values. The RDs and RTs are *only* for route exchange between the PEs running BGP. That is, for the PEs to do MPLS VPN work, they have to exchange routing information with more fields than usual for IPv4 routes; that extra information includes (but is not limited to) the RDs and RTs.

VPN Solutions Center software sets the route distinguisher and route target values, but you can assign your own values if you choose (as described in this section).

You can also override the default RD value set by the VPN Solutions Center software. For instructions, see the "Overriding the Default VRF Name and Route Distinguisher Values" section on page 6-17.

Tip You can change the RD and RT values with the VPN Solutions Center software for a given Provider Administrative Domain *only* when creating a new PAD. You cannot edit the RD and RT values once they are initially set.

By default, the product software assigns the RD values as follows:

- CEs with hub connectivity use `BGP_AS:`*value*.
- CEs with spoke connectivity use `BGP_AS:`*value* `+ 1`

Each spoke uses its own RD value for proper hub and spoke connectivity between CEs; therefore, the VPN Solutions Center software implements a new RD for each spoke that is provisioned.

To assign the Route Distinguisher or Route Target values, follow these steps:

Step 1 From the New Provider Administrative Domain dialog box (see Figure 4-48), choose the **Advanced** tab.

Figure 4-48 The New Provider Administrative Domain Dialog Box

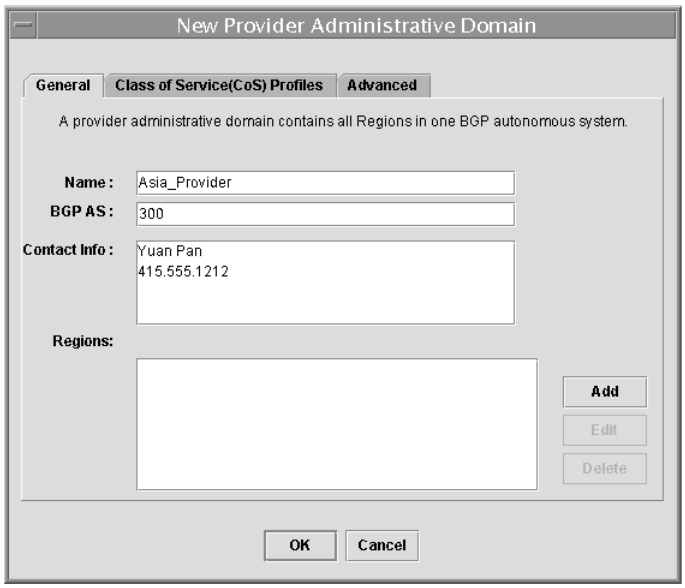

The New Provider Administrative Domain: Advanced dialog box appears (see Figure 4-49), which allows you to alter the default Route Distinguisher and Route Target values.

Figure 4-49 Setting the Route Distinguisher and Route Target Values

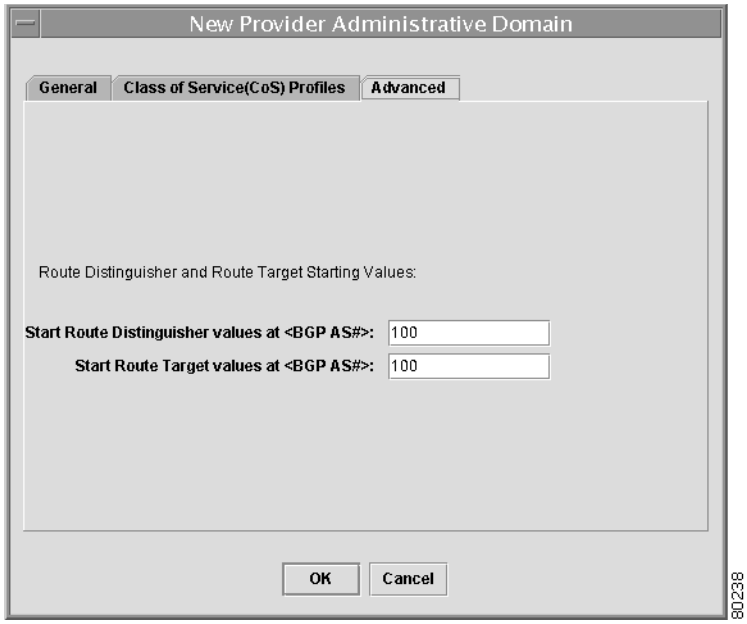

Step 2 *Start Route Distinguisher Values At <BGP AS#>*: Enter the new Route Distinguisher value.

Step 3 *Start Route Target Values At <BGP AS#>*: Enter the new Route Target value.

Step 4 Click **OK**.

Importing Customer Edge Routers into VPN Solutions Center

To import CE router configuration files into VPN Solutions Center, follow these steps:

Step 1 Create a directory of configuration files for a given set of devices and copy the appropriate configuration files into the directory.

Device names within each directory must be unique.

A typical set includes Provider and Customer edge routers (PEs and CEs).

Step 2 From the VPN Console menu, choose **Setup > Create Targets From Configuration Files**.

The opening window for the Import Manager appears (see Figure 4-50).

Figure 4-50 Import Manager: Opening Screen

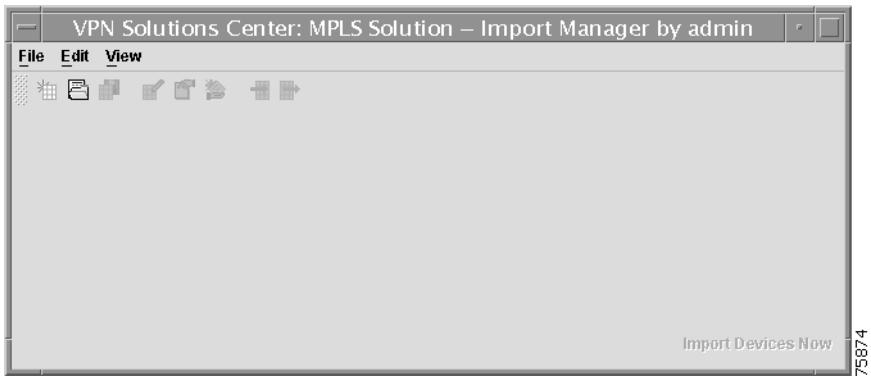

Step 3 From the Import Manager menu bar, choose **File > Import > Import Customer Edge Devices** (see Figure 4-51).

Figure 4-51 Import Customer Edge Devices Menu Option

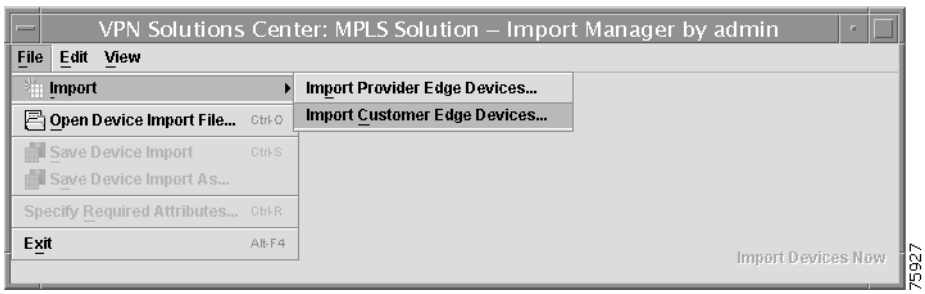

A dialog box appears that lets you change directories to where a set of Cisco IOS configuration files are located (see Figure 4-52).

Figure 4-52 Selecting the Customer Edge Router Configuration Files

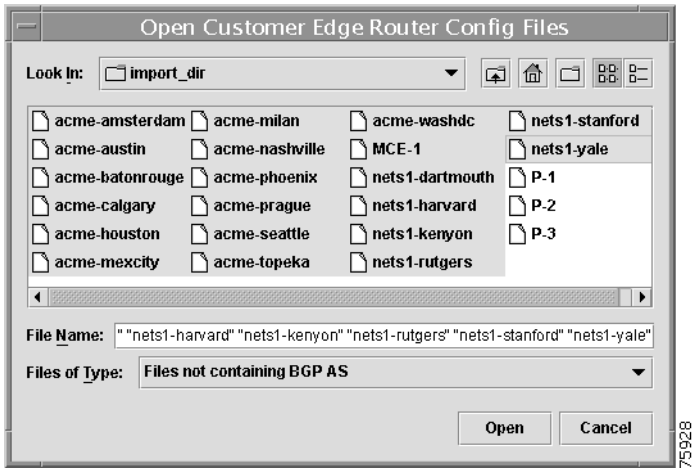

Step 4 Use the Open Customer Edge Router Config Files dialog box to locate and specify the CE router configuration files that you want to import into VPN Solutions Center.

a. *Look In*: Navigate to the directory where the CE configuration files reside.

b. *File Name*: When you select CEs from the displayed list, their filenames are displayed. You can also type in the pertinent CE filenames if you wish.

c. *Files of Type*: When importing CE files, the Import Manager by default lists **Files not containing BGP AS**.

Valid Customer Edge Router configuration files do not include the BGP AS (Border Gateway Protocol Autonomous System) number. With **Files not containing BGP AS** set for the file type, the Import Manager displays CE configuration files (as well as any other router files that are not PEs).

In addition, the *Files of Type* field provides two other file type options:

- **All Files**
- **Files containing BGP AS**

d. When you have selected the CE configuration files you want to import, click **Open**.

The Import Manager displays a dialog box in spreadsheet format that shows the list of CEs you selected for import (see Figure 4-53).

Figure 4-53 CEs Displayed in Import Manager Spreadsheet

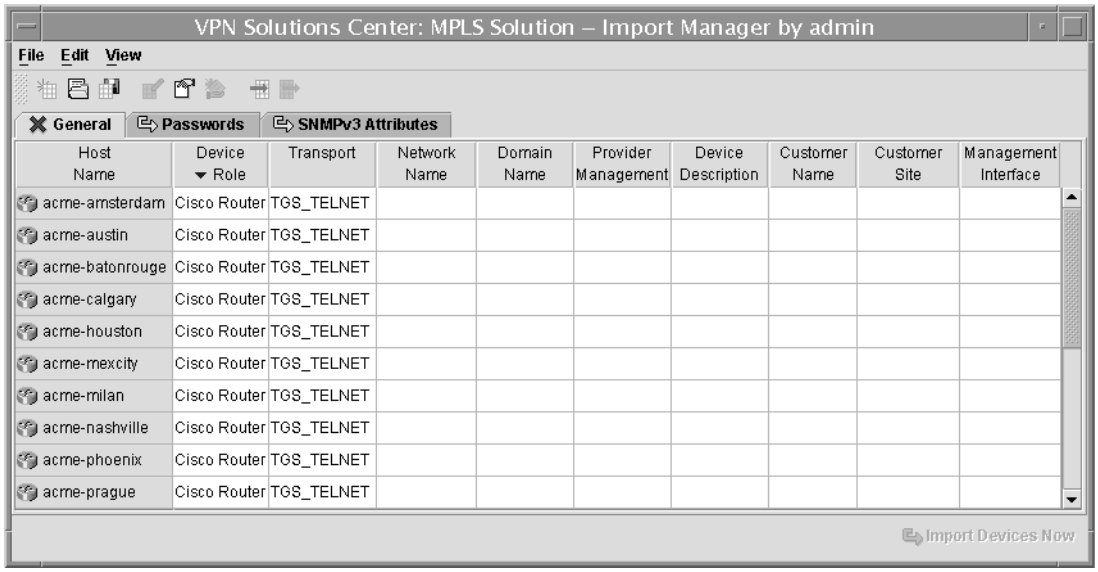

Observe the Import Manager dialog box as shown in Figure 4-53:

- The **General** tab has a red X in the tab title. A red X indicates that some or all of the parameters in the spreadsheets for those tabs must be filled in before you can import the devices into VPN Solutions Center. When these values are filled in, the red X changes into a green arrow, indicating that device description for that spreadsheet is complete and ready for import.
- The other two tabs—**Passwords** and **SNMPv3 Attributes**—display yellow arrows. This indicates that although there are some values that are not defined in those areas, they are considered ready for import.
- The **Import Devices Now** button (in the lower right corner) is temporarily disabled. When the required parameters for the CEs are specified, this button is enabled so that you can then import the PE configuration files into VPN Solutions Center.
- The *Device Role* for each device listed is already set to **Cisco Router**, and the *Transport Mechanism* for each CE router is set by default to **TGS_TELNET**. You can change the transport mechanism for all or some of the devices as necessary.

Specifying the Required Attributes for CEs

Most of the attributes for CEs are optional by default. However, you can specify which attributes are required for provisioning.

To specify the required attributes for CEs, follow these steps:

Step 1 From the Import Manager menu bar, choose **File > Specify Required Attributes**.

The Specify Required Attributes editor for MPLS CE devices appears (see Figure 4-54).

Figure 4-54 Specify Required Attributes Editor

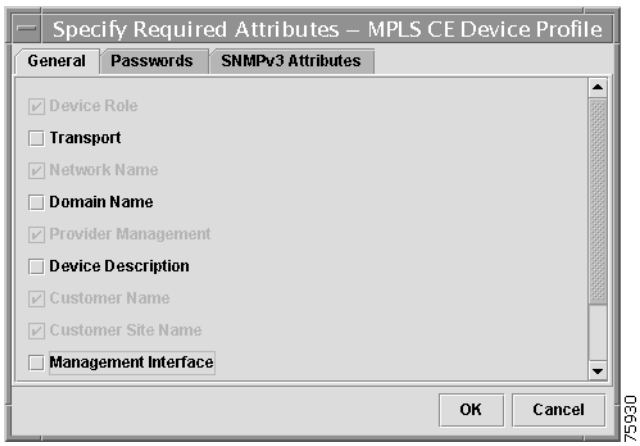

The Specify Required Attributes editor (for CEs) is organized into three tabs:

- *General*

  As shown in Figure 4-54, the General tab includes the following attributes: **Device Role**, **Transport**, **Network Name**, **Domain Name**, **Provider Management**, **Device Description**, **Customer Name**, **Customer Site Name**, and **Management Interface**.

  The **Device Role**, **Network Name**, **Provider Management**, **Customer Name**, and **Customer Site Name** attributes are required by default.

- *Passwords*
- *SNMPv3 Attributes*

Step 2 In the list of attributes in the General tab, select the checkboxes for those additional attributes that you want to be required.

Step 3 Click the **Passwords** tab.

The Specify Required Password Attributes dialog box appears (see Figure 4-55).

Figure 4-55 Specify Required CE Passwords

Step 4 In the list of attributes in the Passwords tab, select the checkboxes for those password attributes that you want to be required.

Step 5 Click the **SNMPv3 Attributes** tab.

The Specify Required SNMPv3 Attributes dialog box appears (see Figure 4-56).

Figure 4-56 Specify Required SNMPv3 Attributes

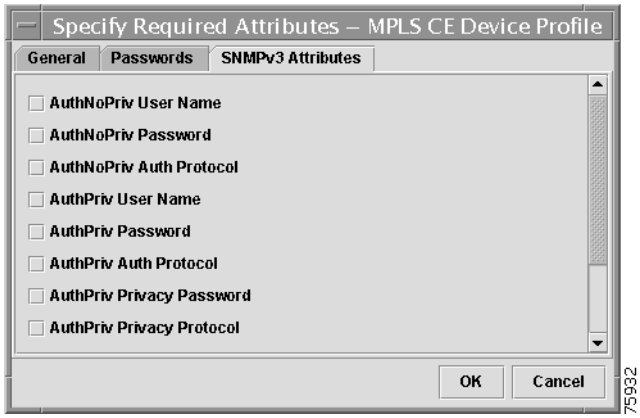

Step 6 In the list of attributes in the SNMPv3 Attributes tab, select the checkboxes for those SNMPv3 attributes that you want to be required.

Step 7 When satisfied with the settings for required CE attributes, click **OK**.

VPN Solutions Center will now require the selected required attributes to be defined in the Import Manager before the set of CE routers can be imported into the VPNSC Repository.

Defining a New VPN Customer Name

If the Customer name for some or all of the CE configuration files you want to import does not yet exist in VPNSC, you must define the Customer name from the Import Manager user interface.

To define a VPN Customer name, follow these steps:

Step 1 With the CEs displayed in the Import Manager dialog box, **double-click** the *Customer Name* cell in one of the CE device rows.

The Customer Name dialog box appears (see Figure 4-57).

Figure 4-57 Customer Name Dialog Box

Step 2 Choose **Create New Customer**, then click **OK**.

The following dialog box appears (see Figure 4-58).

Figure 4-58 Entering the New Customer Name

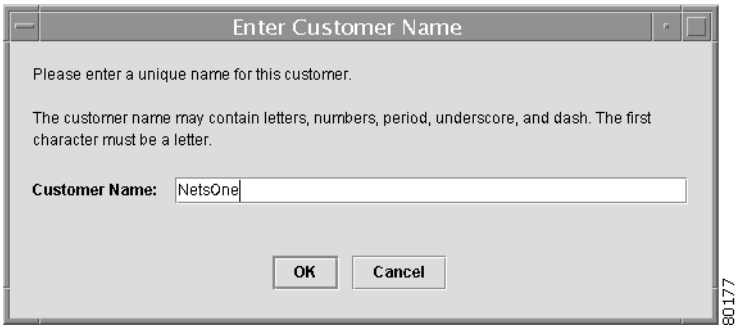

Step 3 Enter the name of the new VPN Customer, then click **OK**.

The new VPN Customer now exists in the Import Manager.

Apply the New Customer Name to the CEs

You can now apply the new Customer name to the appropriate CEs.

Step 1 From the *Customer Name* column in the Import Manager window, select the CEs that are associated with the new Customer.

Step 2 Choose **Edit > Edit Selected Devices**.

The Edit Parameters: Customer Name dialog box appears (see Figure 4-59).

Figure 4-59 Edit Parameters: Customer Name

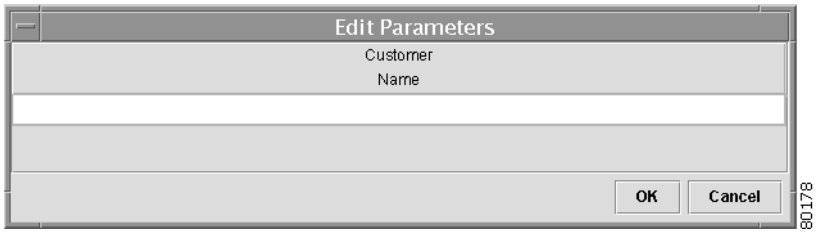

Step 3 Place the cursor anywhere in the empty field, then **double-click**.

The Customer Name dialog box appears. Notice that the Customer name that you created is now listed among the Customer names you can choose from (see Figure 4-60).

Figure 4-60 Choosing the New Customer Name

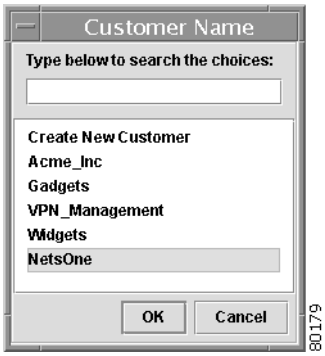

Step 4 Select the name of the new Customer, then click **OK**.

You return to the Edit Parameters dialog box, where the Customer name you selected is displayed.

Step 5 To apply the Customer name to the selected CE routers, click **OK**.

Specifying a Customer Site for Each CE

When Customer names have been assigned to the incoming CEs, you can specify the Customer site for each CE.

When specifying the Customer site, you have five options:

- Create a new site name
- Specify the site name in the format: *Customer Name + Site name + Host name*
- Specify the site name in the format: *Site name + Host name*
- Specify the site name in the format: *Host name + Site name*
- Select one of the existing sites already defined for the current Customer

Creating a New Site Name

If you want to create a site name that does not conform to one of the site name formats, follow these steps:

Step 1 From the Import Manager dialog box, **double-click** the *Customer Site* cell for the router of interest.

The following dialog box appears (see Figure 4-61).

Figure 4-61 Defining a New Customer Site Name

Step 2 Choose **Create New Site**, then click **OK**.

The following dialog box appears (see Figure 4-62).

Figure 4-62 Entering the New Customer Site Name

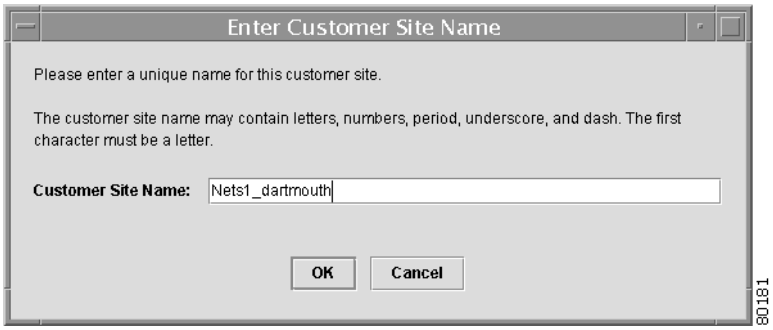

Step 3 Enter the name of the new site, then click **OK**.

The name of the new site is displayed in the *Customer Site* cell.

Specifying a Site Name Format

As explained above, you can specify the site name for one CE or any set of the CEs in one of three predefined formats:

- *Customer Name + Site name + Host name*
- *Site name + Host name*
- *Host name + Site name*

To specify the site name for CEs in a site name format, follow these steps:

Step 1 From the Import Manager dialog box, select all or any subset of the *Customer Site* cells.

Step 2 Choose **Edit > Edit Selected Devices**.

The Edit Parameters: Customer Site Name dialog box appears (see Figure 4-63).

Figure 4-63 Edit Parameters: Customer Site Name

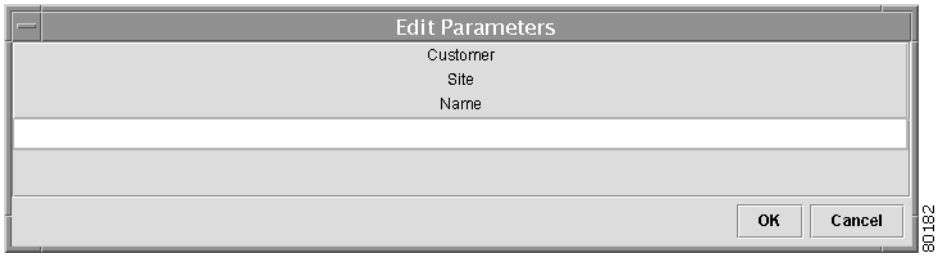

Step 3 Place the cursor anywhere in the empty field, then **double-click**.

The following dialog box appears (see Figure 4-64).

Figure 4-64 Specifying a Customer Site Name Format

Step 4 Select one of the Customer site name formats, then click **OK**.

You return to the Edit Parameters: Customer Site Name dialog box, where the selected site name format is displayed.

Step 5 To accept the site name format for the selected CEs, click **OK**.

The specified site name format is displayed in the Customer Site cells for the selected CEs.

Specifying the Management Status for CE Routers

The Provider Management attribute lets you specify the management status for the CE routers you are importing into VPN Solutions Center—that is, whether the CEs are managed or unmanaged devices. In addition, Customer Edge routers can use the SA Agent data collection utility in a number of different ways. Therefore, the Provider Management attribute lets you specify the SA Agent status for the CEs as well.

As shown in Figure 4-65, the provider management options for CEs are as follows:

- *Managed CE—No SA Agent*: Indicates that the CE is a managed CE that does not employ the SA Agent feature.
- *Managed CE—Regular SA Agent*: Indicates that the CE is a managed CE that has a dual function as a CE and an SA Agent router. That is, while functioning as a CE in the VPN, it is also monitoring traffic response times between CEs in the same VPN.
- *Managed CE—Shadow SA Agent*: Indicates that the designated CE is a managed CE that is actually a PE (in provider space) functioning as an SA Agent device.
- *Managed CE—Management LAN*: Defines a managed router in service provider space that functions as a Management CE (MCE) in a Management VPN.

  The network management subnet is connected to the Management CE (MCE). The MCE *emulates* the role of a CE, but the MCE is a router in provider space that serves as a network operations center gateway router. The MCE is part of a management site as defined in the VPN Solutions Center software.
- *Managed CE—Management LAN, SA Agent*: Defines a router in service provider space as an MCE in a Management VPN that is also functioning as an SA Agent device.
- *Multi-VRF CE—No SA Agent*: Indicates that the CE is a multi-VRF CE that is not collecting data through the SA Agent. For more information on Multi-VRF CEs, see Chapter 11, "Provisioning Multi-VRF CEs in VPN Solutions Center."

- *Multi-VRF CE—Regular SA Agent*: Indicates that the CE is a multi-VRF CE that is collecting data through the SA Agent.
- *Multi-VRF CE—Shadow SA Agent*: Indicates that the CE is a multi-VRF CE that is in provider space.
- *Unmanaged CE—No SA Agent*: Indicates that the CE is an unmanaged CE that does not employ the SA Agent feature.
- *Unmanaged CE—Regular SA Agent*: Indicates that the CE is an unmanaged CE that has a dual function as a CE and an SA Agent router. That is, while functioning as a CE in the VPN, it is also monitoring traffic response times between CEs in the same VPN.
- *Unmanaged CE—Shadow SA Agent*: Indicates that the designated CE is an unmanaged CE that is actually a PE (in provider space) functioning as an SA Agent device.
- *Unmanaged CE—Management LAN*: Defines an unmanaged router in service provider space that functions as a MCE in a Management VPN.
- *Unmanaged CE—Management LAN, SA Agent*: Defines an unmanaged router in service provider space that functions both as a MCE in a Management VPN and as an SA Agent device.

To specify the management status for the CE routers you are importing into your network, follow these steps:

Step 1 From the Import Manager: General tab, locate the Provider Management column, then select the set of *Provider Management* cells for the CEs of interest.

Step 2 Choose **Edit > Edit Selected Devices**.

The Edit Parameters: Provider Management dialog box appears.

Step 3 Place the cursor anywhere in the empty field, then **double-click**.

The following dialog box appears (see Figure 4-65).

Figure 4-65 Specifying the Type of Router Management

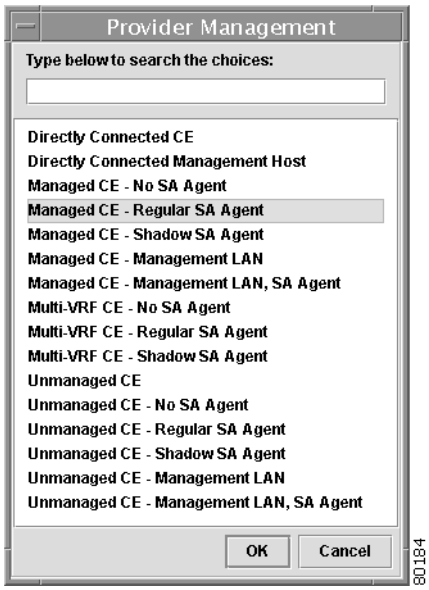

Step 4 Select the router management type of choice, then click **OK**.

You return to the Edit Parameters: Provider Management dialog box, where the management type you selected is displayed (see Figure 4-66).

Figure 4-66 Router Management Type Selected

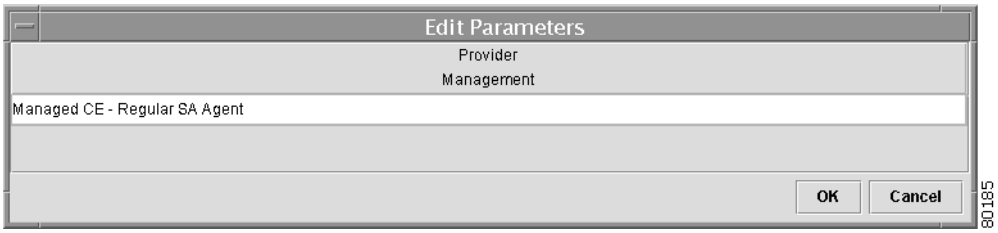

Step 5 If satisfied with your selection, click **OK**.

You return to the Import Manager screen, where the selected *Provider Management* cells display your router management selection.

Specifying the Management Interface for CE Routers

The VPN Solutions Center Network Management Subnet resides inside the service provider network, and communicates with edge routers through an assigned *management interface*. Configuration changes are managed by VPN Solutions Center software and transported to the appropriate edge routers through the management interface.

Tip When setting up network devices in VPN Solutions Center, be sure to have both the DNS-resolvable hostname and an IP address designated as the management interface specified for each device. If VPNSC cannot access a device via DNS, and no routable IP address is specified for the target devices in VPNSC, data collection operations will fail.

You have a choice of two methods when specifying the management interface for CEs:

- Select the management interface for a particular CE router.
- Assign a default management interface to all or a subset of the CEs.

Selecting the Management Interface for a Specific CE

To select the management interface for a specific CE, follow these steps:

Step 1 From the Import Manager: General tab, locate the row for the CE of interest, then select the *Management Interface* cell for that device (see Figure 4-67).

Figure 4-67 CEs Displayed in Import Manager Spreadsheet

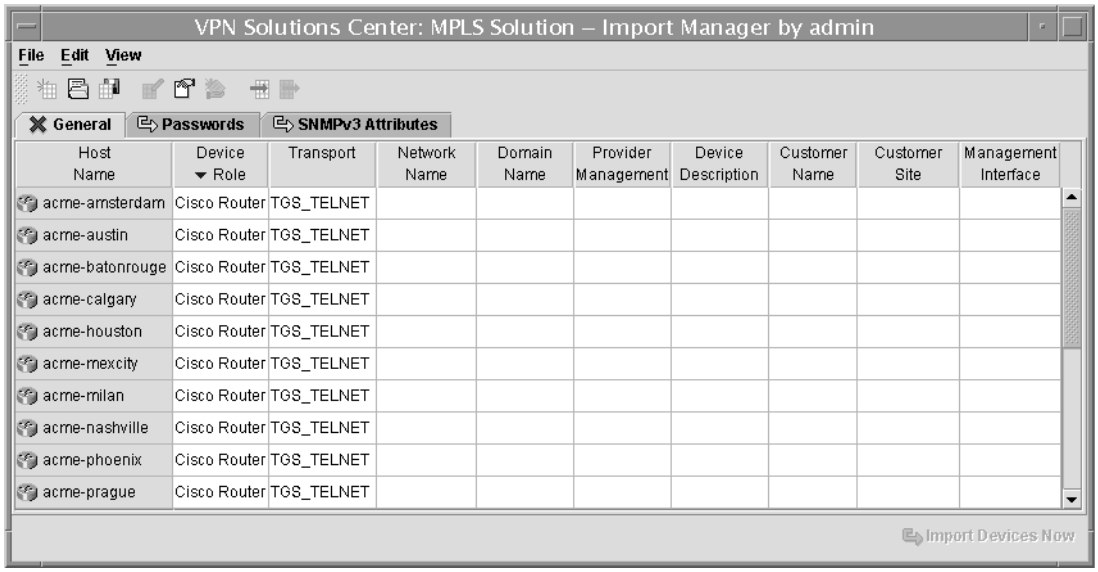

The Management Interface dialog box appears (see Figure 4-68).

Figure 4-68 Selecting the Management Interface

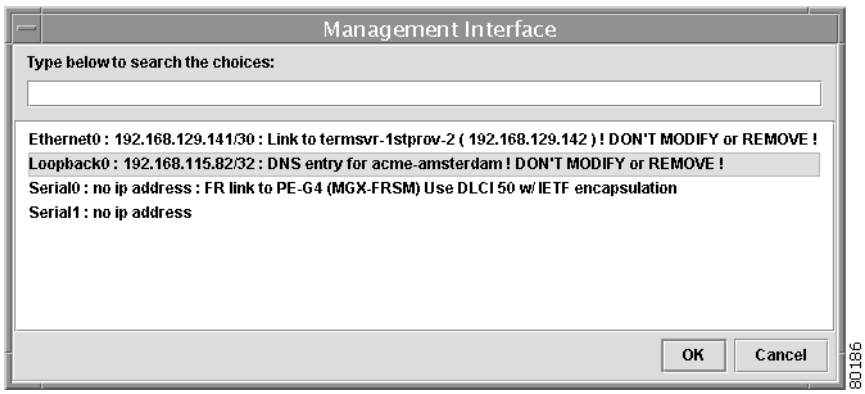

This dialog box displays the list of interfaces VPNSC parsed in the router's configuration file.

Step 2 Select the interface you want to designate as the CE's management interface, then click **OK**.

The interface you selected is assigned to be the router's management interface (see Figure 4-69).

Figure 4-69 Management Interface Specified in the Import Manager

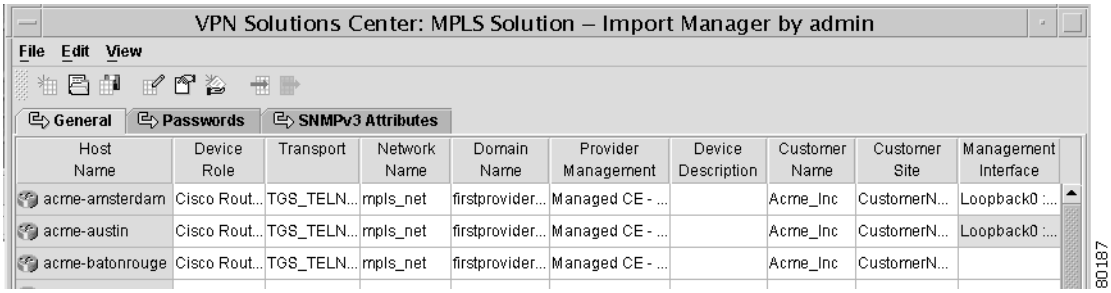

Assigning the Default Management Interface for a Group of CEs

This task assigns a default management interface to the set of CEs that you are importing. Doing so is possible only if the CEs all use the same type of interface(s); for example, if they all use Loopback0 or S0 for their management interface.

In addition, you can specify multiple interfaces: the interface entered first is the interface that VPNSC searches for and selects first. If the first interface is not found, VPNSC searches each CE in the set for the next interface specified in the list.

To set the default management interface for a group of CEs:

Step 1 From the Import Manager: General tab, locate the Management Interface column, then select the set of *Management Interface* cells for the CEs of interest (see Figure 4-67 on page 4-51).

Step 2 Choose **Edit > Edit Default Cell Values**.

The MPLS CE Default Values Editor appears (see Figure 4-70).

Figure 4-70 MPLS CE Default Values Editor: General Parameters

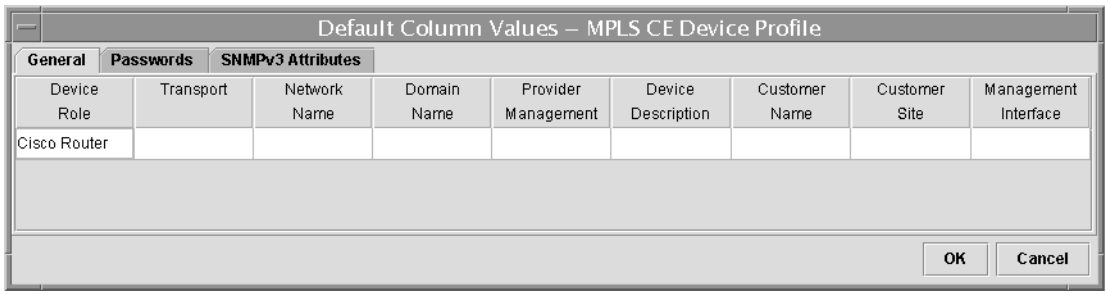

Step 3 Place the cursor in the *Management Interface* cell and **double-click**.

The following dialog box appears (see Figure 4-71).

Figure 4-71 Specifying the Default Management Interfaces

Step 4 Enter one or two default management interface(s) for the set of CEs.

Note Use a *semicolon* (**;**) to separate multiple interface names.

The interface entered first is the interface that VPNSC searches for and selects first. Using the example in Figure 4-71, VPNSC will search each of the PE devices for a Loopback0 interface. If a Loopback0 interface is found on a device, VPNSC assigns the management interface for that device to Loopback0.

If the first interface is not found, VPNSC searches each CE in the set for the next interface specified in the list; in our example, the next interface would be S0. In cases where S0 exists, but Loopback0 does not, VPNSC will assign the S0 interface as the management interface.

Step 5 Click **OK**.

The management interfaces you specified are displayed in the Default Editor's *Management Interface* cell (see Figure 4-72).

Figure 4-72 MPLS CE Default Values Editor: Management Interface Values

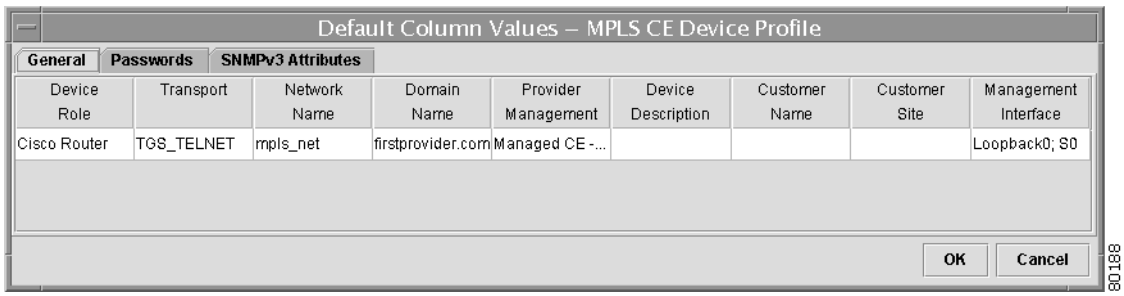

Step 6 Click **OK**.

You return to the Import Manager screen. The next step is to load the management interface values into the Import Manager.

Step 7 In the Import Manager screen, locate the Management Interface column, then select the set of *Management Interface* cells for the CEs of interest.

Step 8 With the appropriate column and routers selected, choose **Edit > Load Default Values to Selected Cells**.

The Import Manager loads the default management interfaces that you specified into the appropriate *Management Interface* cells, including the name of the interface and its IP address (see Figure 4-73).

Figure 4-73 Complete Management Interface Values Loaded

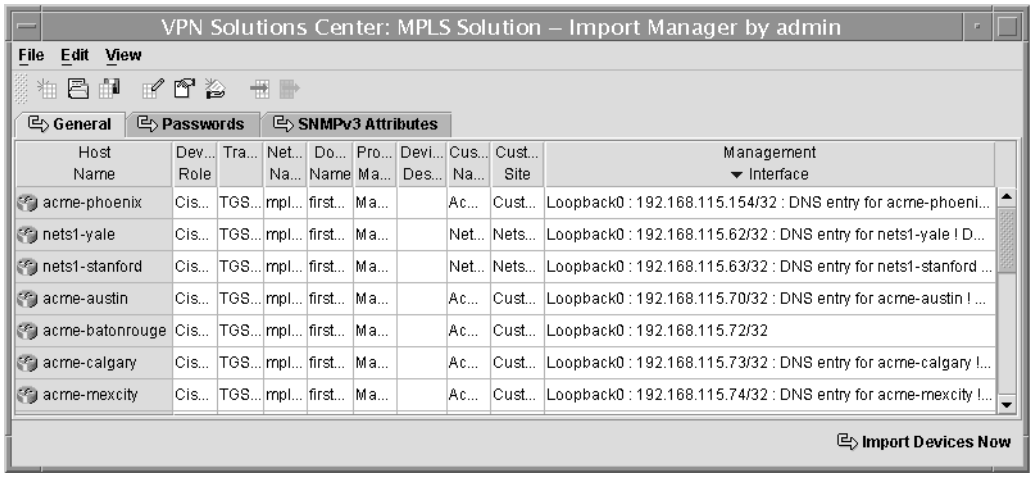

Specifying Password Parameters for CEs

Caution VPN Solutions Center requires that the PEs and managed CEs in the network have a login password (also called the *virtual terminal* password). Data collection operations fail if VPN Solutions Center does not find the login password set on routes it attempts to collect data from.

To specify the passwords for the set of CEs you are importing, follow these steps:

Step 1 From the MPLS CE Import Manager, choose the **Passwords** tab.

The Passwords dialog box appears (see Figure 4-74).

Figure 4-74 Password Parameters in the MPLS CE Import Manager

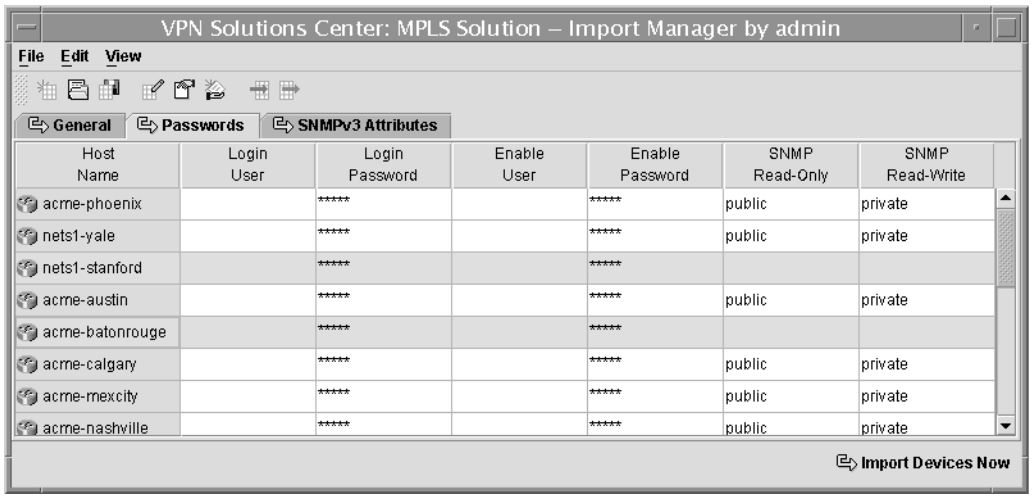

As you can see in this figure, some of the password attributes, such as the login password, enable password, and so on, were picked up from the routers' configuration files and populated into the appropriate columns in the dialog box.

The yellow arrow in the Passwords tab indicates that even though some attributes are not filled in, the passwords for the selected CEs could be imported without further additions or changes. However, if you wish to specify additional password attributes, proceed to the next step.

Step 2 Select the CE routers that you want to edit.

You can select all the devices in the window by choosing **Edit > Select All**.

Step 3 From the Import Manager menu bar, choose **Edit > Edit Selected Devices**.

The Edit Parameters: Passwords dialog box is displayed (see Figure 4-75).

Figure 4-75 Editing CE Passwords

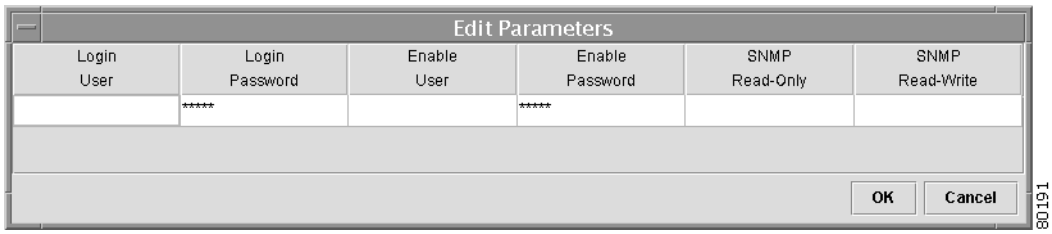

When You Leave Some Values Empty

You may choose to edit some parameters and leave some parameters empty. When you enter the items you want to change and click **OK**, the following prompt is displayed (see Figure 4-76).

Figure 4-76 Confirmation Prompt

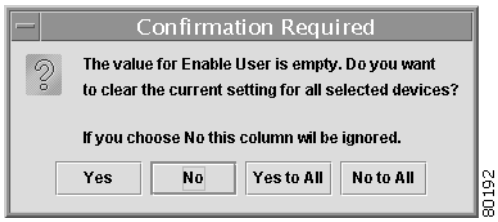

- When you click **No**, the Editor leaves the named column value untouched.
- When you click **Yes**, the Editor enters the changed items you specified in the Edit Parameters dialog and clears the values from all the other cells.

Login User

Step 1 To set the login username for the selected CE routers to be imported into VPNSC, **double-click** the *Login User* cell.

The Login User dialog box is displayed (see Figure 4-77).

Figure 4-77 Entering the CE Login Username

Step 2 Enter the login username here, then click **OK**.

Login Password

The login password is the router's virtual terminal password, which establishes password protection on incoming Telnet sessions.

On a router that is to be accessed by a terminal server, the login password and the console password should be identical. *A terminal server accesses a router via the router's console port.*

Step 1 To set the login password for the selected CE routers to be imported into VPNSC, **double-click** the *Login Password* cell. The Login Password dialog box is displayed (see Figure 4-78).

Figure 4-78 Entering the CE Login Password

Step 2 Enter the login password.

Step 3 In the *Verify Password* field, enter the login password again, then click **OK**.

Enable User

Step 1 To set the enable username for the selected CE routers to be imported into VPNSC, **double-click** the *Enable User* cell. The Enable User dialog box is displayed (see Figure 4-79).

Figure 4-79 Specifying the CE Enable Username

Step 2 Enter the enable username, then click **OK**.

Enable Password

If desired, set the enable password here for the selected CEs.

Step 1 To set the enable password for the selected CE routers to be imported into VPNSC, **double-click** the *Enable Password* cell. The Enable Password dialog box is displayed (see Figure 4-80).

Figure 4-80 Specifying the Enable Password

Step 2 Enter the enable password.

Step 3 In the *Verify Password* field, enter the enable password again, then click **OK**.

SNMP Read-Only Community String

The SNMP Read-Only community string is used to read MIB variables.

Note The SNMP community strings must be set on all the PEs and CEs in the service provider's network; the SNMP settings specified in VPN Solutions Center must match the SNMP string values configured for the routers. For related information, see the "Setting Up SNMPv1 and SNMPv2 on the Routers in the Service Provider Network" section on page 2-4 and the "Setting the SNMPv3 Parameters on the Routers in the Service Provider Network" section on page 2-5.

Step 1 To set the SNMP Read-Only community string for the selected CE routers to be imported into VPNSC, **double-click** the *SNMP Read-Only* cell (see Figure 4-81).

Figure 4-81 Specifying the SNMP Read-Only String

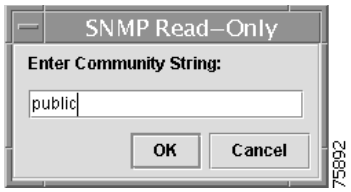

Step 2 Enter the SNMP Read-Only community string, then click **OK**.

SNMP Read-Write Community String

The SNMP Read-Write community string is used to set MIB variables.

Step 1 To set the SNMP Read-Write community string for the selected CE routers to be imported into VPNSC, **double-click** the *SNMP Read-Write* cell (see Figure 4-82).

Figure 4-82 Specifying the SNMP Read-Write String

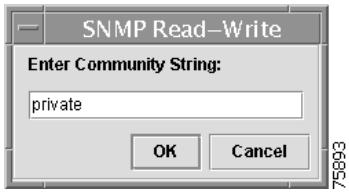

Step 2 Enter the SNMP Read-Write community string, then click **OK**.

Step 3 If satisfied with the default password settings, click **OK**.

Specifying the SNMPv3 Attributes for CEs

Simple Network Management Protocol Version 3 (SNMPv3) is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting packets over the network.

The values you set here must match the actual SNMPv3 values configured on the selected device (see "Setting the SNMPv3 Parameters on the Routers in the Service Provider Network" section on page 2-5).

To specify the SNMPv3 attributes for the set of CEs you are importing, follow these steps:

Step 1 From the MPLS CE Import Manager, choose the **SNMPv3 Attributes** tab.

The SNMPv3 Attributes dialog box appears (see Figure 4-83).

Figure 4-83 SNMPv3 Parameters in the MPLS CE Import Manager

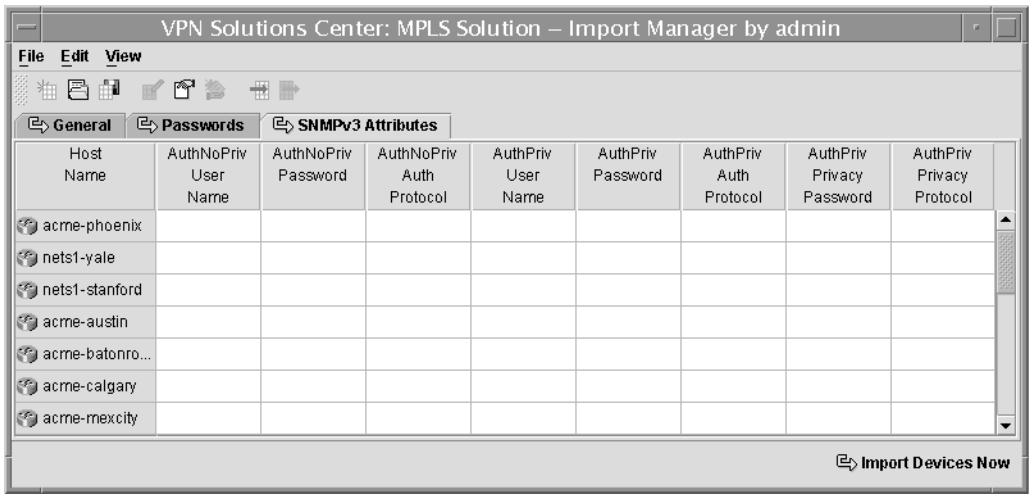

The yellow arrow in the Import Manager's SNMPv3 tab indicates that even though none of the SNMPv3 attributes are filled in, the selected CEs could be imported without further additions or changes to the SNMPv3 area. However, if you wish to specify additional default SNMPv3 attributes, proceed to the next step.

Step 2 Select the CE routers that you want to edit.

You can select all the devices in the window by choosing **Edit > Select All**.

Step 3 From the Import Manager menu bar, choose **Edit > Edit Selected Devices**.

The Edit Parameters: SNMPv3 dialog box is displayed (see Figure 4-84).

Figure 4-84 Editing CE SNMPv3 Parameters

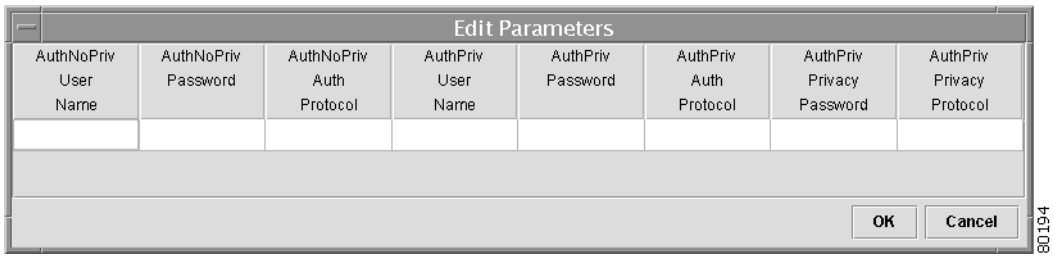

AuthNoPriv User Name

The AuthNoPriv user must have permission to the object identification numbers (OIDs) specified in the security request (that is, write permission for a set request, and read permission for a get request).

Step 1 To set the AuthNoPriv username for the selected CE routers to be imported into VPNSC, **double-click** the *AuthNoPriv User Name* cell.

The AuthNoPriv User Name dialog box is displayed.

Step 2 Enter the AuthNoPriv user name configured on the specified edge device routers, then click **OK**.

AuthNoPriv Password

Step 1 To set the AuthNoPriv password for the selected CE routers to be imported into VPNSC, **double-click** the *AuthNoPriv Password* cell.

The AuthNoPriv Password dialog box is displayed.

Step 2 *Password*: Enter the AuthNoPriv authentication password configured on the specified edge device routers.

Step 3 *Verify Password*: Enter the password again to verify it, then click **OK**.

AuthNoPriv Protocol

The available AuthNoPriv protocol options are **None**, **MD5**, or **SHA**.

- MD5: The MD5 algorithm is intended for digital signature applications, where a large file must be compressed in a secure manner before being encrypted with a private (secret) key under a public-key cryptosystem such as RSA. The MD5 algorithm takes as input a message of arbitrary length and produces as output a 128-bit *fingerprint* or *message digest* of the input.
- SHA: Secure Hash Algorithm. Computes a condensed representation of a message or a data file. When a message of any length is input, the SHA-1 produces a 160-bit output called a *message digest*. The message digest can then be input to the Digital Signature Algorithm (DSA), which generates or verifies the signature for the message. The creator of the digital signature and the verifier of the digital signature must use the same hash algorithm.

Step 1 To set the AuthNoPriv protocol for the selected CE routers to be imported into VPNSC, **double-click** the *AuthNoPriv Protocol* cell.

The AuthNoPriv Protocol dialog box is displayed (see Figure 4-85).

Figure 4-85 Specifying the Default Authentication Protocol

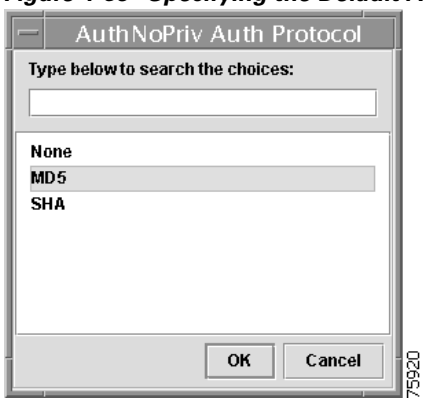

Step 2 Select the authentication protocol from the list, then click **OK**.

AuthPriv User Name

The AuthPriv user must have permission to the object identification numbers (OIDs) specified in the security request (that is, write permission for a set request, and read permission for a get request).

Step 1 To set the AuthPriv username for the selected CE routers to be imported into VPNSC, **double-click** the *AuthPriv User Name* cell.

The AuthPriv User Name dialog box is displayed.

Step 2 Enter the AuthPriv user name configured on the specified edge device routers, then click **OK**.

AuthPriv Password

Step 1 To set the AuthPriv password for the selected CE routers to be imported into VPNSC, **double-click** the *AuthPriv Password* cell.

The AuthPriv Password dialog box is displayed.

Step 2 *Password*: Enter the authentication password configured on the specified edge device routers.

Step 3 *Verify Password*: Enter the password again to verify it, then click **OK**.

AuthPriv Authentication Protocol

The available AuthPriv Authentication protocol options are **None**, **MD5**, or **SHA**.

- **MD5**: The MD5 algorithm is intended for digital signature applications, where a large file must be compressed in a secure manner before being encrypted with a private (secret) key under a public-key cryptosystem such as RSA. The MD5 algorithm takes as input a message of arbitrary length and produces as output a 128-bit *fingerprint* or *message digest* of the input.
- **SHA**: Secure Hash Algorithm. Computes a condensed representation of a message or a data file. When a message of any length is input, the SHA-1 produces a 160-bit output called a *message digest*. The message digest can then be input to the Digital Signature Algorithm (DSA), which generates or verifies the signature for the message. The creator of the digital signature and the verifier of the digital signature must use the same hash algorithm.

Step 1 To set the AuthPriv authentication protocol for the selected CE routers imported into VPNSC, **double-click** the *AuthPriv Auth Protocol* cell.

The AuthPriv Auth Protocol dialog box is displayed.

Step 2 Select the appropriate authentication protocol from the list, then click **OK**.

AuthPriv Privacy Password

The privacy password is the encryption password.

Step 1 To set the AuthPriv privacy password for the selected CE routers imported into VPNSC, **double-click** the *AuthPriv Password* cell.

The AuthPriv Privacy Password dialog box is displayed.

Step 2 *Password*: Enter the AuthPriv privacy (encryption) password configured on the specified edge device routers.

Step 3 *Verify Password*: Enter the encryption password again to verify it, then click **OK**.

AuthPriv Privacy Protocol

Currently, the only AuthPriv privacy protocol supported is **DES-56**.

Data Encryption Standard (DES) encrypts packet data. The Cisco IOS implements the mandatory 56-bit DES-CBC (Cipher Block Chaining) with the explicit initialization vector. Cipher Block Chaining requires an initialization vector to start encryption. The initialization vector is given in the IPsec packet. Triple DES (3DES) adds security by performing the operation three times with different subkeys.

Step 1 To set the AuthPriv privacy protocol for the selected CE routers to be imported into VPNSC, **double-click** the *AuthPriv Privacy Protocol* cell.

The AuthPriv Privacy Protocol dialog box is displayed.

Step 2 Select the **DES-56** authentication protocol, then click **OK**.

Step 3 When you are finished defining the AuthPriv privacy protocol for the selected devices, click **OK**.

You return to the Import Manager dialog box, where the values you specified are displayed.

If Ready, Import the CE Configuration Files

If you wish to specify any default values for the CEs you are importing, then proceed to the next section, "Specifying the Default Values for the CE Routers," and complete that task before importing the CE configuration files.

If you have completed specifying the required attributes for CEs, you may choose to import the CE configuration files into the VPN Solutions Center Repository now.

When you are finished entering all the required attributes, the Import Devices Now button (in the bottom right corner of the Import Manager dialog box) is enabled (see Figure 4-86).

Figure 4-86 Import Now Button Enabled

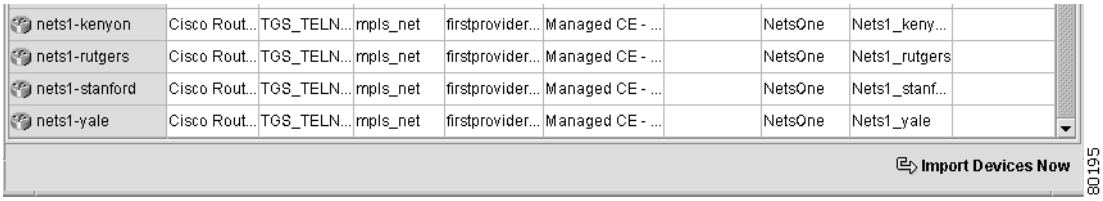

– To import the CEs into VPNSC, click **Import Devices Now**.

The selected configuration files and the additional information specified in this import procedure are imported into VPN Solutions Center Repository.

The new Customers and sites are added to the VPN Console's hierarchy pane.

Specifying the Default Values for the CE Routers

You are now ready to specify any desired default values for the CE parameters. Once you specify the default values, you can apply them to all or any subset of the imported CE devices.

Specifying default values is not a requirement—it is provided as a convenience to the Network Administrator. You always have the option of defining the parameters and attributes of the selected devices without specifying default values.

Step 1 From the Import Manager, choose **Edit > Edit Default Cell Values**.

The MPLS CE Default Values Editor appears (see Figure 4-87).

Figure 4-87 MPLS CE Default Values Editor: General Parameters

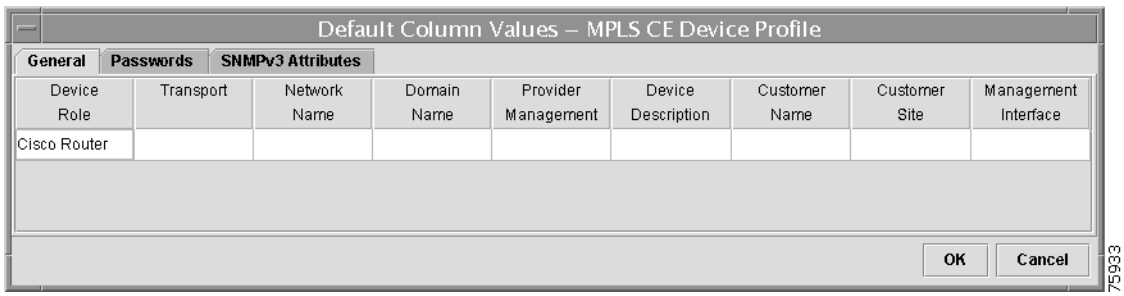

Step 2 To specify the default value for any value, **double-click** the appropriate cell. A value editor is displayed.

Step 3 Select or enter the appropriate default value for each parameter.

Step 4 When you have defined each parameter to your satisfaction, click **OK**, then proceed to the next tab and define the next set of attributes as needed.

Step 5 When all the default attributes are assigned, you will then import the default values into the Import Manager (see the "Importing Default Values Into the Import Manager" section on page 4-67).

Specifying the Default General Parameters for CEs

As you can see in Figure 4-87, you can set the default values for all the selected devices for the following general parameters:

- Device Role (set by default to *Cisco Router*)
- Transport (set by default to *TGS_TELNET*)
- Network Name
- Domain Name
- Provider Management

  See the "Specifying the Management Status for CE Routers" section on page 4-48.
- Device Description (optional)
- Customer Name

  See the "Defining a New VPN Customer Name" section on page 4-44.
- Customer Site

  See the "Specifying a Customer Site for Each CE" section on page 4-46.
- Management Interface

  See the "Assigning the Default Management Interface for a Group of CEs" section on page 4-52.

Default Device Role

When importing CEs into VPNSC, the device role is set by default to **Cisco Router**.

Default Transport Mechanism

The **Transport** parameter configures the method of communication between the VPN Solutions Center workstation and the specified CE routers. The default transport mechanism for MPLS operations is **TGS_TELNET**.

Step 1 To change the default transport mechanism, **double-click** the *Transport* cell.

The Transport dialog box is displayed (see Figure 4-88).

Figure 4-88 Specifying the Default Transport Mechanism

Step 2 From the list of transport mechanisms, choose the configuration file transport method you are using.

- *TGS_Telnet*: The *TGS_TELNET* option is the default transport method for MPLS VPNs.
- *TGS_SSH*: The configuration file transport method for VPN Solutions Center IPsec mode is *TGS_SSH* (Telnet Gateway Server—Secure Shell).
- *TGS_TFTP*: If you choose TGS_TFTP as the default transport method, be sure to enable TFTP (Trivial File Transfer Protocol) on the VPN Solutions Center workstation and on the target routers.

  For details, see the "Enabling TFTP in VPN Solutions Center" section on page 2-7.

Step 3 Click **OK**.

You return to the MPLS CE Default Values Editor, where the selected default transport method is now displayed in the *Transport* cell.

Default Network Name

Step 1 To set the default network for all the CE routers imported into VPNSC, **double-click** the *Network Name* cell.

The Network Name dialog box is displayed (see Figure 4-89).

Figure 4-89 Specifying the Default Network

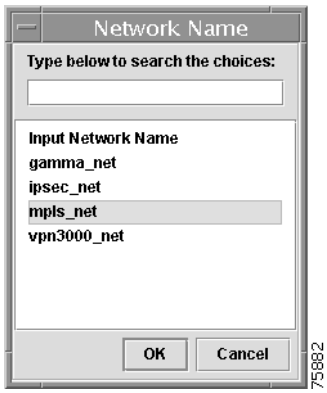

Step 2 From the list of networks, you can either choose the name of one of the networks listed or enter a new network name.

Step 3 Click **OK**.

The network name you specified is displayed in the Default Editor's *Network Name* cell.

Default Domain Name

If you do not specify the domain server name, VPN Solutions Center uses the target name; then, using, DNS (Domain Name System), it performs a domain lookup.

Step 1 To set the default domain for all the CE routers imported into VPNSC, **double-click** the *Domain Name* cell.

The Domain Name dialog box is displayed (see Figure 4-90).

Figure 4-90 Specifying the Default Domain

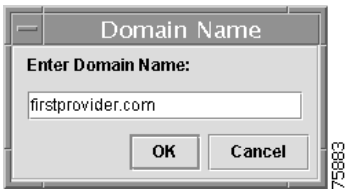

Step 2 Enter the default domain name for the selected CEs.

Step 3 Click **OK**.

The domain name you specified is displayed in the Default Editor's *Domain Name* cell.

Default Provider Management

The Provider Management attribute lets you specify the precise management status—managed CE or unmanaged CE—and the SA Agent status of the CEs in your network. For details on the available router management options, see the "Specifying the Management Status for CE Routers" section on page 4-48.

Step 1 **Double-click** the *Provider Management* cell.

The Provider Management dialog box is displayed (see Figure 4-91).

Figure 4-91 Specifying the Default Router Management Type

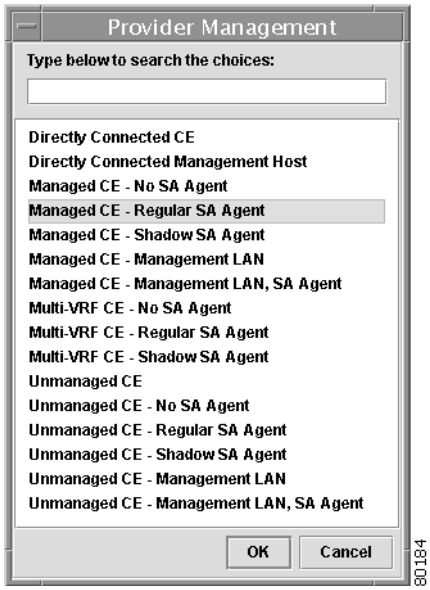

Step 2 Select the appropriate management type for the selected routers, then click **OK**.

Default Device Description

A default description of all the imported CEs may or may not be useful. For this reason, this field is optional. You can enter up to 256 characters in a device description.

If you do wish to enter a default device description for the imported CEs, **double-click** the *Device Description* cell, enter the description in the **Device Description** dialog box, then click **OK**.

Importing Default Values Into the Import Manager

When you have completed defining the default attributes for a set of devices, you must first import those default values into the Import Manager. The default values for each tab are loaded separately. For example, if you defined default values for the General attributes and the Provider attributes, you need to load the default values into both areas.

When all the necessary values are specified in the Import Manager, you can then import the configuration files into the VPN Solutions Center Repository.

To import the default values and import the configuration files, follow these steps:

Step 1 From the MPLS CE Import Manager, choose a tab category for which you created default values.

Step 2 Choose **Edit > Select All**.

This operation selects all the devices listed in the Import Manager dialog box.

Tip If you wish to load the default values to specific devices displayed in the Import Manager dialog box (not all the listed devices), place the cursor on the appropriate host name icon, then press **Ctrl+Click**. Repeat this as necessary to select the desired devices.

Step 3 With the appropriate devices selected, choose **Edit > Load Default Values to Selected Cells**.

The Import Manager loads the default values that you specified into the appropriate cells (see Figure 4-92).

Figure 4-92 Default CE Values Loaded into the Import Manager

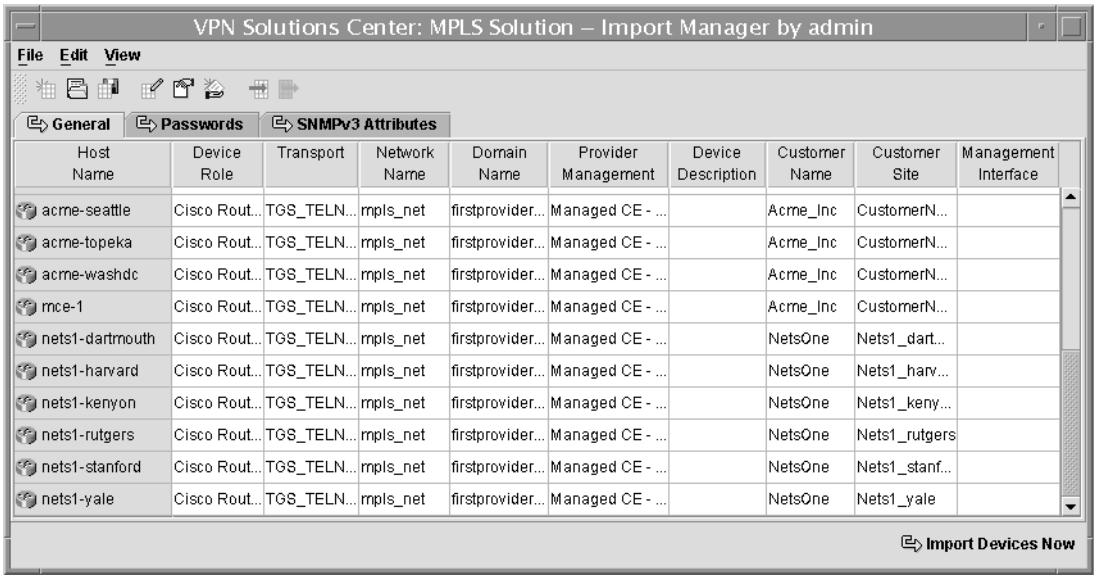

Step 4 Repeat this procedure for each tab category for which you created default values.

Importing the CE Configuration Files Into the Repository

When you are finished entering all the required attributes, the Import Devices Now button (in the bottom right corner of the Import Manager dialog box) is enabled (see Figure 4-93).

Figure 4-93 Import Now Button Enabled

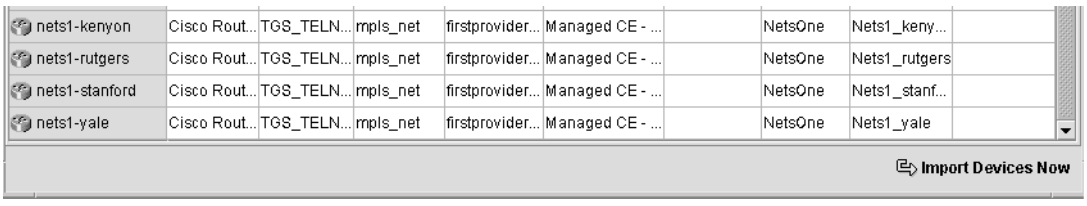

You are now ready to import the CE devices into the VPN Solutions Center Repository.

- To import the CEs into VPNSC, click **Import Devices Now**.

The selected configuration files and the additional information specified in this import procedure are imported into VPN Solutions Center Repository.

The new Customers and sites are added to the VPN Console's hierarchy pane.

What's Next?

Once you have successfully imported the PEs and CEs into VPN Solutions Center, the Network Administrator can proceed to define the necessary VPNs (if they have not already been defined in VPNSC). For details, see the "Defining a New VPN in the VPNSC Software" section on page 5-2.

Before Network Operators can provision VPN service requests, the Network Administrator must create and organize Service Request Profiles:

- See the "Creating Service Request Profiles" section on page 5-9.
- See the "Administering Service Request Profiles" section on page 5-14.

If you are ready to provision service requests, refer to Chapter 6, "Provisioning MPLS VPN Service Requests."

The sections below address various configuration file maintenance and administrative tasks:

- Editing a Device's Configuration File, page 4-70
- About the Download and Version Console, page 4-72
- Downloading a Previous Version of a Configuration File, page 4-73
- Using the Download Console, page 4-76
- Running IOS Commands from the VPN Console, page 4-79

Editing a Device's Configuration File

VPN Solution Center provides a mechanism for you to edit a configuration file that has been imported into VPNSC. Once you edit the configuration file, you can use the Download Console to download the edited file to the appropriate device.

To edit an imported configuration file, follow these steps:

Step 1 From the VPN Console hierarchy pane, expand the Device Inventory folder until you can see the list of networks.

Step 2 Select the name of the pertinent network, then **double-click**.

The Network window appears, displaying the list of devices in the selected network.

Step 3 Select the name of the device whose configuration file you want to edit.

Step 4 From the Network window, choose **Actions > Edit Latest Configuration File**.

The Configuration File Editor window appears, displaying the most recent version of the configuration file for the selected device (see Figure 4-94).

Figure 4-94 Configuration File Editor

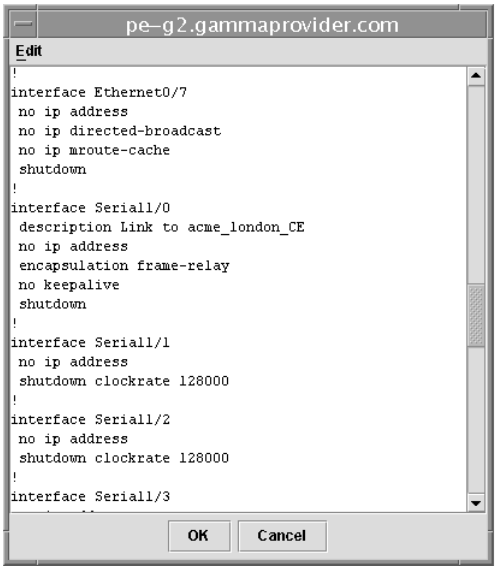

Step 5 Make the changes you wish to enter into the configuration file.

The Edit menu provides the following standard editing options:

- **Cut (Ctrl-X)**
- **Copy (Ctrl-C)**
- **Paste (Ctrl-P)**
- **Import**

  The **Import** option allows you to import a text file into the current configuration file.

Step 6 When finished editing the file, click **OK**.

VPNSC saves your changes to the Repository.

Updating the VPNSC Device Definition

If you make changes in a device's configuration file that are different from the information entered into VPN Solution Center's device definition, you must update the VPNSC device definition with the changes.

Step 1 In the VPN Console, open the Networks folder, then select the pertinent network.

Step 2 **Double-click** the selected network.

The Network window appears, displaying the names and roles of the devices in the network.

Step 3 From the Network window, select the name of the device whose configuration file you edited.

Step 4 Choose **Actions > Edit Target**.

The Edit Target dialog box appears (see Figure 4-95).

Figure 4-95 The Edit Target Dialog Box

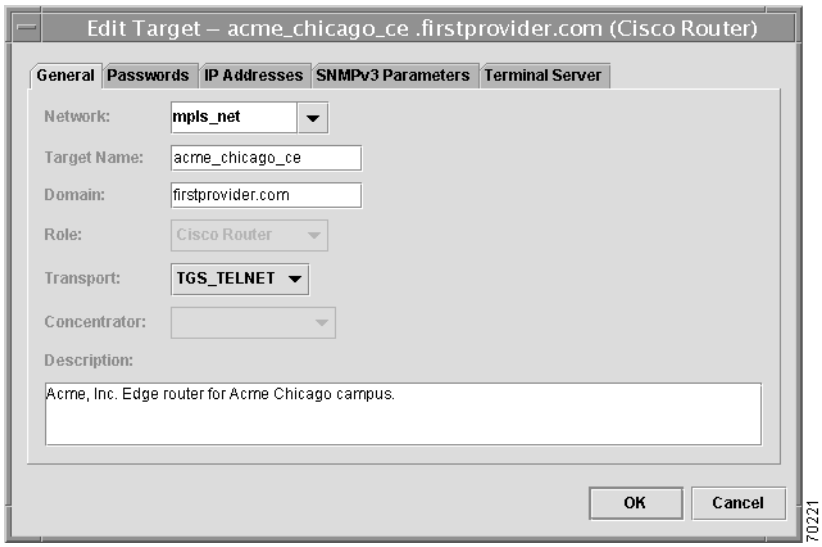

Step 5 Update the corresponding device information as necessary so that it matches the changes you made in the configuration file (for example, if you added, removed, or modified an IP address), then click **OK**.

Step 6 Download the updated configuration file to the target device as described in the "Downloading a Previous Version of a Configuration File" section on page 4-73.

About the Download and Version Console

VPN Solutions Center provides a mechanism called the *Download and Version Console* that allows you to download configuration files—or any set of IOS commands—to one or more routers.

- The *Download Console* adds the IOS commands to the router's existing configuration file. You can choose to download a configuration file or a set of commands to the router's running configuration or to the router's startup configuration.
- The *Version Console* provides a way to retrieve all the previous configuration version of a selected device. From the Version Console, you can view the contents of any version of a configuration file, download a version to the device, or copy any portion of the configuration version to another file.

To bring up the Download and Version Console:

From the VPN Console menu bar, choose **Tools > Download and Version Console**.

The Download Console appears (see Figure 4-96).

Figure 4-96 The Download Console

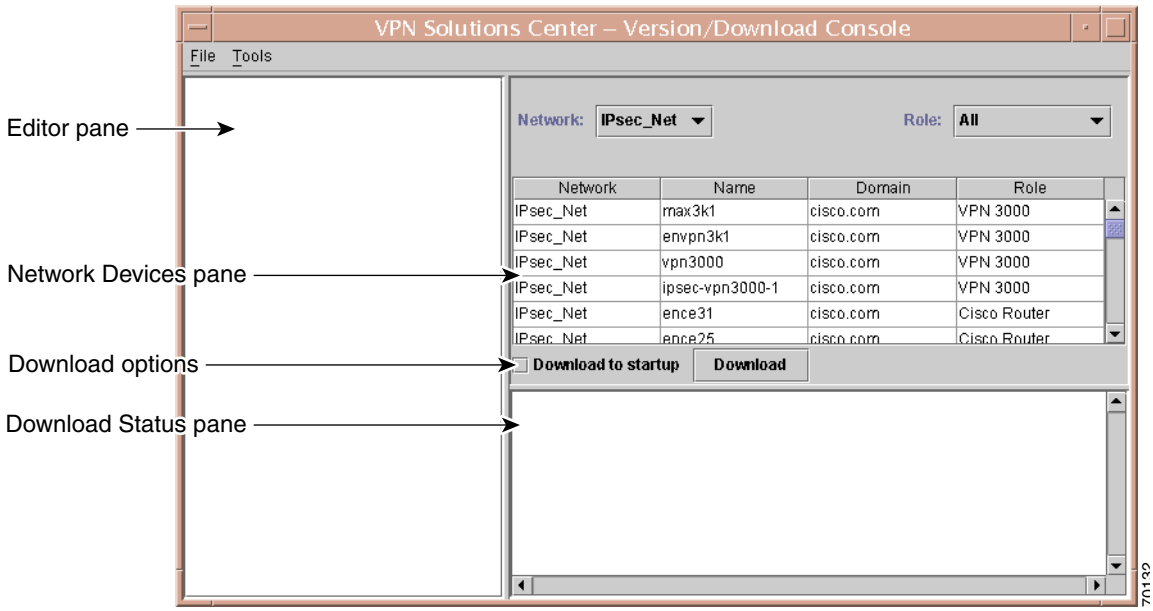

As shown in Figure 4-96, the Download Console consists of three major functional areas:

- *Editor pane*

  In this pane, you can enter commands directly or import commands from a file. You can then modify the text displayed in the Editor pane as necessary by using the standard keyboard commands to cut (**Ctrl-X**), copy (**Ctrl-C**), or paste (**Ctrl-V**) the text.

- *Network Devices pane*

  In this pane, you specify the network and role to view and select the network devices you want to download commands to.

- *Download Status pane*

  In this pane, the Download Console displays the status of each of the commands you download to the selected devices.

Downloading a Previous Version of a Configuration File

The following procedure describes a typical scenario in which you need to download a previous version of a device's configuration file to that device. To perform that operation, you must use the Version Console to retrieve the desired previous version of the configuration file, and then download the version of the configuration file to the device. You can download the file to either the device's running configuration or startup configuration.

To retrieve a previous version of a configuration file stored on the VPN Solutions Center workstation, follow these steps:

Step 1 Invoke the Download Console by choosing choose **Tools > Download and Version Console**.

Step 2 From the Download Console menu bar, choose **Tools > Version**.

The Version Console appears (see Figure 4-97).

Figure 4-97 The Version Console

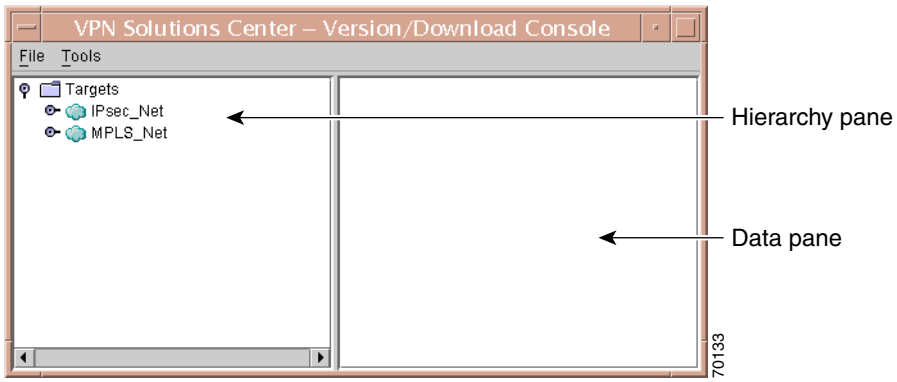

You can return to the Download Console at any time by choosing **Tools > Download**.

In the hierarchy pane on the left, the Version Console initially displays the list of existing networks. It organizes the configuration files by networks and their associated devices.

Step 3 Expand the Version Console hierarchy until you can see the router icons and the names of the routers in the pertinent network.

Step 4 Select the router icon for the router that contains the configuration file of interest, then **double-click**.

As shown in Figure 4-98, the Version Console displays a list of all the versions of the router's configuration file.

The versions are displayed according to the dates and times the configuration files were collected, and organized in descending order with the most recent version listed first, the next most recent version listed second, and so on.

Figure 4-98 Versions of a Configuration File Listed

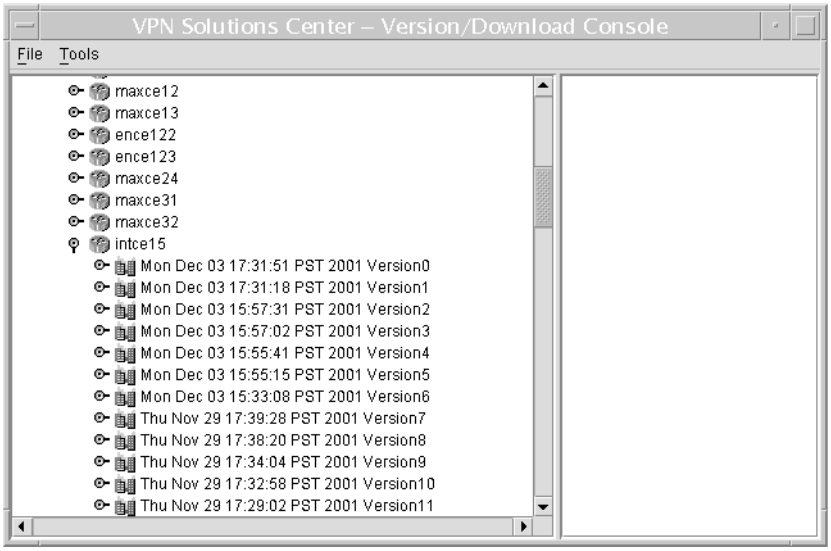

Step 5 Select the version of the configuration file that you're interested in.

Step 6 With the configuration file version of interest selected, **right-click**.

The menu shown in Figure 4-99 appears. From this menu, you can do any of the following operations:

- Open the selected file to view it.
- Download the selected file to the router's running configuration.
- Download the selected file to the router's startup configuration

Figure 4-99 Configuration File Version Menu

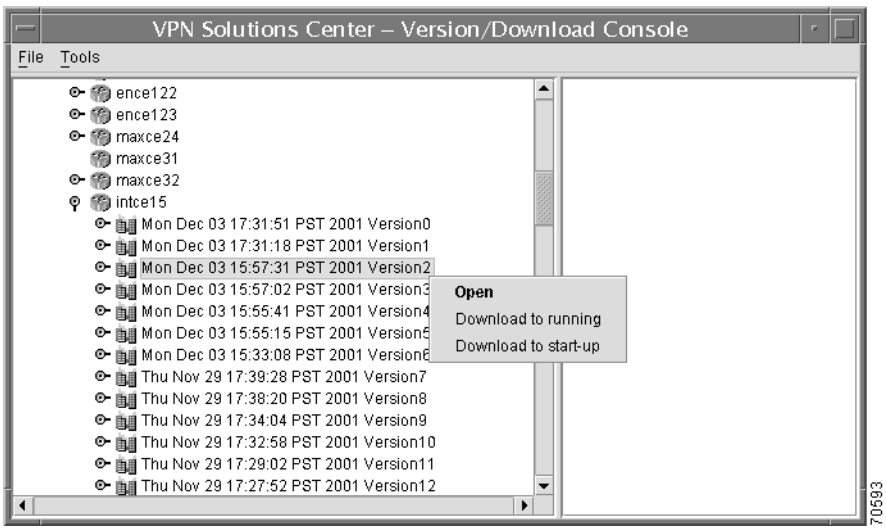

Viewing the Contents of a Configuration File Version

Step 7 To view the contents of the configuration file version, choose **Open**.

The configuration file contents are displayed in the data pane (see Figure 4-100).

Figure 4-100 Contents of a Configuration File Displayed

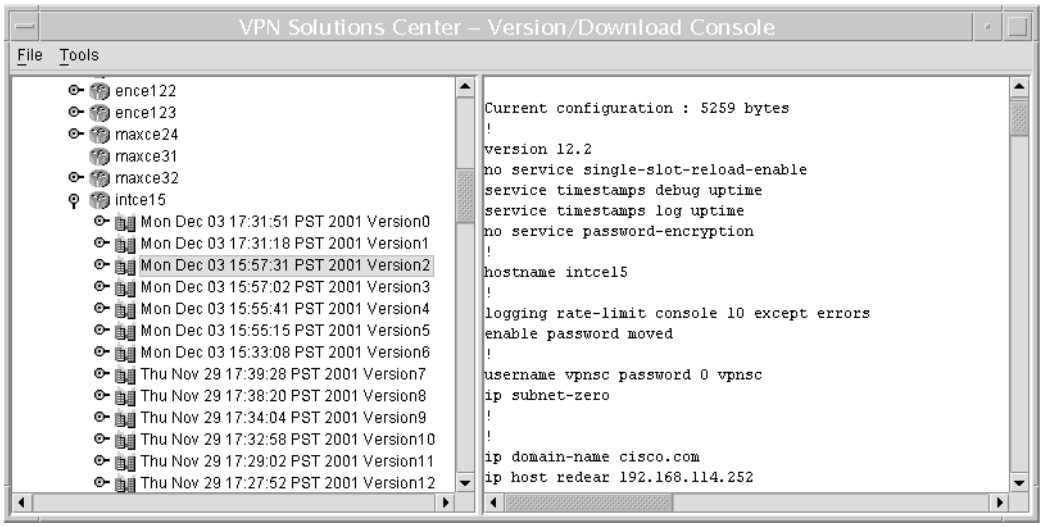

Downloading the Selected Version

You can download the selected version of the configuration file to either the router's running configuration or the startup configuration.

Of course, you should take care whenever you download a configuration to startup to make sure that you want to replace the current startup configuration.

Note Downloading a configuration file to the startup configuration uses TFTP as the transport mechanism. Make sure that TFTP is configured properly on the router, the TFTP server, and VPN Solutions Center before downloading to startup.

Step 8 To download the configuration file version to the router, select the version, then **right-click**.

Step 9 From the menu, choose as follows:

- To download the selected version of the file to the router's running configuration, choose **Download to running**.
- To download the selected version of the file to the router's startup configuration, choose **Download to start-up**.

You receive a confirmation message, asking if you want to proceed.

Step 10 Click **Yes** to proceed with the download operation.

Using the Download Console

You can use the Download Console to perform the following tasks:

- Download commands to one or more devices.
- Import a text file or a configuration file to the Download Console to download to one or more devices.

Downloading Commands to Multiple Devices

The Download Console provides a mechanism to download device commands (for Cisco IOS devices or non-IOS devices such as the Cisco VPN 3000 concentrator) to one or more devices at once.

VPN Solutions Center downloads the commands to the selected devices. These commands are added to the existing configuration on the selected devices.

To download commands to multiple devices, follow these steps:

Step 1 From the VPN Console menu bar, choose **Tools > Download and Version Console**.

The Download Console appears (see Figure 4-101).

Figure 4-101 The Download Console

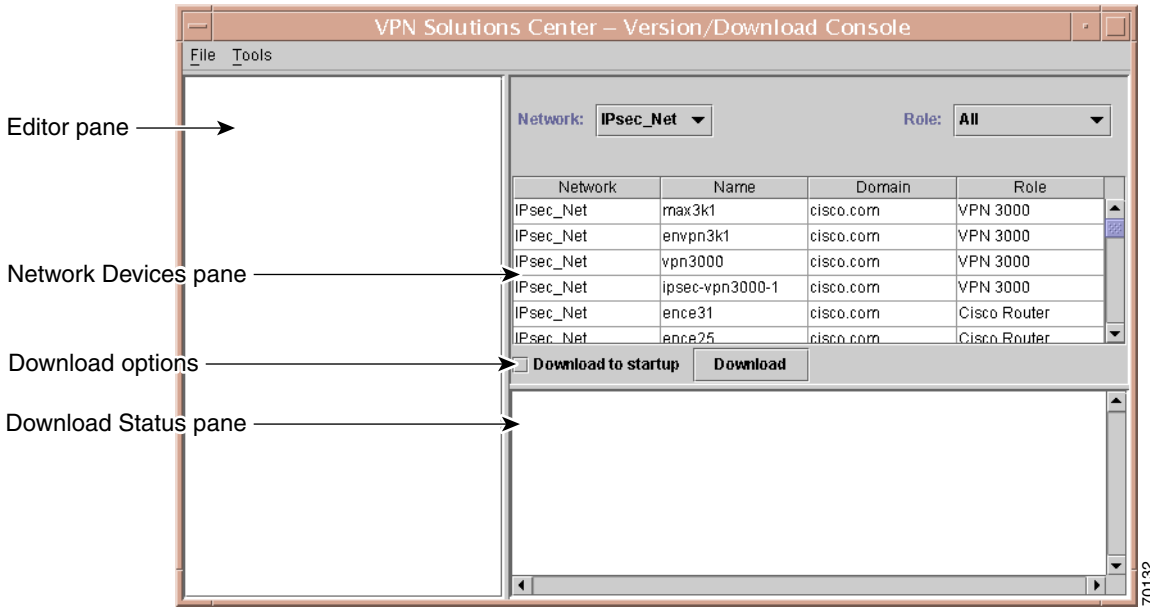

Step 2 Place the cursor into the Editor pane.

Step 3 Enter one or more commands (one command per line).

Step 4 In the Network Devices pane, select the devices you want to download the commands to (see Figure 4-102).

Figure 4-102 Commands Entered and Devices Selected

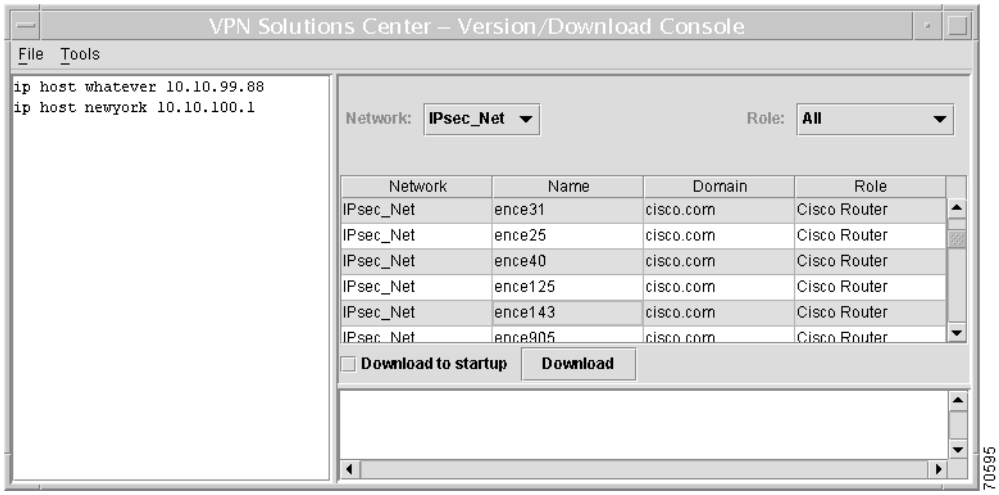

You can edit the text displayed in the Editor pane as necessary by using the standard keyboard commands to cut (**Ctrl-X**), copy (**Ctrl-C**), or paste (**Ctrl-V**) the text.

Step 5 Specify whether you want to download the commands to the target routers' running configuration or the startup configuration:

- To download the commands to the *running configuration* on each of the selected devices, click **Download**.
- To download the commands to the *startup configuration* on each of the selected devices, check the **Download to startup** checkbox, then click **Download** (see Figure 4-103).

Figure 4-103 Choosing the Download to Startup Option

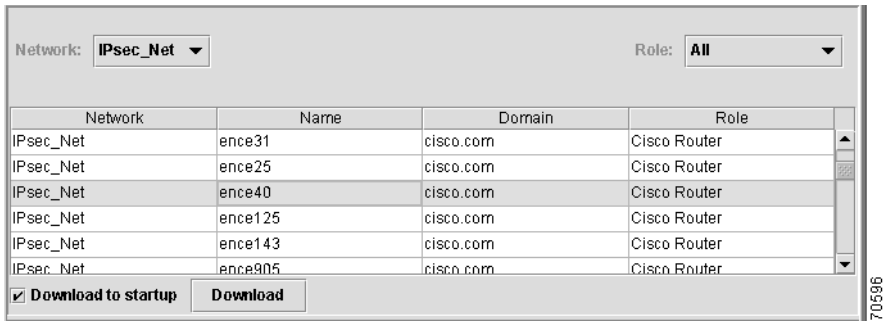

VPN Solutions Center downloads the commands displayed in the Editor pane to the selected devices. These commands are added to the existing configuration on the selected devices.

The Download Status pane (on the lower right) displays the status of each of the commands downloaded to the devices (see Figure 4-104).

Figure 4-104 Status of Downloaded Commands Displayed

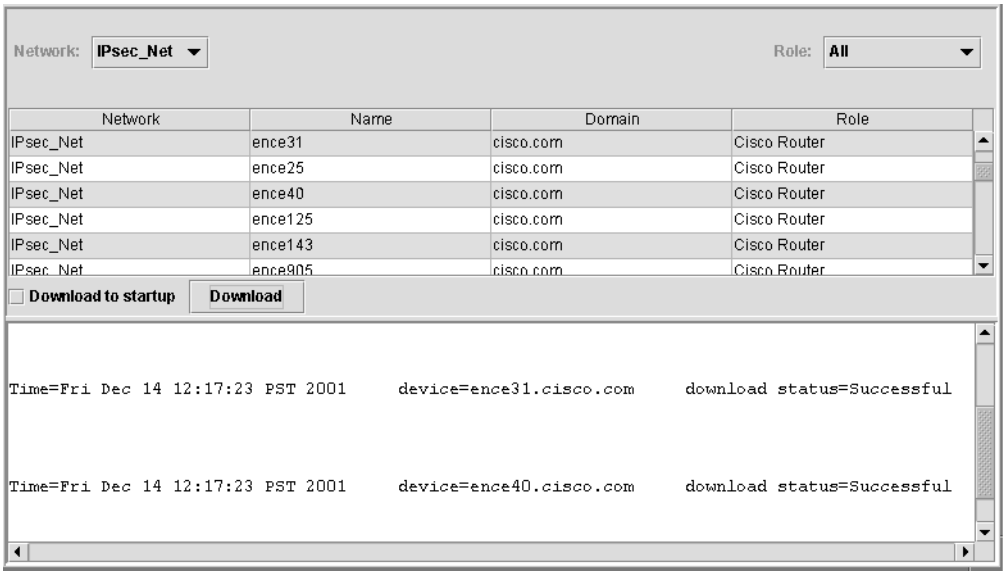

Use the scroll bar in the Download Status pane to view the status information that the routers return in response to the downloaded commands.

Step 6 To exit from the Download Console, choose **File > Exit**.

Importing a Text File or Configuration File to the Download Console

You can import a text file (consisting of commands) or a configuration file to the Download Console. You can then modify the file as needed and download the contents of the file to one or more devices.

To import a text file or configuration file to the Download Console, follow these steps:

Step 1 Bring up the Download Console by choosing **Tools > Download and Version Console** from the VPN Console menu bar.

Step 2 From the Download Console menu (shown in Figure 4-105), choose **File > Import**.

Figure 4-105 Download Console Menu

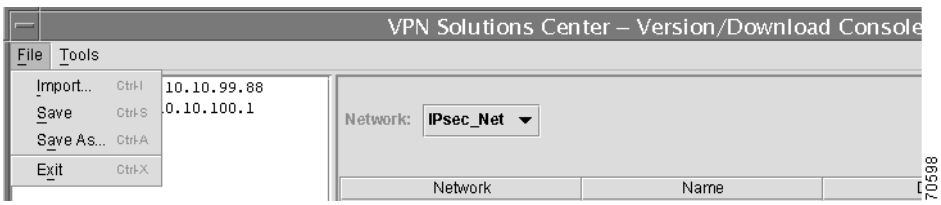

The Open dialog box appears.

Step 3 In the Open dialog box, navigate to the location of the text file or configuration file that you want to import.

Step 4 Select the name of the file, then click **Open**.

The contents of the selected file is displayed in the Editor pane.

Step 5 If necessary, edit the file to make it ready for downloading to a router.

You can edit the text displayed in the Editor pane as necessary by using the standard keyboard commands to cut (**Ctrl-X**), copy (**Ctrl-C**), or paste (**Ctrl-V**) the text.

Step 6 In the Network Devices pane, select the target devices.

Step 7 Specify whether you want to download the commands to the target routers' running configuration or the startup configuration:

- To download the commands to the *running configuration* on each of the selected devices, click **Download**.
- To download the commands to the *startup configuration* on each of the selected devices, check the **Download to startup** checkbox, then click **Download** (see Figure 4-106).

Figure 4-106 Choosing the Download to Startup Option

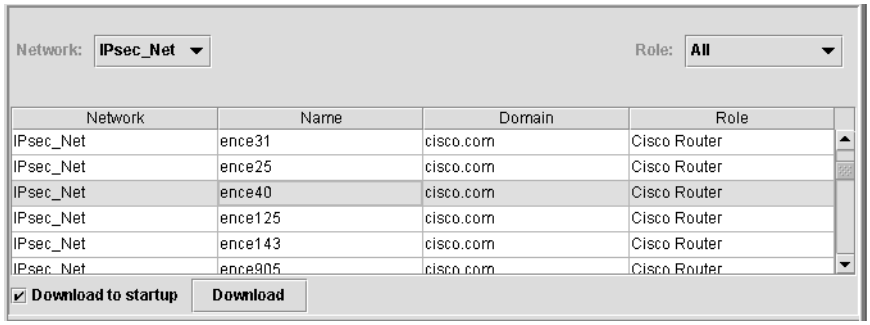

The Download Status pane displays the status of each of the commands downloaded to the devices.

Step 8 If you want to save the contents of the text file, you have two options:

- If you want to save the contents of the file to the original text file, choose **File > Save**.
- If you want to save the contents of the file to another filename or location, choose **File > Save As**.

Step 9 To exit from the Download Console, choose **File > Exit**.

Running IOS Commands from the VPN Console

You can run Cisco IOS commands on a router's command line by using VPN Solutions Center's Exec Command feature. This feature makes it easy to run commands on multiple routers at once. The Exec Command Console puts you in Enable mode, thus you can run any IOS commands that are executable in Enable mode.

(For sending configuration mode commands, use the Download Console as described in the "Using the Download Console" section on page 4-76.)

Executing commands in this way does not change the router's configuration file. VPN Solutions Center simply runs the commands you enter and returns the command's response, just as it does when communicating with a router through a console.

To execute an IOS command on a router:

Step 1 From the VPN Console menu bar, choose **Tools > Exec Command**.

The Cisco VPN Solutions Center browser appears. If the browser is not already running, you must log in.

Step 2 In the Netscape Password dialog box, enter the VPN Solutions Center administrative user name and password, then click **OK**.

The VPN Solutions Center Exec Command Console page appears (see Figure 4-107).

Figure 4-107 The Exec Command Console

You can run commands in the Exec Command Console in either of two ways:

- Specifying a command input file that contains a set of valid Cisco IOS commands.

  The command input file must be a text file. There is no practical limit to the number of commands that be included in a command input file.

- Entering commands manually in the Commands pane.

Step 3 From the Network drop-down menu, choose the name of the network that the target router resides in.

The routers in the selected network are displayed in the window below the *Network* field.

Step 4 From the list of routers, select one or more routers on which you want to run the command.

Step 5 To run IOS commands from a command input file:

a. Enter the path and name of the command input file in the *Input file* field.

   You can also specify the name and path for the command input file by clicking **Browse** and selecting the file in its directory.

b. Click **Load File**.

c. Click **Send**.

Step 6 To enter commands manually, in the Commands pane, enter the commands you want to run, then click **Send**.

If you need to erase the contents of the Commands pane, click **Clear**. Then reenter the commands as needed.

The lower pane displays the output from the command you entered for each device you selected.

CHAPTER 5

Creating MPLS VPNs and Administering Service Request Profiles

This chapter describes two Network Administrator tasks:

1. How to define MPLS VPNs using the VPN Solutions Center software.
2. How to use the MPLS Service Request Editor to create and organize Service Request Profiles.

The main **VPN** topics presented in this chapter are as follows:

The main **Service Request Profile** topics presented in this chapter are as follows:

Defining a New VPN in the VPNSC Software

You have defined the network elements, defined the Provider Administrative Domains, and imported the device configuration files into VPN Solutions Center. created the VPN customer definition. The final stage of setting up is to define the VPNs in the service provider network.

Note This procedure does not implement the VPN in the network; it only defines the existing VPN within the VPN Solutions Center software.

To define the VPN, follow these steps:

Step 1 From the VPN Console menu, choose **Setup > New VPN Definition**.

Figure 5-1 Selecting the PAD for a New VPN

Step 2 From the drop-down list in the Select Provider Administrative Domain dialog box (as shown in Figure 5-1), select the Provider Administrative Domain for the VPN, then click **OK**.

The New VPN Definition dialog box appears (see Figure 5-2).

Figure 5-2 Defining a New VPN

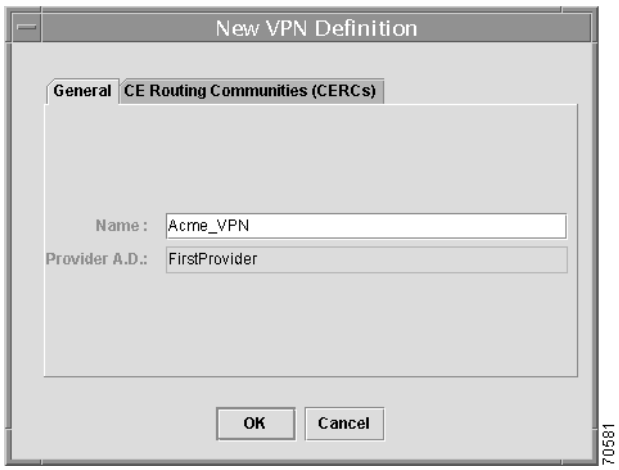

Step 3 Enter the name of the new VPN and click **OK**.

You return to the VPN Console window, which now displays the new VPN name under the VPNs folder. This is all that is required to complete the VPN definition. However, you may want to define one or more CE Routing Communities for this VPN. If so, proceed to the next section.

Defining CE Routing Communities

Whenever you create a VPN, the VPN Solutions Center software creates one default CE routing community (CERC) for you. This means that until you need advanced customer layout methods, you will not need to define new CERCs. Up to that point, consider a CERC as standing for the VPN itself—they are identical.

Tip CERCs should be defined only with 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, the provisioning software does the rest, assigning route target values and VRF tables to arrange exactly the connectivity the customer requires.

You can use the Topology tool to double-check the CERC memberships and resultant VPN connection status.

For more information on CERCs, see the "CE Routing Communities" section on page 1-18.

To define a new CE Routing Community (CERC) for a VPN, follow these steps:

Step 1 From the VPN Console hierarchy pane, expand the VPNs folder so that you can see the list of VPNs defined in VPNSC.

Step 2 Select the name of the VPN, then **right-click** (see Figure 5-3).

Figure 5-3 Opening a VPN

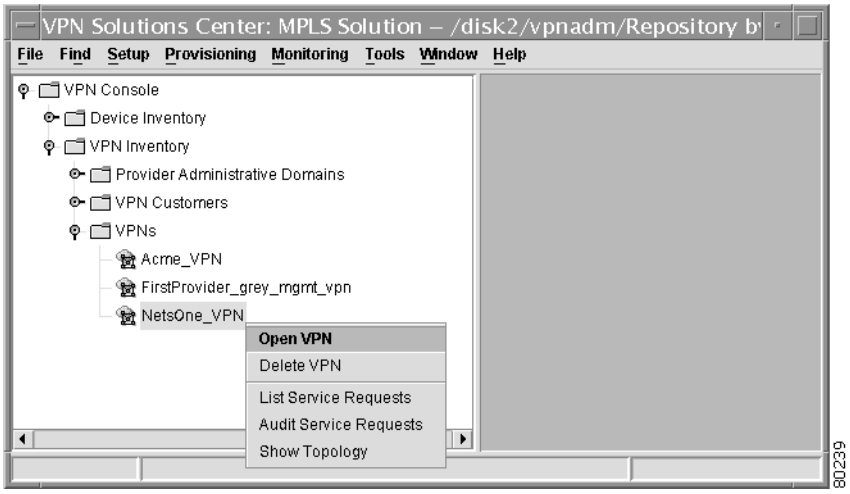

The Edit VPN Definition: General tab dialog box appears. The General tab shows the name of the selected VPN and its associated Provider Administrative Domain.

Step 3 Choose the **CE Routing Communities (CERCs)** tab (see Figure 5-4).

Figure 5-4 Editing a CE Routing Community

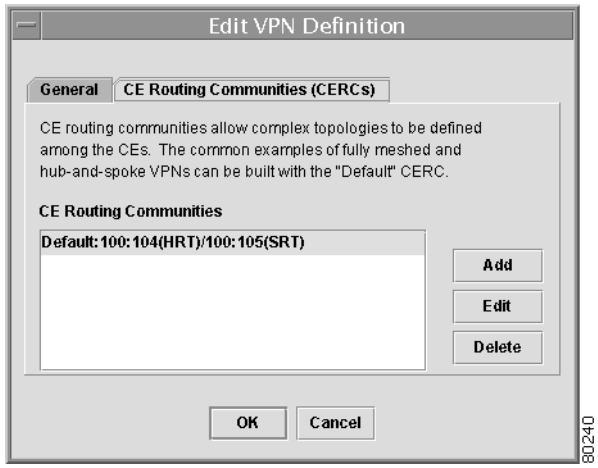

As you can see in Figure 5-4, VPNSC displays the default CERC. It shows both the *hub route target* (HRT) and the *spoke route target* (SRT).

Step 4 From the CE Routing Communities (CERCs) tab, click **Add**.

The Add CE Routing Community dialog box appears.

Figure 5-5 Add CE Routing Community Dialog Box

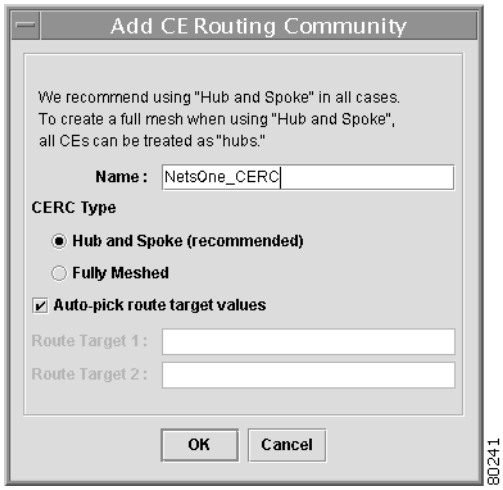

Step 5 Complete the fields as required for the VPN:

a. *Name*: Enter the name of the CERC.

b. *CERC Type*: Specify the CERC type: *Hub and Spoke* or *Fully Meshed.*

c. *Auto-pick route target values*: Choose to either let VPN Solutions Center 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 once they have been defined in the VPN Solutions Center software.

Step 6 When you have finished entering the information in the Add CE Routing Community dialog box, click **OK**.

The new CERC is added to the VPN definition.

Deleting a CE Routing Community Definition

You cannot delete a CERC from the VPNSC software if there are active service requests using the CERC.

To delete a CERC definition, follow these steps:

Step 1 From the VPN Console hierarchy view, expand the **VPNs** folder.

The list of VPNs is displayed.

Step 2 Select the VPN that the CERC is defined in, then **right-click**.

The VPN menu appears (see Figure 5-6).

Figure 5-6 The VPN Menu

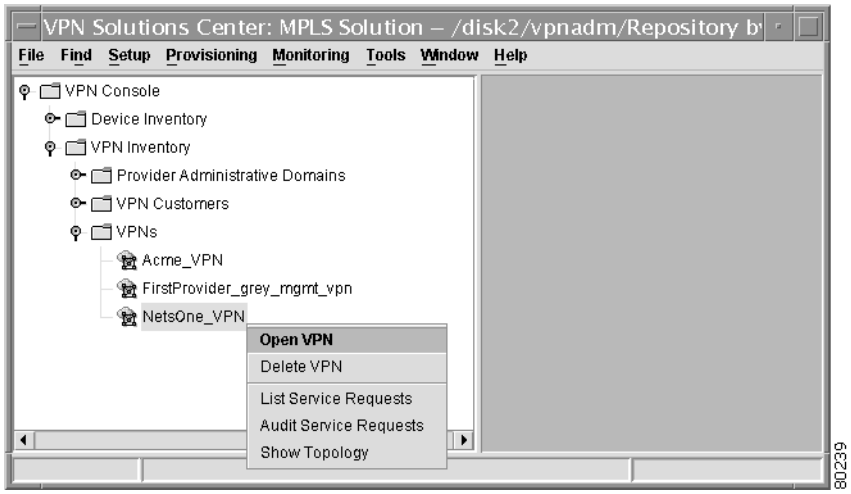

Step 3 Choose **Open VPN**.

The Edit VPN Definition: General tab dialog box appears. The General tab shows the name of the selected VPN and its associated Provider Administrative Domain.

Step 4 Choose the **CE Routing Communities (CERCs)** tab

Step 5 The Edit VPN: CE Routing Communities dialog box appears (see Figure 5-7).

Figure 5-7 List of Available CERCs Displayed

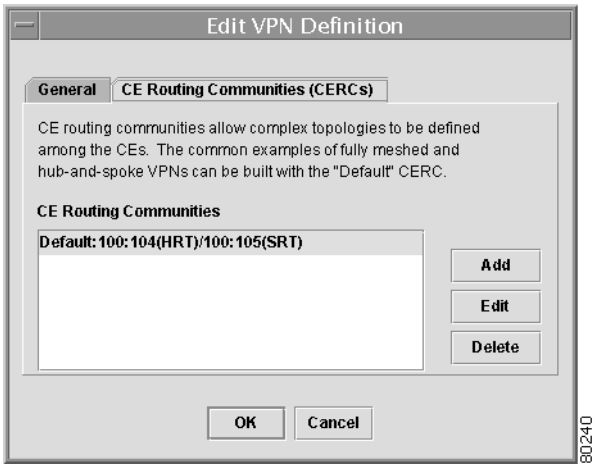

Step 6 To delete a CERC from the currently selected VPN, click **Delete**.

If the CERC you wish to delete has active service requests, you receive a warning that the CERC is not deletable.

Step 7 Click **OK**.

About Service Request Profiles

A service request profile serves as a type of template for MPLS service request attributes that are common across multiple service requests. By using service request profiles, you need only specify once the various parameters that are the same for similar service requests. Also, service request profiles make it simple to make modifications to existing service request profiles so that you can use existing profiles to quickly create variations of existing service request profiles.

In the VPN Solutions Center 2.2 release, the service profile attributes are stored as name/value string pairs in a plain text file that is placed in the Profiles subdirectory. The Profiles subdirectory is in the Repository directory.

Network administrators can organize service request profiles into whatever hierarchical categories are most useful to the service provider. For example, you could organize service request profiles by Customer and by the services provided for each Customer:

- **Customer A Service Request Profiles**
  - ATM Service Request Profile
  - Cable Service Request Profile
- **Customer B Service Request Profiles**
  - BGP Service Request Profile
  - Ethernet Service Request Profile

A service request profile must have a filename extension of **.profile**. There are no other restrictions regarding service request profile filenames. The name of a service request profile is the pathname of the profile relative to the Profiles subdirectory.

When you create a service request using a service request profile, VPNSC loads the profile into memory; the service request profile attributes are expanded and stored with the service request object itself. Thus, any changes made to the service request profile attributes for a given service request does not effect the profile itself, unless you choose to store the changes to the service request profile.

- You create and modify service request profiles with the **Profile Editor**, which is described in the next section, "Creating Service Request Profiles."
- You administer and organize service request profiles with the **Profile Manager** (see the "Administering Service Request Profiles" section on page 5-14).

MPLS Attributes and Their Corresponding Documentation

Table 5-1 lists all the MPLS attributes that are available from the MPLS Profile Editor and the section and page where you can find documentation on each attribute.

Table 5-1 MPLS Attributes in Service Request Profiles

MPLS Attributes	For Details, Refer To...
Profile	
Description	Profile Description, page 5-21
Interfaces	
PE Interface	PE Interface Name, page 5-22
PE Interface Name	PE Interface Name, page 5-22
Interface Description	PE Interface Description, page 5-22
Shutdown PE Interface	Shutdown PE Interface, page 5-22
CE Interface	
CE Interface Name	CE Interface Name, page 5-23
Interface Description	CE Interface Description, page 5-23
Encapsulations	
PE Encapsulation	PE Encapsulation, page 5-24
CE Encapsulation	CE Encapsulation, page 5-24
DLCI	
PE DLCI	PE DLCI, page 5-25
CE DLCI	CE DLCI, page 5-25
VLAN ID	
PE VLAN ID	PE VLAN ID, page 5-25
CE VLAN ID	CE VLAN ID, page 5-25
ATM Circuit Identifiers (vdc:vpi:vci)	
PE ATM Circuit Identifiers	PE ATM Circuit Identifiers, page 5-25
CE ATM Circuit Identifiers	CE ATM Circuit Identifiers, page 5-26
Tunnel Address	
Tunnel Source Address	Tunnel Source Address (PE), page 5-27

Table 5-1 MPLS Attributes in Service Request Profiles (continued)

MPLS Attributes	For Details, Refer To...
Tunnel Destination Address	Tunnel Destination Address (CE), page 5-28
Cable Helper Addresses	
Cable Maintenance Interface	Cable Helper Addresses, page 5-29
Cable Helper Addresses	
Secondary Addresses	Secondary Addresses, page 5-31
Routing Information	
Routing Protocol	Routing Information, page 5-31
Give Only Default Routes to CE	Give Only Default Routes to CE, page 5-33
Redistribute Static	Redistribute Static (BGP and RIP), page 5-35
Redistribute Connected	Redistribute Connected (BGP Only), page 5-35
Static	Static Protocol Chosen, page 5-33
Advertised Routes to CE	Advertised Routes to CE, page 5-33
Routes to Reach Other Sites	Routes to Reach Other Sites, page 5-34
RIP	RIP Protocol Chosen, page 5-35
Redistribute OSPF	Redistribute OSPF, page 5-35
OSPF Process ID for RIP	OSPF Process ID for RIP, page 5-35
Redistributed Protocols	Redistributed Protocols, page 5-36
BGP	BGP Protocol Chosen, page 5-37
BGP AS ID	BGP AS ID, page 5-37
Neighbor Allow-AS In	Neighbor Allow-AS In, page 5-37
Neighbor AS Override	Neighbor AS Override, page 5-37
Redistributed Protocols	Redistributed Protocols, page 5-38
OSPF	OSPF Protocol Chosen, page 5-39
Redistribute RIP	Redistribute RIP, page 5-39
OSPF Process ID on PE	OSPF Process ID on PE, page 5-39
OSPF Process ID on CE	OSPF Process ID on CE, page 5-39
OSPF Area Number on PE	OSPF Area Number on PE, page 5-39
OSPF Area Number on CE	OSPF Area Number on CE, page 5-40
Redistributed Protocols	Redistributed Protocols, page 5-40
Interface Addresses	Interface Addresses, page 5-41
IP Numbering Scheme	IP Numbering Scheme, page 5-41
Extra CE Loopback Required	Extra CE Loopback Required, page 5-42
Automatically Assign IP Address	Automatically Assign IP Address, page 5-43
PE Interface Address	PE Interface Address, page 5-43
CE Interface Address	CE Interface Address, page 5-43
Extra CE Loopback Address	Extra CE Loopback Address, page 5-44

Table 5-1 MPLS Attributes in Service Request Profiles (continued)

MPLS Attributes	For Details, Refer To...
VRF Maps	VRF Maps, page 5-44
Export Map	Export Map, page 5-45
Import Map	Import Map, page 5-45
Maximum Routes	Maximum Routes, page 5-45
NetFlow	NetFlow, page 5-46
Turn on NetFlow Accounting	Turn on NetFlow Accounting, page 5-46
Templates	Templates, page 5-46
PE Template	PE Template, page 5-46
CE Template	CE Template, page 5-49

Creating Service Request Profiles

You can create a new service request profile or edit an existing profile with the **Profile Editor**. The Profile Editor presents service request attributes in logical categories.

To create a new service request profile or edit an existing profile, follow these steps:

Step 1 From the VPN Console, choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor is displayed (see Figure 5-8).

Figure 5-8 The MPLS Service Request Editor

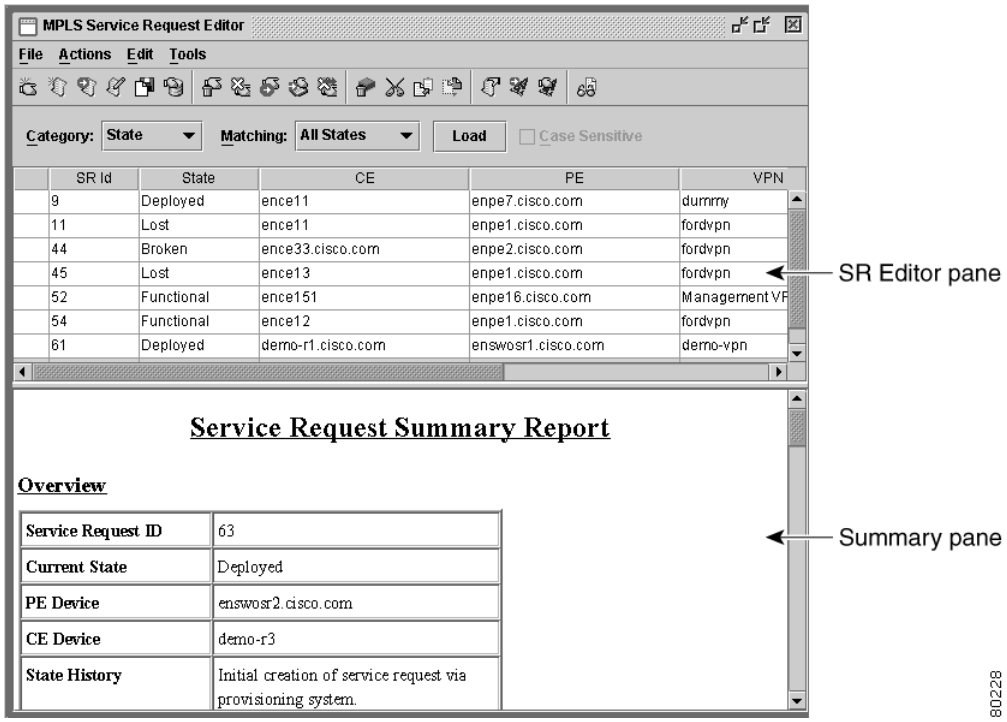

For a detailed description of the Service Request Editor, see the "Adding a Service for a PE-CE Link" section on page 6-6.

Step 2 From the Service Request Editor menu bar, choose **File > New Profile**.

The Profile Editor appears (see Figure 5-9).

Figure 5-9 Profile Editor

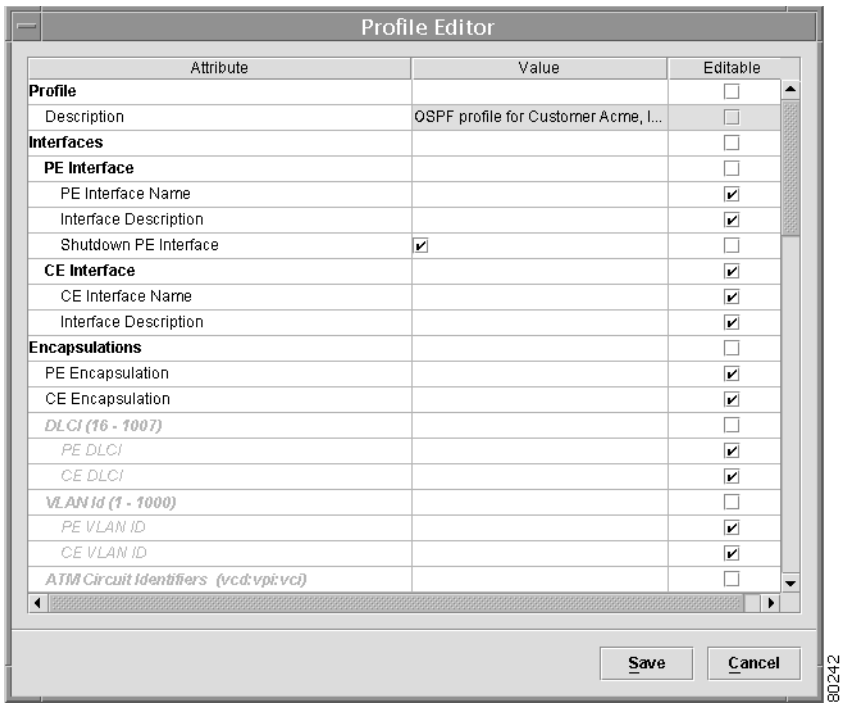

The Profile Editor presents three columns: **Attribute**, **Value**, and **Editable**:

- **Attribute**

  The Attribute column displays the name of the attribute categories (for example, "Interfaces") and the attributes for each category. For a list of all the attributes available from the MPLS Profile Editor, see Table 5-1 on page 5-7.

- **Value**

  For the details on specifying values for any MPLS attribute, see the "Specifying the MPLS Attributes for a Service Request Profile" section on page 5-20.

  You can enter values for each value listed in the Attribute column by **double-clicking** the pertinent cell for the attribute you want to edit. When an attribute is available for editing, the attribute option is displayed. When an attribute is not available for editing, it is displayed in a gray italic font. (The attribute categories, such as "Interfaces" and "Encapsulations" do not have associated values.)

  The type of editor that is invoked when you double-click an attribute cell 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 is displayed. 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.

Note that 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 Profile 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 Administrator to indicate the attributes that are likely to change across multiple service requests. When certain attributes are checked as editable, only those attributes will be made available to the Network Operator when creating or modifying service requests with that service request profile.

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

Step 3 Enter the values for each attribute that are common across multiple service requests.

Step 4 If desired, select the attributes that you want to be editable by the Network Operator when creating service requests.

Step 5 When satisfied with the results, click **Save**.

The Save Profile As dialog box appears (see Figure 5-10).

Figure 5-10 Saving a Service Request Profile

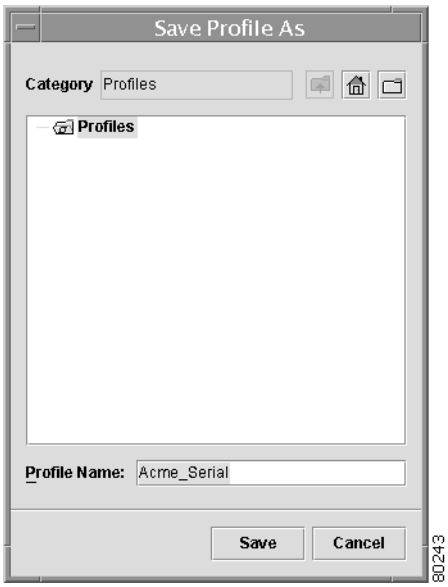

Step 6 Enter the name of the service request profile, then click **Save**.

The new profile is added to the set of service request profiles. You return to the MPLS Service Request Editor.

Opening or Editing an Existing Service Profile

To open or edit an existing service request profile, follow these steps:

Step 1 If the MPLS Service Request Editor is not up, choose **Provisioning > Add VPN Service to CE** from the VPN Console.

The MPLS Service Request Editor is displayed.

Step 2 From the MPLS Service Request Editor, choose **File > Edit Profile**.

The Open Profile dialog box appears (see Figure 5-11).

Figure 5-11 Opening a Service Request Profile

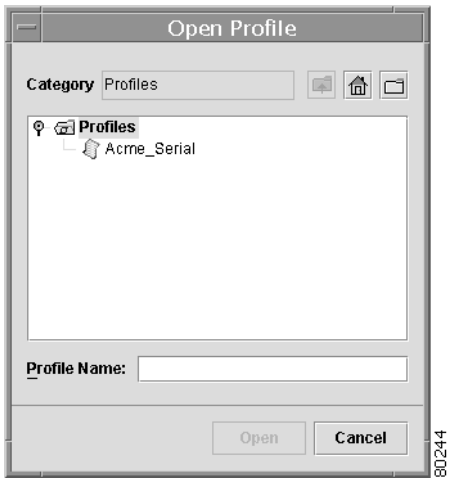

Step 3 From the Open Profile dialog box, select the profile you want to open, then click **Open**.

The Profile Editor displays the selected service request profile (see Figure 5-12).

Figure 5-12 Editing a Selected Service Request Profile

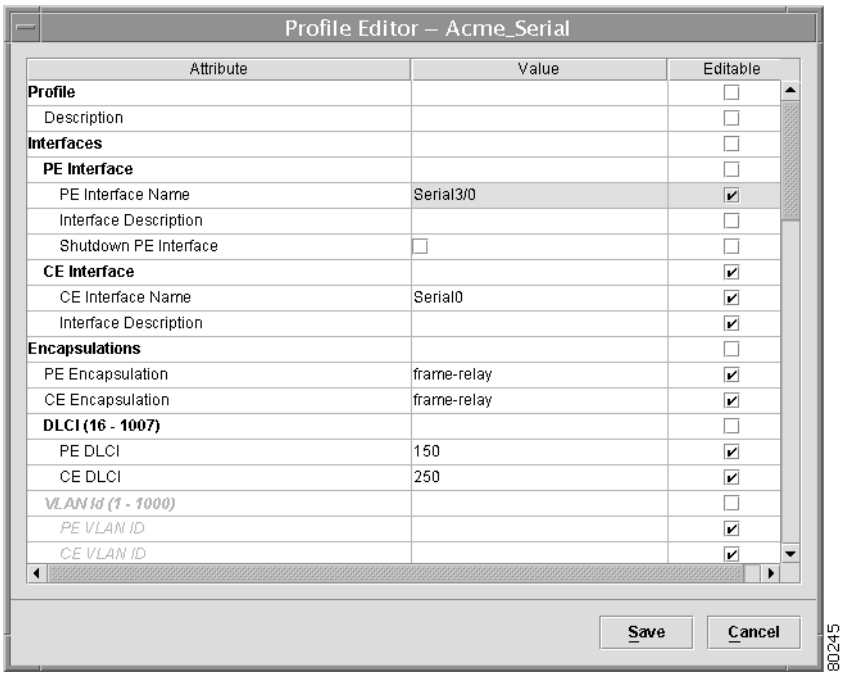

Step 4 Make any necessary changes to the selected service request profile, then click **Save**.

For details on specifying any of the MPLS attributes, see the "Specifying the MPLS Attributes for a Service Request Profile" section on page 5-20.

Administering Service Request Profiles

The Profile Manager provides a mechanism for Network Administrators to manage all the service request profiles for the VPN Solutions Center installation. The Profile Manager allows you to create, edit, delete, rename, and move service request profiles, as well as create profile categories.

To administer service request profiles, follow these steps:

Step 1 If the MPLS Service Request Editor is not up, choose **Provisioning > Add VPN Service to CE** from the VPN Console.

The MPLS Service Request Editor is displayed (see Figure 5-13).

Figure 5-13 The MPLS Service Request Editor

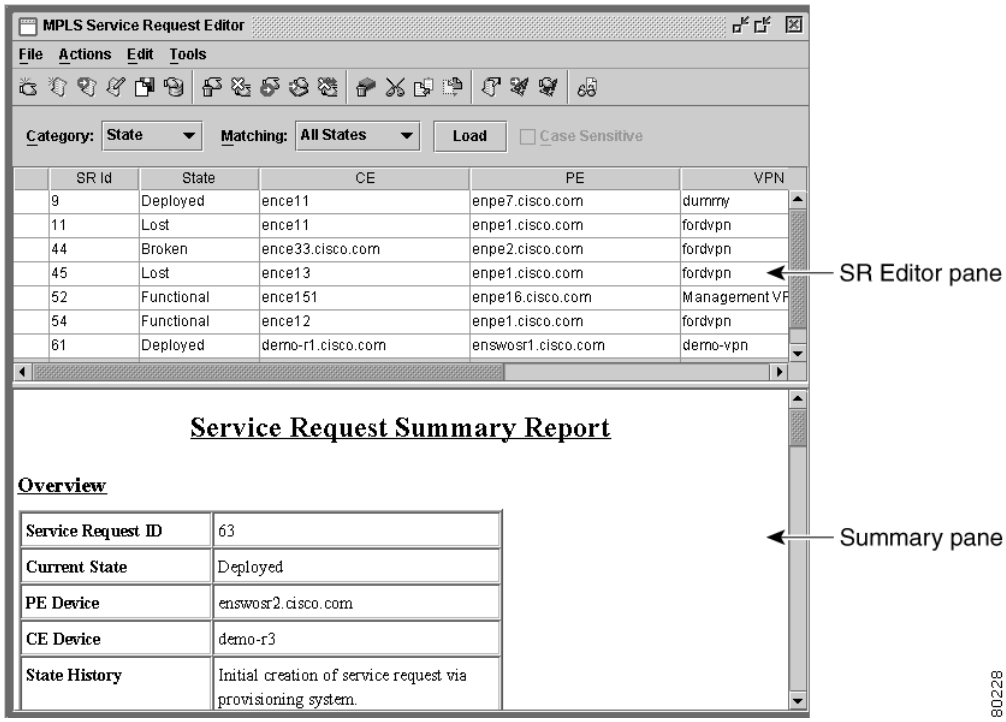

Step 2 From the MPLS Service Request Editor, choose **Tools > Profile Manager**.

The Profile Manager dialog box appears (see Figure 5-14).

Figure 5-14 The Profile Manager

From the Profile Manager, you can perform the following tasks:

- Creating and Renaming a New Service Request Profile Category, page 5-15
- Creating a New Service Request Profile, page 5-17
- Editing a Service Request Profile, page 5-18
- Moving a Service Request Profile, page 5-19
- Deleting a Service Request Profile, page 5-19

Creating and Renaming a New Service Request Profile Category

You can use the Profile Manager to organize your service request profiles into categories that are logical and useful for the service provider. For example, profiles could be organized at the top level by Customer, with each Customer category containing service request profiles for the types of VPN services provided for that Customer:

- **Customer A Service Request Profiles**
  - ATM Service Request Profile
  - Cable Service Request Profile
- **Customer B Service Request Profiles**
  - BGP Service Request Profile
  - Ethernet Service Request Profile

When you click the Create New Category icon (see Figure 5-15 on page 5-16), the Profile Manager creates the new profile category and places it *under the category that is currently selected.*

The New Category is automatically named “NewCategory.” You must then rename the new category to the desired name.

To create a new service profile category:

Step 1 Invoke the Profile Manager by choosing **Tools > Profile Manager** from the MPLS Service Request Editor.

The Profile Manager dialog box appears.

Step 2 Select the category that you want to place the new category under.

Step 3 Click the **Create New Category** icon.

A new category named "NewCategory" appears in the Profile Manager tree (see Figure 5-15).

Figure 5-15 A New Profile Category Created

Step 4 With "NewCategory" still selected, click **Rename**.

The Rename Profile dialog box appears (see Figure 5-16).

Figure 5-16 Renaming a New Category

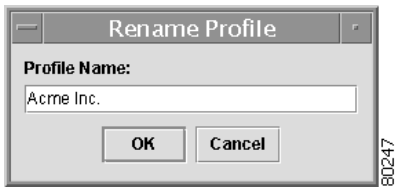

Step 5 Enter the name of the new category, then click **OK**.

The only valid characters are alphanumeric characters, the period, underscore, and hyphen.

The new service profile category is displayed in the Profile Manager tree (see Figure 5-17).

Figure 5-17 A New Profile Category Renamed

Step 6 If finished with the current Profile Manager session, click **Close**.

You return to the MPLS Service Request Editor.

Creating a New Service Request Profile

You can create a service request profile from two locations in the VPN Solutions Center user interface: the VPN Console Setup menu (see "Creating Service Request Profiles" section on page 5-9) or the Profile Manager.

To create a new service request profile from the Profile Manager:

Step 1 Invoke the Profile Manager by choosing **Tools > Profile Manager** from the MPLS Service Request Editor.

The Profile Manager dialog box appears.

Step 2 Select the service request category in which you want to place the profile, then click **Create**.

The Profile Editor dialog box appears (see Figure 5-18).

Figure 5-18 Profile Editor

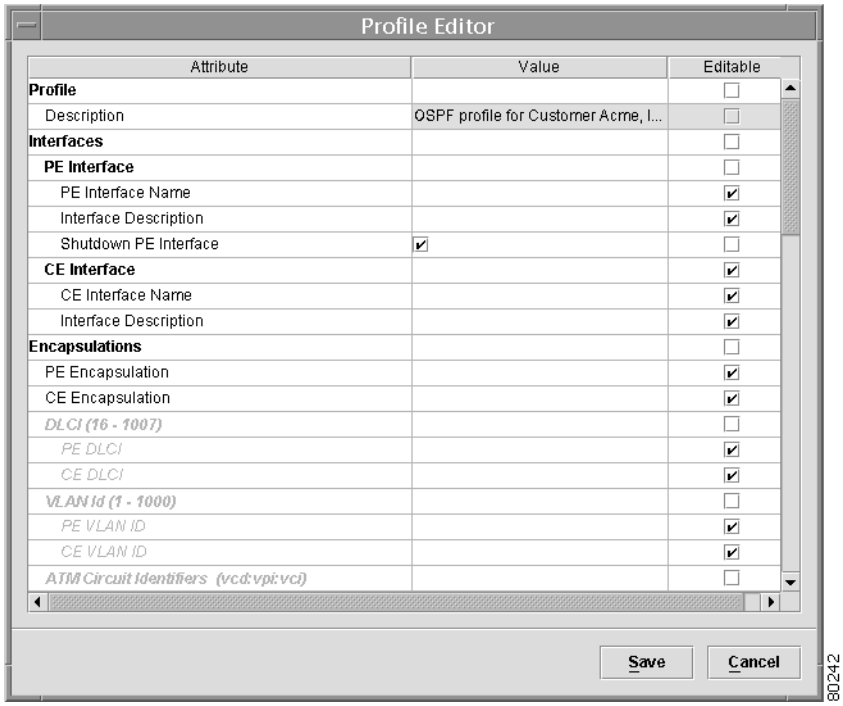

Step 3 Enter the values for each attribute that are common across multiple service requests.

Step 4 If desired, select the attributes that you want to be editable by the Network Operator when creating service requests.

Step 5 When satisfied with the results, click **Save**.

The Save Profile As dialog box appears.

Step 6 Enter the name of the service request profile, then click **Save**.

The new profile is added to the set of service request profiles. You return to the MPLS Service Request Editor.

Editing a Service Request Profile

To edit a service request profile, follow these steps:

Step 1 Invoke the Profile Manager by choosing **Tools > Profile Manager** from the MPLS Service Request Editor.

The Profile Manager dialog box appears.

Step 2 Select the service request profile you want to modify, then click **Edit**.

The Profile Editor displays the selected service request profile.

Step 3 Make any necessary changes in the profile, then click **Save**.

You return to the Profile Manager.

Step 4 Continue with additional service profile tasks or click **Close** to close the Profile Manager.

Moving a Service Request Profile

You may need to move a service request profile to another location in the Profile Manager tree. To move a profile, follow these steps:

Step 1 Invoke the Profile Manager by choosing **Tools > Profile Manager** from the MPLS Service Request Editor.

The Profile Manager dialog box appears.

Step 2 Select the service request profile you want to move, then click **Move**.

The Select Destination Profile Category dialog box appears (see Figure 5-19).

Figure 5-19 Selecting a Destination Profile Category

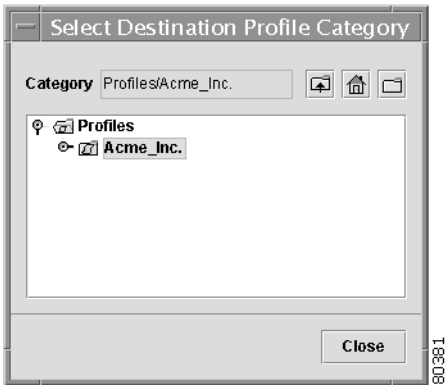

Step 3 Select the profile category that you want to move the selected service request profile to, then click **Close**.

The selected service request profile is moved to its new location in the Profile Manager tree.

Deleting a Service Request Profile

To delete a service request profile, follow these steps:

Step 1 Invoke the Profile Manager by choosing **Tools > Profile Manager** from the MPLS Service Request Editor.

The Profile Manager dialog box appears.

Step 2 Select the service request profile you want to delete, then click **Delete**.

The selected service request profile is deleted from the Profile Manager tree.

Specifying the MPLS Attributes for a Service Request Profile

This section describes the details you need to know to specify the MPLS attributes in the MPLS Service Request Profile Editor. For a complete list of MPLS attributes that are available in the Profile Editor, see Table 5-1 on page 5-7.

To open or edit an existing service request profile, follow these steps:

Step 1 If the MPLS Service Request Editor is not up, choose **Provisioning > Add VPN Service to CE** from the VPN Console.

The MPLS Service Request Editor is displayed.

Step 2 From the MPLS Service Request Editor, choose **File > Edit Profile**.

The Open Profile dialog box appears (see Figure 5-20).

Figure 5-20 Opening a Service Request Profile

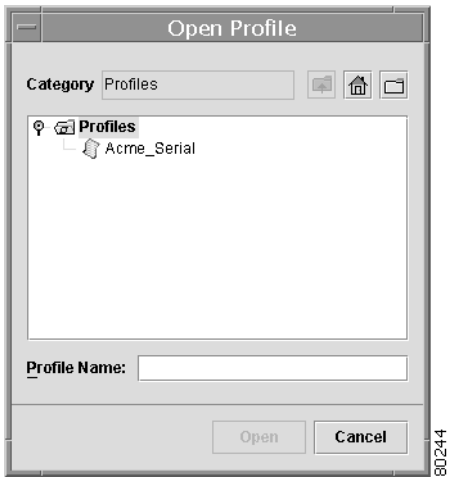

Step 3 From the Open Profile dialog box, select the profile you want to open, then click **Open**.

The Profile Editor displays the selected service request profile.

Profile Description

To enter the description of the service request profile:

Step 1 In the Profile Editor's Value column, **double-click** the *Profile Description* field.

The Profile Description dialog box appears (see Figure 5-21).

Figure 5-21 Entering the Profile Description

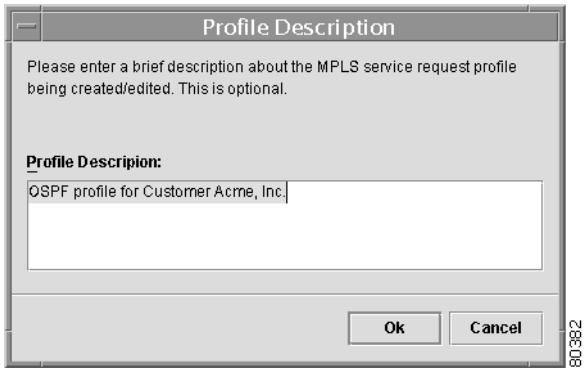

Step 2 Enter a description of the service request profile, then click **OK**.

Interfaces

VPN Solutions Center supports the following interface types (for both PEs and CEs):

- ATM (Asynchronous Transfer Mode)
- Cable
- Ethernet
- Fast Ethernet
- FDDI (Fiber Distributed Data Interface)
- Gigabit Ethernet
- Gigabit Ethernet WAN

  To enter this interface type in the Profile Editor, type `ge-wan`
- HSSI (High Speed Serial Interface)
- Loopback
- PoS (Packet over Sonet)
- Serial
- Tunnel

For information regarding the protocol encapsulations for each of the interface types, see Table 5-2 on page 5-23.

PE Interface

PE Interface Name

Click into the *PE Interface Name* cell, then enter the PE interface name.

The Profile Editor checks to see if the interface name is valid. If it is not, you receive an error message:

```
Profile Error: Unknown or unsupported interface name entered for PE Interface Name,
"name_entered." Please enter a valid interface name.
```

Click **OK** to exit the message and return to the Profile Editor. Then reenter a valid interface name.

PE Interface Description

To enter the description of the PE interface:

Step 1 In the Profile Editor's Value column, **double-click** the PE Interface *Interface Description* field.

The PE Interface Description dialog box appears (see Figure 5-22).

Figure 5-22 Entering the PE Interface Description

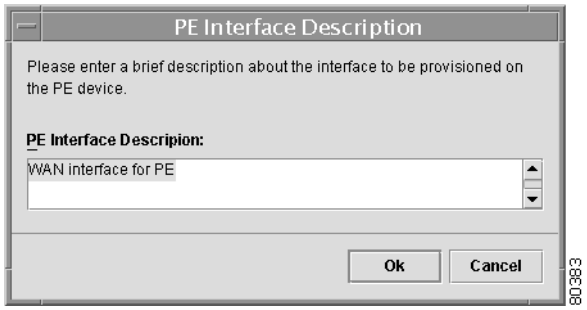

Step 2 Enter the description for the PE interface, then click **OK**.

Shutdown PE Interface

When you enable the **Shutdown PE Interface** checkbox, the specified PE interface is configured in a shut down state.

CE Interface

CE Interface Name

Click into the *CE Interface Name* cell, then enter the CE interface name.

The Profile Editor checks to see if the interface name is valid. If it is not, you receive an error message:

```
Profile Error: Unknown or unsupported interface name entered for CE Interface Name,
"name_entered." Please enter a valid interface name.
```

Click **OK** to exit the message and return to the Profile Editor.

CE Interface Description

To enter the description of the CE interface:

Step 1 In the Profile Editor's Value column, **double-click** the CE Interface *Interface Description* field.

The CE Interface Description dialog box appears (see Figure 5-23).

Figure 5-23 Entering the CE Interface Description

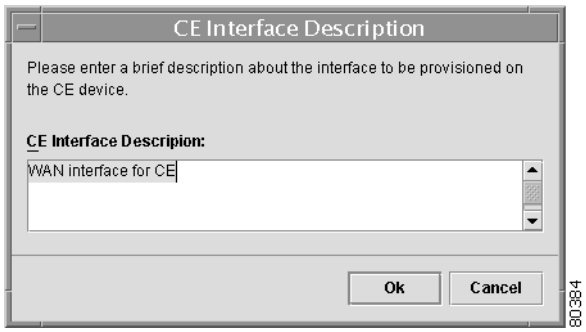

Step 2 Enter the description for the CE interface, then click **OK**.

Encapsulations

Encapsulation occurs when data is wrapped in a particular protocol header. For example, Ethernet data is wrapped (encapsulated) in a specific Ethernet header before transit over a network.

Table 5-2 shows the protocol encapsulations available for each of the supported interface types.

Table 5-2 Interface Types and Their Protocol Encapsulations

Interface Type	Encapsulations
ATM	ATM
Cable	Default frame
Ethernet	Default frame
Fast Ethernet	Default frame, ISL (Inter-Switch Link), 802.1q

Table 5-2 Interface Types and Their Protocol Encapsulations

Interface Type	Encapsulations
FDDI (Fiber Distributed Data Interface)	Default frame
Gibabit Ethernet	Default frame, ISL (Inter-Switch Link), 802.1q
Gigabit Ethernet WAN	Default frame, ISL (Inter-Switch Link), 802.1q
HSSI (High Speed Serial Interface)	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.
Loopback	Default frame
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)
Tunnel	GRE (Generic Routing Encapsulation)

PE Encapsulation

In the *PE Encapsulation* cell, specify the encapsulation method for the link.

The protocol encapsulations that are available depend on which interface type is specified in the *PE Interface Name* cell.

- The *ISL (Inter-Switch Link) Protocol* is used to connect two VLAN-capable Ethernet, Fast Ethernet, or Gigabit Ethernet devices. The packets on the ISL link contain a standard Ethernet, FDDI, or Token-Ring frame, and the VLAN information associated with that frame.
- The *802.1Q protocol* establishes a standard method for tagging Ethernet frames with VLAN membership information. The 802.1Q protocol was developed 1) to allow for the segmentation of large switched networks into smaller segments so that broadcast and multicast traffic does not use excessive bandwidth, and 2) to provide a higher level of security between segments of internal networks.
- When you enter a tunnel interface name, the encapsulation is automatically set to **GRE** (Generic Routing Encapsulation).

For a comprehensive list of the protocol encapsulations for each interface type, see Table 5-2 on page 5-23.

CE Encapsulation

In the *CE Encapsulation* cell, specify the encapsulation method for the link.

The protocol encapsulations that are available depend on which interface type is specified in the *CE Interface Name* cell.

For a comprehensive list of the protocol encapsulations for each interface type, see Table 5-2 on page 5-23.

DLCI

A DLCI (Data-Link Connection Identifier) is a value that specifies a PVC (permanent virtual circuit) or a SVC (switched virtual circuit) in a Frame Relay network. In the basic Frame Relay specification, DLCIs are locally significant—that is, connected devices might use different values to specify the same connection. In the LMI (Local Management Interface) extended specification, DLCIs are globally significant—they specify individual devices).

When an interface encapsulation type is set to Frame Relay or Frame Relay-IETF, the Profile Editor's DLCI cells are enabled.

PE DLCI

Click into the *PE DLCI* cell, then enter the DLCI number for the PE.

CE DLCI

Click into the *CE DLCI* cell, then enter the DLCI number for the CE.

VLAN ID

When you specify the protocol encapsulation for the PE or CE as either **802.1q** or **ISL**, the VLAN (Virtual LAN) ID cells in the Profile Editor are enabled. The valid values are any integer from 1 to 1005.

PE VLAN ID

Click into the *PE VLAN ID* cell, then enter the VLAN ID for the PE.

CE VLAN ID

Click into the *CE VLAN ID* cell, then enter the VLAN ID for the CE.

ATM Circuit Identifiers

ATM standards define two types of ATM connections: *virtual path connections* (VPCs), which contain *virtual circuit connections* (VCCs). A virtual circuit, which is the basic unit, carries a single stream of cells, in order, from user to user. A collection of virtual circuits can be bundled together into a virtual path connection. An ATM network also uses virtual paths internally to bundle virtual circuits together between switches. Two ATM switches can have many different virtual circuit connections between them, each belonging to a different user.

PE ATM Circuit Identifiers

If you chose **ATM** as the encapsulation method, the ATM circuit ID information for the PE and CE is enabled in the Profile Editor.

The PE and CE are connected by a virtual path. Each virtual path can carry up to 65,536 virtual circuits (simultaneous connections). Each virtual circuit is identified by the 16-bit virtual circuit ID, which uniquely identifies the ATM connection.

To enter the ATM circuit identifiers for the PE:

Step 1 In the Profile Editor's Value column, **double-click** the *PE ATM Circuit Identifiers* cell.

The PE ATM Circuit Identifiers dialog box appears (see Figure 5-24).

Figure 5-24 Entering the PE's ATM Circuit Identifiers

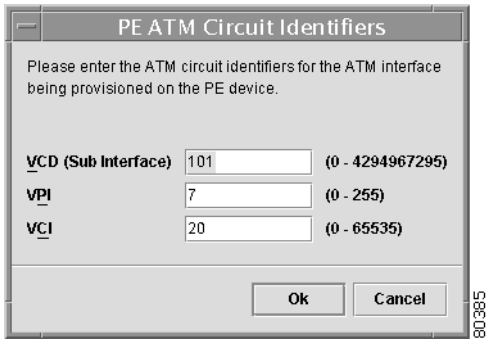

Step 2 Enter the appropriate values for the ATM circuit:

a. *VCD (PE Subinterface number)*: Enter the subinterface number for the ATM circuit on the PE.

b. *PE VPI*: Enter the Virtual Path ID (VPI) for the PE. The VPI identifies the virtual path for this connection.

c. *PE VCI*: Enter the Virtual Circuit ID (VCI) for the PE. The VCI identifies the virtual circuit within the specified virtual path connection.

Step 3 When satisfied with the settings, click **OK**.

The values you entered are displayed in the *PE ATM Circuit Identifiers* cell in the format:

VCD_number:*VPI_number*:*VCI_number*

For example: **101:7:20**

CE ATM Circuit Identifiers

To enter the ATM circuit identifiers for the CE:

Step 1 In the Profile Editor's Value column, **double-click** the *CE ATM Circuit Identifiers* cell.

The CE ATM Circuit Identifiers dialog box appears (see Figure 5-25).

Figure 5-25 Entering the CE's ATM Circuit Identifiers

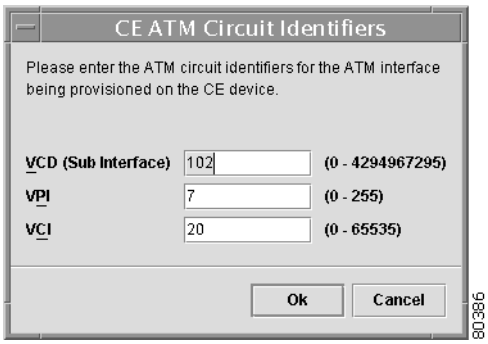

Step 2 Enter the appropriate values for the ATM circuit:

a. *VCD (CE Subinterface number)*: Enter the subinterface number for the ATM circuit on the CE.

b. *CE VPI*: Enter the Virtual Path ID (VPI) for the CE. The VPI identifies the virtual path for this connection.

c. *CE VCI*: Enter the Virtual Circuit ID (VCI) for the PE. The VCI identifies the virtual circuit within the specified virtual path connection.

Step 3 When satisfied with the settings, click **OK**.

The values you entered are displayed in the *CE ATM Circuit Identifiers* cell in the format:

VCD_number:*VPI_number*:*VCI_number*

For example: **102:7:20**

Tunnel Address

A *tunnel interface* is a logical interface used when the PE and CE are not connected directly, but are instead connected through an IPv4 network via a GRE (Generic Routing Encapsulation) tunnel.

By default, the logical interface type is *not* available in the VPNSC provisioning user interface. To include the Logical interface type, you must set the following property in the *csm.properties* file to **true**:

netsys.svrc.showLogicalInterfaces.unix=false

Tunnel Source Address (PE)

The tunnel addresses in the Profile Editor are enabled when the PE and CE interface names are set to a tunnel interface; for example, **tunnel0**.

The *tunnel source address* is the tunnel address on the PE.

The *logical tunnel interface* is used when the PE and CE are not connected directly, but are instead connected through an IPv4 network via a GRE (Generic Routing Encapsulation) tunnel.

To enter the tunnel source address on the PE:

Step 1 In the Profile Editor's Value column, **double-click** the *Tunnel Source Address* cell.

The Tunnel Source Address dialog box appears (see Figure 5-26).

Figure 5-26 Entering the Tunnel Source Address on the PE

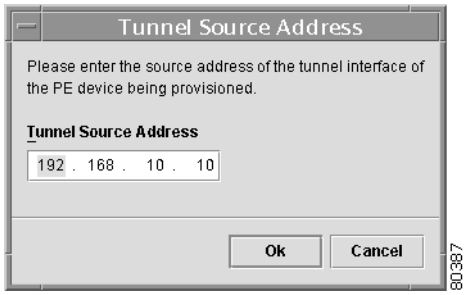

Step 2 Enter the IP address for the tunnel source address on the PE, then click **OK**.

Tunnel Destination Address (CE)

The tunnel addresses in the Profile Editor are enabled when the PE and CE interface names are set to a tunnel interface; for example, **tunnel0**.

The *tunnel destination address* is the tunnel address on the PE.

The *logical tunnel interface* is used when the PE and CE are not connected directly, but are instead connected through an IPv4 network via a GRE (Generic Routing Encapsulation) tunnel.

To enter the tunnel destination address on the CE:

Step 1 In the Profile Editor's Value column, **double-click** the *Tunnel Destination Address* cell.

The Tunnel Source Destination dialog box appears (see Figure 5-27).

Figure 5-27 Entering the Tunnel Destination Address on the CE

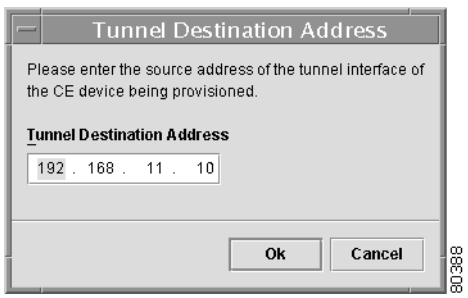

Step 2 Enter the IP address for the tunnel source address on the CE, then click **OK**.

Cable Helper Addresses

In the cable subscriber environment, several thousand subscribers share a single physical interface. Configurations with multiple logical subinterfaces are a vital part of the MPLS VPN network over cable. You can configure multiple subinterfaces and associate a specific VRF (VPN Routing/Forwarding table) with each subinterface. You can split a single physical interface (the cable plant) into multiple subinterfaces, where each subinterface is associated with a specific VRF. Each ISP requires access on a physical interface and is given its own subinterface. The MSO administrator can define subinterfaces on a cable physical interface and assign Layer 3 configurations to each subinterface.

The MPLS VPN approach of creating VPNs for individual ISPs or customers requires subinterfaces to be configured on the cable interface. One subinterface is required for each ISP. The subinterfaces are tied to the VRF tables for their respective ISPs.

For details on implementing cable services with VPN Solutions Center (MPLS), see Chapter 9, "Provisioning MPLS VPN Cable Services."

When you specify the PE interface as a *cable* interface, the Profile Editor's *Cable Helper Address* cells are enabled.

Cable Maintenance Interface

When configuring MPLS VPNs for cable services, you must configure the cable maintenance subinterface on the PE. The cable maintenance interface is the means by which the cable device retrieves its own IP address. For this reason, the maintenance subinterface must be configured before cable services provisioning can take place. See the "Provisioning the Cable Maintenance Subinterface" section on page 9-9.

When you enable the Cable Maintenance Interface checkbox, the following takes place:

- The *Secondary Addresses* cell is disabled.
- When you provision a service request using a service request profile with the Cable Maintenance Interface enabled, the service request provisions a .1 subinterface as the cable maintenance interface. This is one-time operation for one major interface.

Tip The cable maintenance interface is provisioned separately—that is, it requires its own service request.

The cable maintenance interface (nnn.nnn.nnn.nnn.1) is assigned to the cable helper address specified in the next section, "Cable Helper Addresses."

When configuring the maintenance interface, be sure to set the Helper Type to **Modem** (see Figure 5-28 below).

Cable Helper Addresses

A *cable helper address* is the IP address of the DHCP server in the Multiple Service Operator (MSO) network. You can specify three types of maintenance helper addresses:

- Modem helper address

  The IP address of the DHCP server in the MSO's network. The modem helper address assigns the IP address of the cable modem interface. A cable modem helper address specifies that only cable modem UDP broadcasts are forwarded.

- Host helper address

  The IP address of the DHCP server of the Internet Service Provider (ISP) to which the customer belongs. A cable host helper address specifies that only cable host UDP broadcasts are forwarded.

- Both

  Cable helper address that is both the host and modem address.

To specify the cable helper address(es):

Step 1 In the Profile Editor's Value column, **double-click** the *Cable Helper Addresses* cell.

The Cable Helper Addresses dialog box appears (see Figure 5-28).

Figure 5-28 Entering a Cable Helper Address

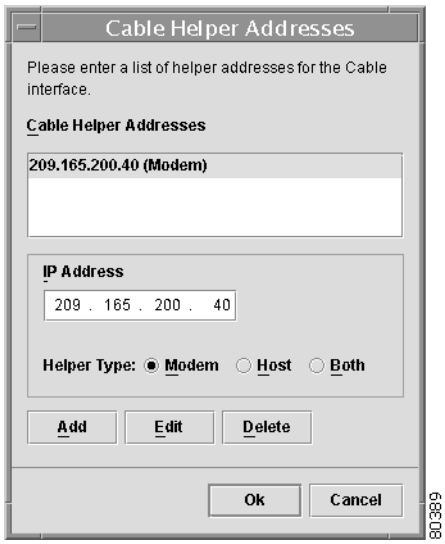

Step 2 *IP Address*: Enter the IP address for the cable helper address on the PE.

Step 3 *Helper Type*: Specify the type of helper address: **Modem**, **Host**, or **Both**.

Step 4 Click **Add**.

Step 5 To add additional helper addresses, repeat steps 2, 3, and 4.

You can edit or delete any helper addresses that have been added to the list.

Step 6 When you are finished adding helper addresses, click **OK**.

Secondary Addresses

Secondary addresses are IP addresses that are used for routing packets to the host devices connected to the cable modem. All the host devices on that cable subnet can use a single secondary address.

To specify the cable secondary address(es):

Step 1 In the Profile Editor's Value column, **double-click** the *Secondary Addresses* cell.

The Secondary Addresses dialog box appears (see Figure 5-29).

Figure 5-29 Entering a Cable Secondary Address

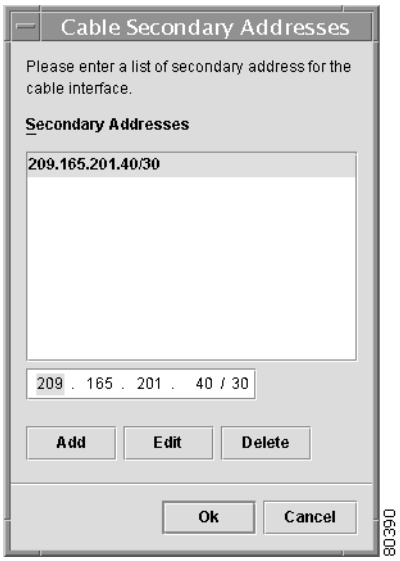

Step 2 *IP Address*: Enter the IP address and subnet mask for the secondary address on the PE.

Step 3 Click **Add**.

Step 4 To add additional secondary addresses, repeat steps 1, 2, and 3.

You can edit or delete any secondary addresses that have been added to the list.

Step 5 When you are finished adding helper addresses, click **OK**.

Routing Information

In this section of the Service Request Profile, you specify the routing protocol for the PE-CE link.

Routing Protocol

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

- **Static** (for specifying a static route)
- **RIP** (Routing Information Protocol)

- **BGP** (Border Gateway Protocol)
- **OSPF** (Open Shortest Path First)
- **No Routing** (to specify parameters for cable services).

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

Step 1 **Double-click** the *Routing Protocol* cell.

The drop-down list shown in Figure 5-30 is displayed:

Figure 5-30 Selecting the Routing Protocol

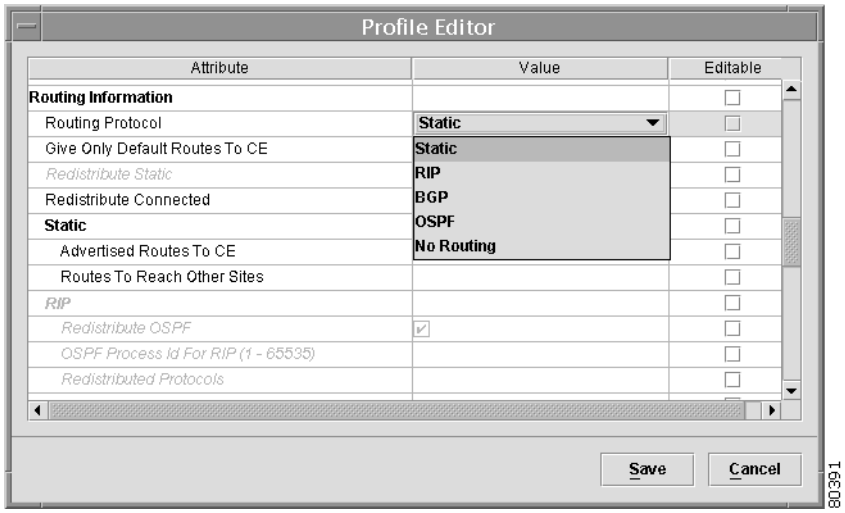

Step 2 Select the appropriate routing protocol for the link.

The related cells for the selected protocol are enabled in the Profile Editor.

Step 3 Complete the information required for the selected protocol.

Give 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.

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 1) route all packets to unknown destinations to the Internet, or 2) 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 pack

Static Protocol Chosen

When you select **Static** as the protocol, two cells are enabled: *Give Only Default Routes to CE* and *Redistribute Connected*.

Give Only Default Routes to CE

When you enable the **Give only default routes to CE** option with static route provisioning on the PE-CE link, VPN Solutions Center creates a default route on the CE that points to the PE. The VRF static route to the CE's 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.

Redistribute Connected (BGP Only)

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.

Tip When joining the management VPN and you are using IP numbered addresses on the link, you must enable the **Redistribute Connected** option.

Static Routes

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.

Advertised Routes to CE

When you specify the advertised routes, you create a list of static routes to put on the PE that describe all of the address space in the CE's site.

To specify the local routes that should be advertised from the selected CE:

Step 1 **Double-click** the *Advertised Routes to CE* cell.

The Advertised Routes dialog box is displayed (see Figure 5-31).

Figure 5-31 Specifying Local Routes to be Advertised

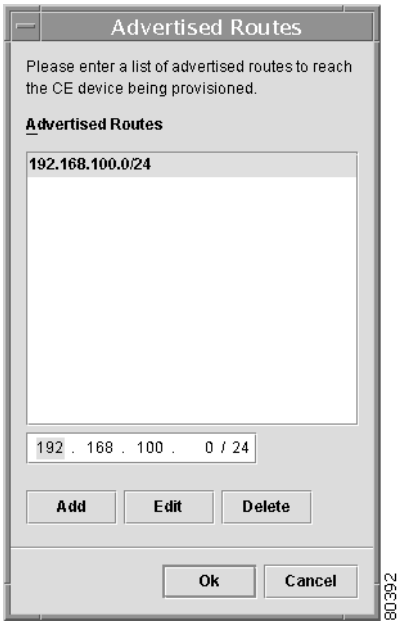

Step 2 Enter the IP address and subnet mask of the first local route to be advertised.

Step 3 Click **Add**.

The address is added to the list of advertised routes.

Step 4 Click **OK**.

You return to the Profile Editor, where the specified route is displayed in the *Advertised Routes to CE* cell.

Step 5 To add additional advertised routes, repeat steps 1 through 4.

Routes to Reach Other Sites

When you specify the routes to reach other destinations, you create a list of those static routes to put on the CE that describe the remote destination networks in the VPN that you want the CE to reach.

To specify the routes that reach specific remote destinations in the VPN:

Step 1 **Double-click** the *Routes to Reach Other Sites* cell.

Step 2 In the dialog box that appears, enter the IP address and subnet mask of the first remote destination.

Step 3 Click **Add**.

The address is added to the list of remote destinations.

Step 4 Click **OK**.

You return to the Profile Editor, where the specified route is displayed in the *Routes to Reach Other Sites* cell.

Step 5 To add additional remote destination addresses, repeat steps 1 through 4.

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.

Give Only Default Routes to 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.

When you enable the **Give Only Default Routes to CE** option for RIP, VPN Solutions Center 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's RIP routing protocol. The RIP routes on the PE to the CE's 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. This option should *not* be used if the CE's site needs a default route for any reason, such as having a separate Internet feed.

Redistribute Static (BGP and RIP)

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).

Redistribute Connected (BGP Only)

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.

Redistribute OSPF

When you enable the **Redistribute OSPF (RIP only)** option for RIP, VPNSC imports the OSPF routes into the PE that is running RIP.

OSPF Process ID for RIP

Click into the *OSPF Process ID for RIP* cell, then enter the OSPF process ID.

Redistributed Protocols

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 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 from:

Step 1 **Double-click** the *RIP: Redistributed Protocols* cell.

The Redistributed Protocols dialog box is displayed (see Figure 5-32).

Figure 5-32 Specifying Protocols Redistributed Into RIP

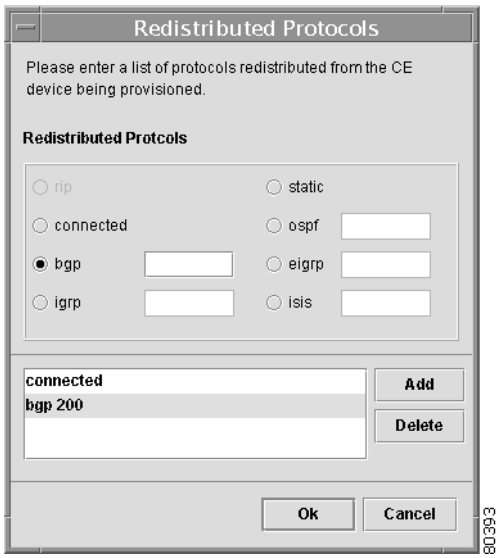

Step 2 Select the routing protocol to import routing information from:

- **Connected**

  Redistribute Connected 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.

- **BGP (Border Gateway Protocol)**

  If you are importing BGP-learned routes into RIP, select **bgp**, then enter the BGP autonomous system (AS) number in the field provided.

- **IGRP (Interior Gateway Routing Protocol)**

  If you are importing IGRP-learned routes into RIP, select **igrp**, then enter the IGRP autonomous system (AS) number in the field provided.

- **Static**

  If you are importing the static routes into RIP, select **static**.

- **OSPF (Open Shortest Path First)**

  If you are importing OSPF-learned routes into RIP, select **ospf**, then enter the OSPF process number in the field provided.

- **EIGRP (Enhanced IGRP)**

  If you are importing EIGRP-learned routes into RIP, select **eigrp**, then enter the EIGRP autonomous system (AS) number in the field provided.

- **IS-IS (Intermediate System-to-Intermediate System)**

  If you are importing IS-IS-learned routes into RIP, select **isis**, then enter the ISIS tag number in the field provided.

Step 3 Click **Add**.

The selected redistributed protocol is added to the list displayed in the Redistributed Protocols dialog box.

Step 4 To add additional redistributed protocols, repeat steps 2 and 3.

Step 5 When finished defining the redistributed protocols, click **OK**.

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.

BGP AS ID

Click in the *BGP AS ID* cell, then enter the BGP autonomous system number for the customer's BGP network.

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

Neighbor Allow-AS In

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.

Neighbor AS Override

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.

Redistributed Protocols

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 from:

Step 1 **Double-click** the *BGP: Redistributed Protocols* cell.

The Redistributed Protocols dialog box is displayed (see Figure 5-33).

Figure 5-33 Specifying Protocols Redistributed Into BGP

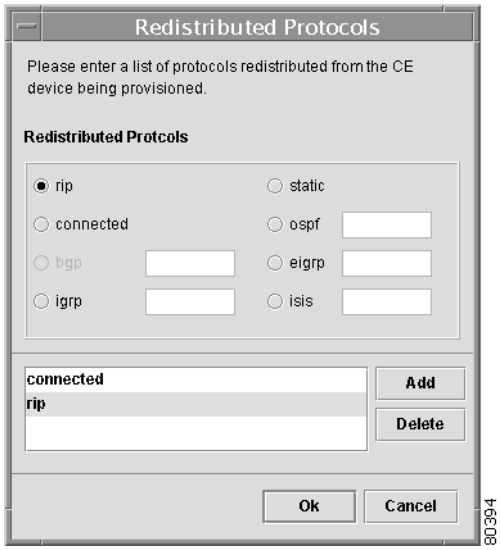

Step 2 Select the routing protocol to import routing information from:

- **RIP**

  If you are importing RIP routes into BGP, select **rip**.

- **Connected**

  Redistribute Connected 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.

- **IGRP (Interior Gateway Routing Protocol)**

  If you are importing IGRP-learned routes into BGP, select **igrp**, then enter the IGRP autonomous system (AS) number in the field provided.

- **Static**

  If you are importing the static routes into BGP, select **static**.

- **OSPF (Open Shortest Path First)**

  If you are importing OSPF-learned routes into OSPF select **ospf**, then enter the OSPF process number in the field provided.

After the VPN customer routes have been placed into the receiving VRF, they must be advertised to other PEs via MP-iBGP. This behavior is not automatic, so redistribution between OSPF and BGP is required.

- **EIGRP (Enhanced IGRP)**

  If you are importing EIGRP-learned routes into BGP, select **eigrp**, then enter the EIGRP autonomous system (AS) number in the field provided.

- **IS-IS (Intermediate System-to-Intermediate System)**

  If you are importing IS-IS-learned routes into BGP, select **isis**, then enter the ISIS tag number in the field provided.

Step 3 Click **Add**.

The selected redistributed protocol is added to the list displayed in the Redistributed Protocols dialog box.

Step 4 To add additional redistributed protocols, repeat steps 2 and 3.

Step 5 When finished defining the redistributed protocols, click **OK**.

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.

Redistribute RIP

When you enable the **Redistribute RIP (OSPF only)** option for OSPF, VPNSC redistributes the RIP routes into the PE that is running OSPF.

OSPF Process ID on PE

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

OSPF Process ID on CE

The OSPF process ID is a unique value assigned for each OSPF routing process within a single router—this 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.

OSPF Area Number on PE

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.

OSPF Area Number on CE

An *area* in OSPF terms is a grouping of contiguous OSPF networks and hosts. OSPF areas are logical subdivisions of OSPF autonomous systems. The topology of each area is invisible to entities in other areas, and each maintains its own topological database.

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

Redistributed Protocols

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 routing information from:

Step 1 **Double-click** the *OSPF: Redistributed Protocols* cell.

The Redistributed Protocols dialog box is displayed (see Figure 5-34).

Figure 5-34 Specifying Protocols Redistributed Into OSPF

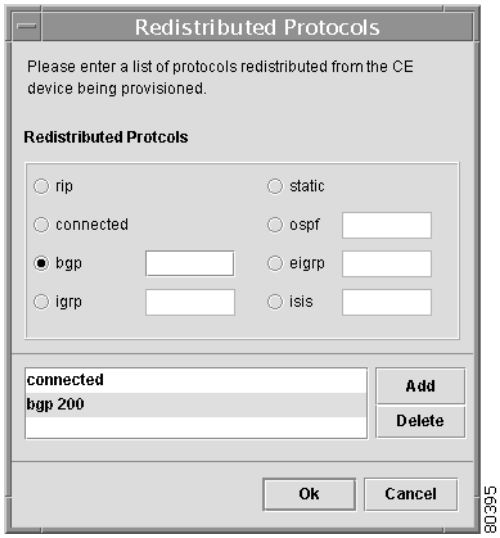

Step 2 Select the routing protocol to import routing information from:

- **RIP**

  If you are importing RIP routes into OSPF, select **rip**.

- **Connected**

  Redistribute Connected 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.

- **BGP (Border Gateway Protocol)**

  If you are importing BGP-learned routes into BGP, select **bgp**, then enter the BGP autonomous system (AS) number in the field provided.

- **IGRP (Interior Gateway Routing Protocol)**

  If you are importing IGRP-learned routes into OSPF, select **igrp**, then enter the IGRP autonomous system (AS) number in the field provided.

- **Static**

  If you are importing the static routes into OSPF, select **static**.

- **EIGRP (Enhanced IGRP)**

  If you are importing EIGRP-learned routes into BGP, select **eigrp**, then enter the EIGRP autonomous system (AS) number in the field provided.

- **IS-IS (Intermediate System-to-Intermediate System)**

  If you are importing IS-IS-learned routes into OSPF, select **isis**, then enter the ISIS tag number in the field provided.

Step 3 Click **Add**.

The selected redistributed protocol is added to the list displayed in the Redistributed Protocols dialog box.

Step 4 To add additional redistributed protocols, repeat steps 2 and 3.

Step 5 When finished defining the redistributed protocols, click **OK**.

Interface Addresses

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 VPN Solutions Center 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.

IP Numbering Scheme

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

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 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**, VPN Solutions Center creates a static route for the PE-CE link.

When you choose **IP unnumbered**, VPN Solutions Center automatically creates a loopback interface (unless a loopback interface already exists with the correct attributes). For related information, see the next section, "Using an Existing Loopback Interface Number."

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 VPN Solutions Center software to use the indicated addresses.

- **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, VPN Solutions Center 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, VPNSC uses those addresses (and does not allocate them from the address pool). If the IP addresses are not present in the configuration file, VPNSC picks IP addresses from a /30 subnet point-to-point IP address pool.

Using an Existing Loopback Interface Number

On each PE, there is one loopback interface number per VRF for interfaces using IP unnumbered addresses. By default, VPN Solutions Center software assigns a loopback number associated with a particular loopback address.

However, if a service provider wants VPN Solutions Center 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 this example of a router configuration file:

```
interface Loopback0
description Provisioned by VPN-SC
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.

Note Unlike standard interfaces, when loopback interfaces are provisioned in VPNSC, 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.

Extra CE Loopback Required

Even though a numbered IP address does not require a loopback address, VPN Solutions Center software provides the option to specify **Extra CE Loopback 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 (see the "Extra CE Loopback Required" section on page 5-42).

Automatically Assign IP Address

If you choose **IP unnumbered** and also enable the **Automatically Assign IP Address** check box, VPN Solutions Center 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, VPN Solutions Center 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, VPNSC uses those addresses (and does not allocate them from the address pool). If the IP addresses are not present in the configuration file, VPNSC picks IP addresses from a /30 subnet point-to-point IP address pool.

PE Interface Address

If you do not automatically assign IP addresses, enter the IP address and subnet mask for the PE interface.

Step 1 From the Profile Editor, **double-click** the *PE Interface Address* cell.

The PE Interface Address dialog box appears (see Figure 5-35).

Figure 5-35 Entering the PE's IP Address

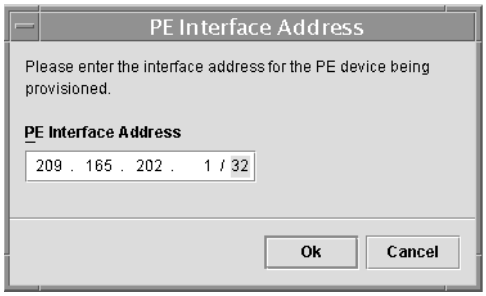

Step 2 Enter the PE's IP address and subnet mask, then click **OK**.

The IP address is entered into the *PE Interface Address* cell.

CE Interface Address

If you do not automatically assign IP addresses, enter the IP address and subnet mask for the CE interface.

Step 1 From the Profile Editor, **double-click** the *CE Interface Address* cell.

The CE Interface Address dialog box appears (see Figure 5-36).

Figure 5-36 Entering the CE's IP Address

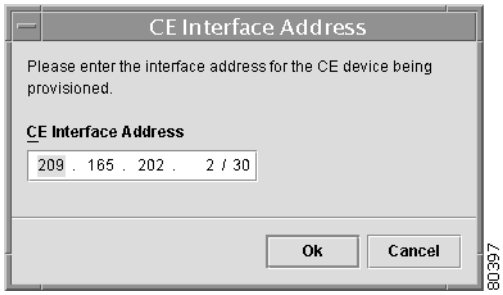

Step 2 Enter the CE's IP address and subnet mask, then click **OK**.

The IP address is entered into the *CE Interface Address* cell.

Extra CE Loopback Address

This option places an IP address on a CE router that is not tied to any physical interface. If you enable the **Extra CE Loopback Required** checkbox**,** the *Extra CE Loopback Address* cell is enabled, and you can enter the IP address for the extra CE loopback interface.

Step 1 From the Profile Editor, **double-click** the *Extra CE Loopback Address* cell.

The Extra CE Loopback Address dialog box appears (see Figure 5-37).

Figure 5-37 Entering the CE's Loopback Address

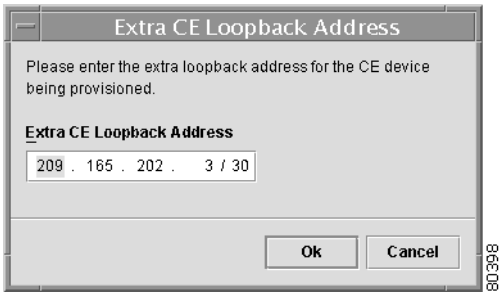

Step 2 Enter the IP address and subnet mask for the CE's extra loopback interface, then click **OK**.

The IP address is entered into the *Extra CE Loopback Address* cell.

VRF Maps

The mechanism by which the originating PE can selectively advertise certain routes with a route target that is different than the other routers that it is also advertising from the same VRF is the *export map*.

In some networks, the service providers base the connectivity toward their customers on static routes configured on PE routers. If the customer is a VPN customer, the route is configured within the context of the customer VRF. To advertise these static routes to other members of the VPN or to other VPNs,

they must be redistributed into MP-iBGP (Multi-Protocol internal BGP). In this environment, the PE configuration can be simplified significantly by using a tag that can be associated with the static route to signify the set of VPNs into which the route should be exported.

When the static routes toward the VPN customer are redistributed into MP-iBGP, a route map is used to match the static route tag and set a corresponding standard BGP community. When the standard BGP community is set, it can then be used to specify which route target should be used for the route when it is exported through MP-iBGP to other PEs.

Export Map

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

The *Export 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 VPN Solutions Center software to define a management VPN, the software 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.

For information on the **route-map** command, refer to the Cisco IOS documentation on IP routing protocol-independent commands.

Import Map

Import Map: Enter the name of the import map.

The *Import 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 this PE to exclude the route.

For command reference details on the **import map** command, see the "import map" section on page B-4.

For information on customizing the VRF name or the RD values, see the "Overriding the Default VRF Name and Route Distinguisher Values" section on page 6-17.

Maximum Routes

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

NetFlow

Turn on NetFlow Accounting

Enable this option if you are using CNS/PerfE (CNS Performance Engine).

Templates

VPN Solutions Center provides a way to integrate a template with VPN Solutions Center configlets. For information on creating and employing VPN Solutions Center templates, see Chapter 10, "Provisioning with the VPN Solutions Center Template Manager."

For a given customer edge router, you specify the following:

- Template name
- Template data file name
- Whether the Template configuration file should be appended or prepended to the VPN Solutions Center configlet
- Whether the Template configuration file is active or inactive for downloading to the edge device

The template data files are tightly linked with its corresponding template. You can use a data file and its associated template to create a template configuration file. The template configuration file is merged with (either appended to or prepended to) the VPN Solutions Center configlet. VPN Solutions Center downloads the combined VPN Solutions Center configlet and template configuration file to the edge device router.

- You can download a template configuration file to a router. For details, see the "Provisioning a Template Configuration File Directly to a Router" section on page 10-24.
- You can apply the same template to multiple edge routers, assigning the appropriate template data file for each device. Each template data file includes the specific data for a particular device (for example, the management IP address or host name of each device).

PE Template

Specify which template data file you want associated with the PE in the current link:

Step 1 From the Profile Editor, **double-click** the *PE Template* cell.

The initial PE Templates dialog box appears (see Figure 5-38).

You first specify which template data file you want associated with the PE in the current link.

Figure 5-38 Initial Dialog Box for Integrating a PE Template

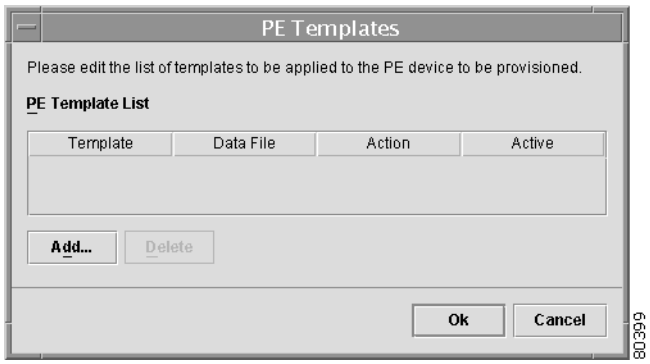

Step 2 From the PE Templates dialog box, click **Add**.

The Template Chooser dialog box appears (see Figure 5-39).

Figure 5-39 Choosing the Template Data Files for the PE

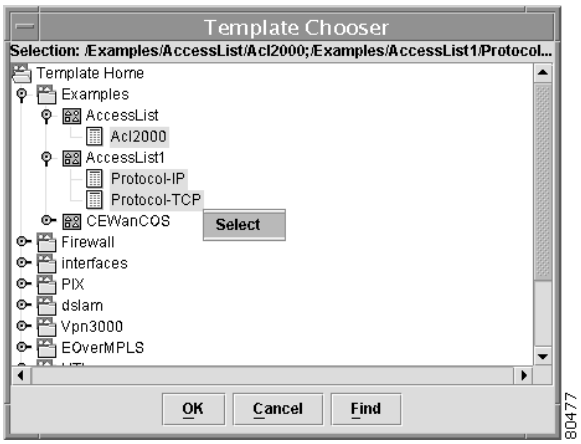

Step 3 Expand the Template Home hierarchy until you can see the pertinent template names and their data files.

The template data files are tightly associated with its corresponding template. You can use a data file and its associated template to create a template configuration file. The template configuration file is merged with (either appended to or prepended to) the VPN Solutions Center configlet. VPN Solutions Center downloads the combined configlet to the customer edge router.

Step 4 Select one or more template data files.

To select multiple template data files:

a. Select the first data file of interest.

b. To select additional template data files, press **Ctrl+Click**.

c. When the template data files are selected, **right-click**.

The **Select** option appears (see Figure 5-39).

d. Choose **Select**.

e. Click **OK**.

You return to the PE Templates dialog box, which now displays the current template selection (see Figure 5-40).

Figure 5-40 PE Templates Added to the Service Request Profile

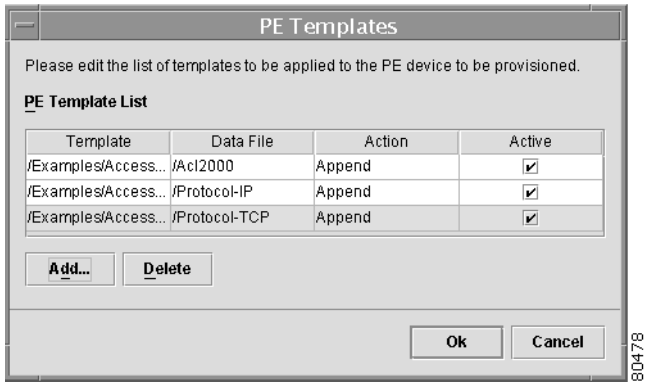

Determining the Placement and Active Status of the PE Template Data File

The **Action** column in the dialog box lets you specify where the template configuration file is placed in the VPN Solutions Center configlet—either *prepended* (placed at the front of the configlet) or *appended* (placed at the end of the configlet).

The **Active** column lets you determine whether you want the template configuration file to be merged with the VPN Solutions Center configlet and downloaded to the target router.

Step 5 To specify the placement of the template configuration file, click the *Action* field for the appropriate template, then choose **Append** (the default) or **Prepend**.

- If you choose **Append**, the template configuration file is appended to (that is, placed at the end of) the VPN Solutions Center configlet prior to being downloaded to the target provider edge router.
- If you choose **Prepend**, the template configuration file is prepended to (that is, placed at the beginning of) the VPN Solutions Center configlet prior to being downloaded to the target provider edge router.

Step 6 Specify the Active status of the template configuration file.

- If you enable the Active checkbox, the template configuration file is merged with the VPN Solutions Center configlet and downloaded to the target router.
- If you disable the Active checkbox, the template configuration file is not merged with the VPN Solutions Center configlet.

Step 7 When the PE templates fields are set to your satisfaction, click **OK**.

The templates you selected are added to the Service Request Profile.

CE Template

Specify which template data files you want associated with the CE:

Step 1 From the Profile Editor, **double-click** the *CE Template* cell.

The initial CE Templates dialog box appears (see Figure 5-41).

You first specify which template data file you want associated with the CE in the current link.

Figure 5-41 Initial Dialog Box for Integrating a Template

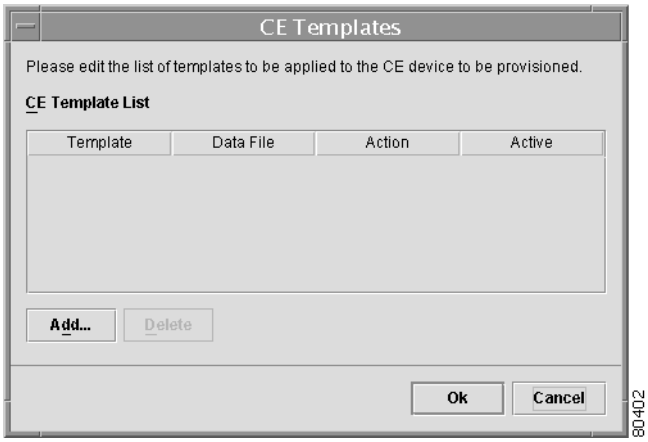

Step 2 From the CE Templates dialog box, click **Add**.

The Template Chooser dialog box appears (see Figure 5-42).

Figure 5-42 Choosing the Template Data Files for the CE

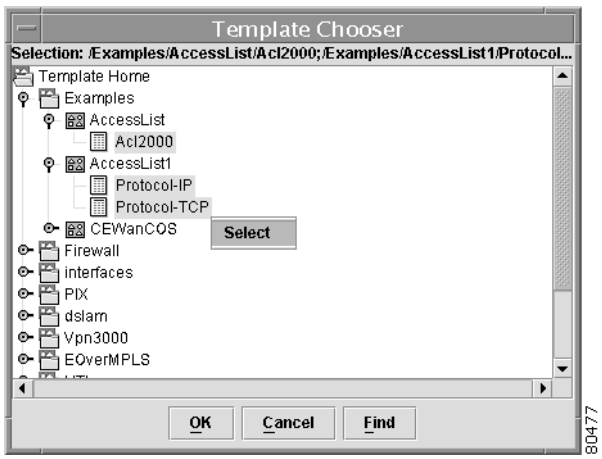

Step 3 Expand the Template Home hierarchy until you can see the pertinent template name and its data files.

The template data files are tightly associated with its corresponding template. You can use a data file and its associated template to create a template configuration file. The template configuration file is merged with (either appended to or prepended to) the VPN Solutions Center configlet. VPN Solutions Center downloads the combined configlet to the customer edge router.

Step 4 Select one or more template data files.

To select multiple template data files:

a. Select the first data file of interest.

b. To select additional template data files, press **Ctrl+Click**.

c. When the template data files are selected, **right-click**.

The **Select** option appears (see Figure 5-42).

d. Choose **Select**.

e. Click **OK**.

You return to the CE Templates dialog box, which now displays the current template selection (see Figure 5-43).

Figure 5-43 CE Templates Added to the Service Request Profile

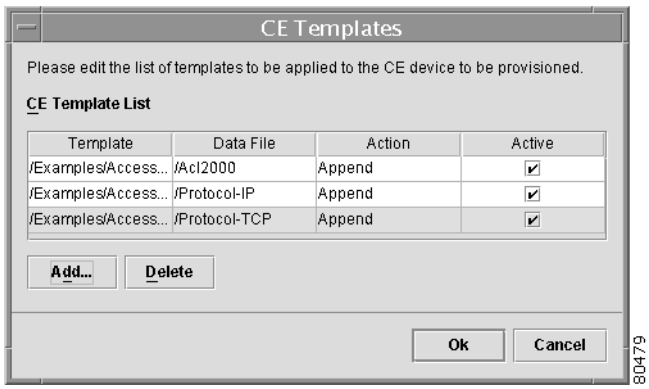

Determining the Placement and Active Status of the CE Template Data File

The **Action** column in the dialog box lets you specify where the template configuration file is placed in the VPN Solutions Center configlet—either *prepended* (placed at the front of the configlet) or *appended* (placed at the end of the configlet).

The **Active** column lets you determine whether you want the template configuration file to be merged with the VPN Solutions Center configlet and downloaded to the target router.

Step 5 To specify the placement of the template configuration file, click the *Action* field for the appropriate template, then choose **Append** (the default) or **Prepend**.

- If you choose **Append**, the template configuration file is appended to (that is, placed at the end of) the VPN Solutions Center configlet prior to being downloaded to the target customer edge router.
- If you choose **Prepend**, the template configuration file is prepended to (that is, placed at the beginning of) the VPN Solutions Center configlet prior to being downloaded to the target customer edge router.

Step 6 Specify the Active status of the template configuration file.

- If you enable the Active checkbox, the template configuration file is merged with the VPN Solutions Center configlet and downloaded to the target router.
- If you disable the Active checkbox, the template configuration file is not merged with the VPN Solutions Center configlet.

Step 7 When the CE templates fields are set to your satisfaction, click **OK**.

The templates you selected are added to the Service Request Profile.

What's Next?

When you have defined the VPNs and created the Service Request Profiles that you will need for your Customers, you are ready to provision service requests. Refer to the next chapter, Chapter 6, "Provisioning MPLS VPN Service Requests," for the procedures in VPN Solutions Center to create and deploy the service requests for your network.

CHAPTER 6

Provisioning MPLS VPN Service Requests

The focus of the VPN Solutions Center product is the service provided for a customer on the link between the customer's CE and the provider's PE. This chapter describes how you create a service request in the VPN Solutions Center software, as well as how to modify and delete service requests.

Finally, this chapter tells you how to check on a service request's status and find out what went wrong if the request failed.

The main topics presented in this chapter are as follows:

Service Request Summary

The service model is the centerpiece of service provisioning. With the service model, the VPN Solutions Center software can capture the specified VPN service provisioning request, analyze the validity of the request, and audit the provisioning results.

The service provider operators take all service request information from their customers. VPN Solutions Center can assist the operator in making entries because the product has customer information such as the VPN information, the list of the assigned PEs and CEs, and so forth.

The VPN Console steps the operator through the process and simplifies the task of provisioning the CE and PE by automating most of the tasks required to set up an MPLS VPN.

Figure 6-1 shows a high-level diagram of the relationships and movement among VPN Solutions Center service request states.

Figure 6-1 Service Request States: Movement and Relationships

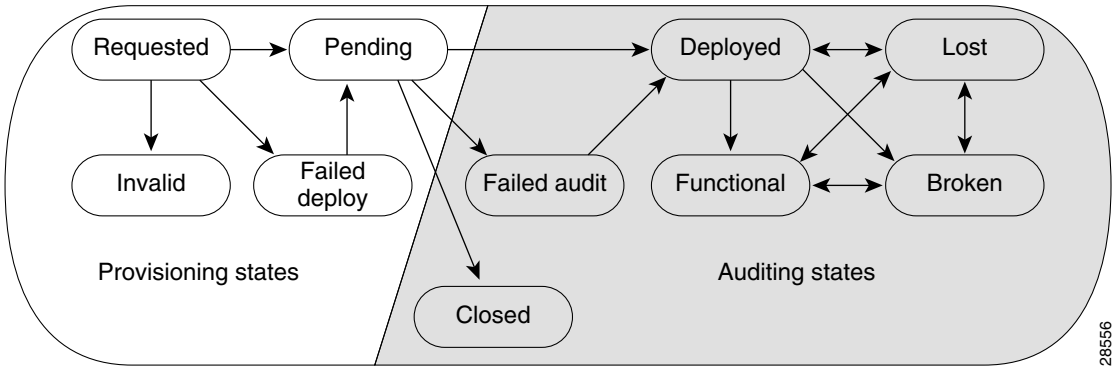

The sections below describe each of the service request states and their transition sequences.

Definitions of VPN Solutions Center Service Request States

Table 6-1 describes the functions of each VPN Solutions Center service request state. They are listed in alphabetical order.

Table 6-1 Summary of VPN Solutions Center Service Request States

Service Request Type	Description
Broken	While the router is correctly configured, the service is unavailable (due to a broken cable or Layer 2 problem, for example). A service request moves to Broken if the Auditor finds the routing and forwarding tables for this service, but they do not match the service intent.
Closed	A service request moves to Closed if the service request should no longer be used during the provisioning or auditing process. A service request moves to the Closed state only upon a successful audit of a remove request. VPN Solutions Center does not remove a service request from the database to allow for extended auditing. Only a specific administrator action results in service requests being removed.
Deployed	A service request moves to Deployed if the configlet commands have been verified as found in the router configuration file. Deployed indicates that the configuration file on the router matches the information specified in the VPNSC service request.
Failed Audit	The Failed Audit state indicates that VPNSC downloaded the configlet to the router successfully, but the service request did not pass the audit. Therefore, the service did not move to either the Functional or Deployed state. The Failed Audit state is initiated from the Pending state. Once a service request is deployed successfully, it cannot reenter the Failed Audit state (except when the service request is redeployed).

Table 6-1 Summary of VPN Solutions Center Service Request States (continued)

Service Request Type	Description
Failed Deploy	After provisioning occurred, the service request failed to download the configuration updates to the router. A service request moves to Failed Deploy if the Telnet Gateway Server (TGS) detected an error during the deployment process. If TGS is not being used to download configuration updates, and VPNSC is simply exporting configuration updates to a directory, there is no way to distinguish between a service request in the Failed Deploy and Pending states. The cause for a Failed Deploy status is that TGS reports that either the upload of the initial configuration file from the routers failed or the download of the configuration update to the routers failed (due to lost connection, faulty password, etc.). If the configuration updates are exported to a directory, the service request cannot move into a Failed Deploy state.
Functional	A service request moves to Functional when the Auditor finds the VPN routing and forwarding tables (VRF) for this service and they match with the service intent. This state requires that both the configuration file audit and the routing audit are successful.
Invalid	Indicates that the service request information is incorrect in some way. A service request moves to Invalid if the request was either internally inconsistent or not consistent with the rest of the existing network/router configurations (for example, no more interfaces were available on the router). The Provisioning Driver cannot generate configlets to service this request.
Lost	A service request moves to Lost when the Auditor cannot find a configuration-level verification of intent in the router configuration files. The service request was deployed, but now some or all router configuration information is missing. A service request can move to the Lost state *only* when the service request had been Deployed or Functional.
Pending	A service request moves to Pending when the Provisioning Driver determines that the request looks consistent and was able to generate the required configlets for this request. Pending indicates that the service request has generated the configlets and the configlets are successfully downloaded to the routers. The Auditor regards pending service requests as new requests and begins the audit. If the service has been freshly provisioned and not yet audited, it is not an error (pending audit). However, if an audit is done and the service is still pending, it is in an error state.
Requested	If the service is newly entered and not yet deployed, it is not an error. However, if a Deploy is done and it remains Requested, the service is in an error state.

Service Request State Transition Sequences

Table 6-2 on page 6-4 and Table 6-3 on page 6-5 show the state transition paths for VPN Solutions Center service requests. The beginning state of a service request is listed in the first column; the states that service requests transition to are displayed in the heading row.

For example, to use Table 6-2 to trace the state of a Pending service request to Functional, find "**Pending**" in the first column and move to your right until you find "**Functional**" in the heading. You can see that for a service request to move from Pending to Functional, a successful routing audit must take place.

Table 6-2 shows the service request transitions from *Requested* to *Lost*.

Table 6-2 State Transition Paths for VPN Solutions Center Service Requests (Part 1)

Service Request States	Requested	Pending	Failed Audit	Deployed	Functional	Lost
Requested	No transition to Requested	Successful service request deployment	No transition to Failed Audit	No transition to Deployed	No transition to Functional	No transition to Lost
Pending	No transition to Requested	—Successful service request deployment —Audit with error	Audit is not successful	Audit is successful	Routing audit is successful	No transition to Lost
Failed Audit	No transition to Requested	Successful service request redeployment	No transition to Failed Audit	Audit is successful	Routing audit is successful	No transition to Lost
Deployed	No transition to Requested	Successful service request redeployment	No transition to Failed Audit	Audit is successful	Routing audit is successful	Audit found error
Functional	No transition to Requested	Successful service request redeployment	No transition to Failed Audit	No transition to Deployed	Routing audit is successful	Audit found error
Lost	No transition to Requested	Successful service request redeployment	No transition to Failed Audit	Audit is successful	Routing audit is successful	Audit found error
Broken	No transition to Requested	Successful service request redeployment	No transition to Failed Audit	No transition to Deployed	Routing audit is successful	Audit found error
Invalid	No transition to Requested	Successful service request redeployment	Redeployment caused service request error	No transition to Deployed	No transition to Functional	No transition to Lost

Table 6-2 State Transition Paths for VPN Solutions Center Service Requests (Part 1) (continued)

Service Request States	Requested	Pending	Failed Audit	Deployed	Functional	Lost
Failed Deploy	No transition to Requested	Successful service request redeployment	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Deployed	No transition to Functional	No transition to Lost
Closed	No transition to Requested	No transition to Pending	No transition to Failed Audit	No transition to Deployed	No transition to Functional	No transition to Lost

Table 6-3 shows the service request transitions from *Broken* to *Closed*.

Table 6-3 State Transition Paths for VPN Solutions Center Service Requests (Part 2)

Service Request States	Broken	Invalid	Failed Deploy	Closed
Requested	No transition to Broken	Deploy Service Request error	Deployment failed	No transition to Closed
Pending	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	Removal of the service request is successful
Failed Audit	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Deployed	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Functional	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Lost	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Broken	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Invalid	No transition to Broken	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Failed Deploy	No transition to Broken	Redeploy service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Closed	No transition to Broken	No transition to Invalid	No transition to Failed Deploy	No transition to Closed

How VPNSC Accesses Network Devices

When VPN Solutions Center attempts to access a router, it uses the following algorithm:

1. Check to see if a terminal server is associated with the device, and if this is the case, VPNSC uses the terminal server to access the device.
2. If there is no terminal server, VPN Solutions Center looks for the management interface on the device.
3. If there is no management interface, it tries to access the device using the fully-qualified domain name (hostname plus domain name).

Note If any step in the VPN Solutions Center device-access algorithm fails, the entire device access operation fails—there is no retry or rollover operation in place. For example, if there is a terminal server and VPNSC encounters an error in attempting to access the target device through the terminal server, the access operation fails at that point. With the failure of the terminal server access method, VPNSC does not then attempt to find the management interface to access the target device.

Overview of the Service Request Process

Provisioning a VPN provides a method to build a service for site-to-site connectivity between a provider edge router and a customer edge router. It includes the following steps:

1. From the VPN Console, define a service request to add VPN service between a CE and PE.
2. Schedule to download the new configuration to the CE and PE pairs.
3. Use the reports available from the Provisioning menu to verify the service requests and view configlets.

The first step in provisioning a VPN is to define a *service request*. A service request defines through whom (the provider edge router) and to whom (the customer edge router) the service is provided. In this procedure, you specify the information required to provision the link between the PE and CE.

Adding a Service for a PE-CE Link

A service request is an instance of service contract between a customer edge router (CE) and a provider edge router (PE). The service request user interface asks you to enter several parameters, including the specific interfaces on the CE and PE routers, routing protocol information, and IP addressing information.

You can also integrate a VPN Solutions Center template with a service request. You can associate one or more templates to the CE and the PE.

To add one or more VPN services, follow these steps:

Step 1 From the VPN Console, choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor is displayed (see Figure 6-2).

Figure 6-2 The MPLS Service Request Editor

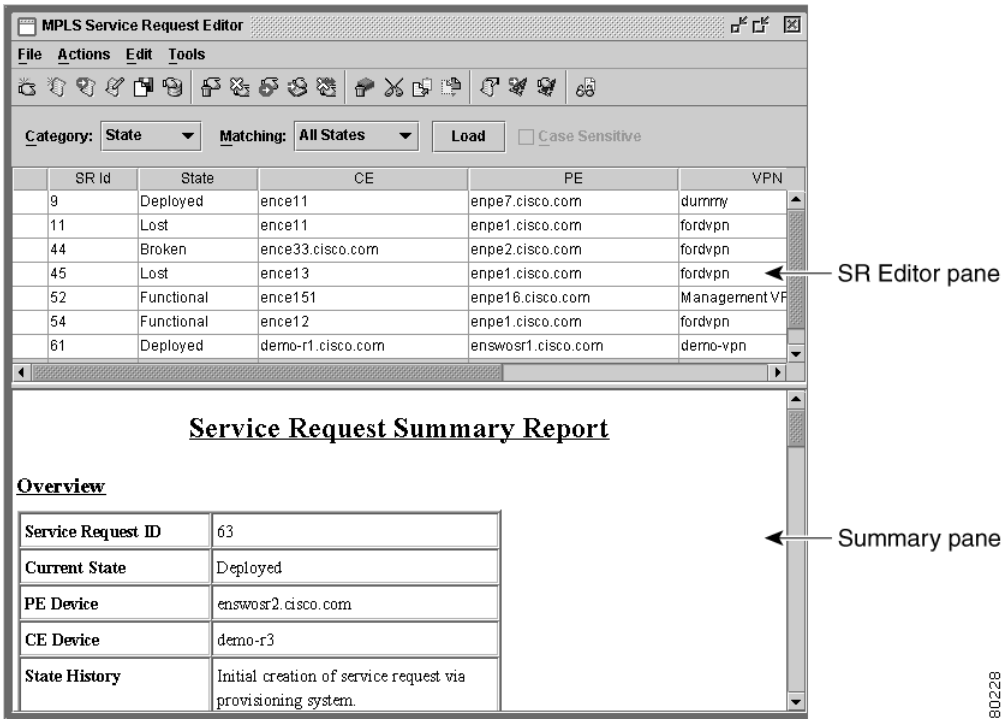

As you can see in Figure 6-2, the MPLS Service Request Editor consists of the following elements:

- Menu bar
- Tool bar
- Editor pane

  The Editor pane is where service request information is listed,
- Summary pane

Step 2 From the MPLS Service Request Editor menu bar, choose **File > New Service Request(s)**.

The Open Profile dialog box appears (see Figure 6-3).

Figure 6-3 Selecting a Service Request Profile

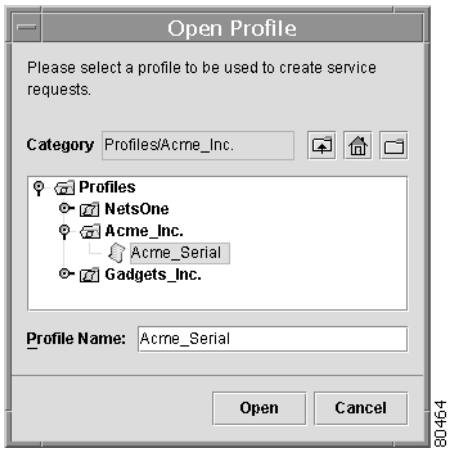

Step 3 From the list of service request profiles, select the appropriate profile for this service.

Step 4 Click **Open**.

The CE Device Chooser appears (see Figure 6-4).

Selecting the CEs for the Service Request

Figure 6-4 Selecting the CEs

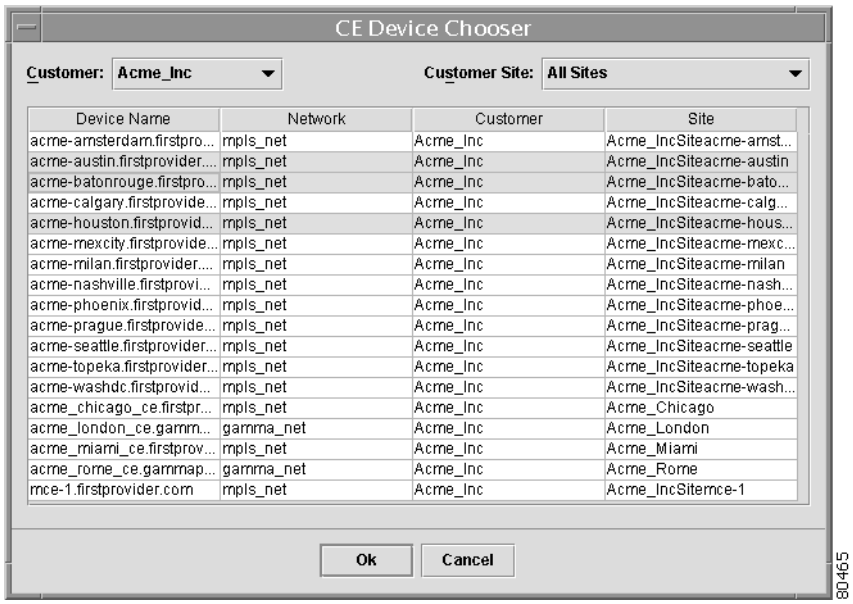

Step 5 Select the CE(s) for this service request.

a. *Customer*: From the Customer drop-down list, select the name of the VPN Customer.

b. *Customer Site*: From the Customer Site drop-down list, select the name of the site you want to see, or choose **All Sites** to see the list of all the sites for the selected Customer.

c. Select one or more CEs from the CE Device Chooser.

To select multiple CEs, press **Ctrl+Click** for each additional CE.

d. Click **OK**.

You return to the Service Request Editor, where the CE information is now displayed (see Figure 6-5).

Figure 6-5 CEs Selected for the Service Request

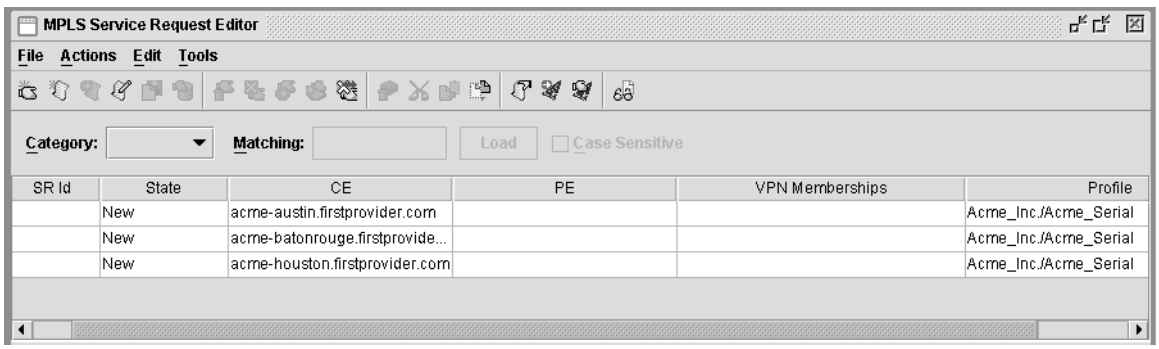

Selecting the PE for the Service Request

Now that you have specified one or more CEs for the service request, the next step is to specify the PE for the PE-CE link. You can specify multiple CEs for a single PE.

Step 6 In the Editor pane of the MPLS Service Request Editor, select one or more CEs that you want to associate with a PE.

Step 7 Choose **Actions > Set PE**.

The PE Device Chooser appears (see Figure 6-6).

Figure 6-6 Selecting the PE

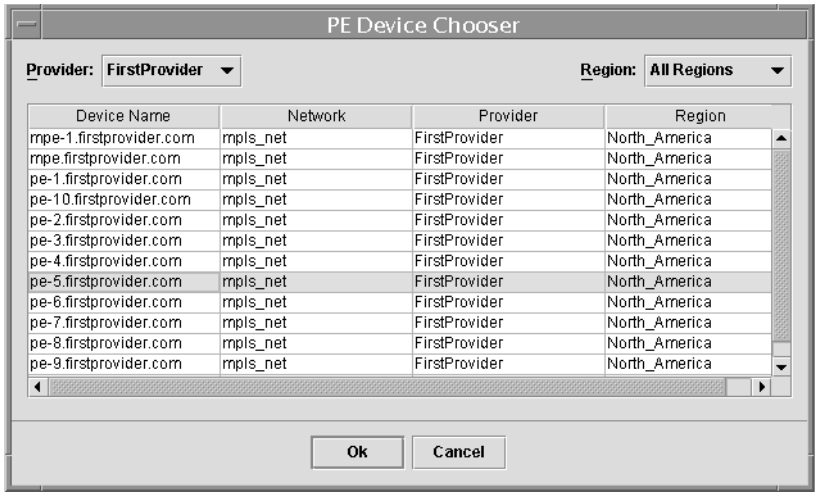

Step 8 Select the PE for this service request.

a. *Provider*: From the Provider drop-down list, select the name of the service provider.

b. *Region*: From the Region drop-down list, select the name of the region the PE is in, or choose **All Regions** to see the list of all the PEs for the selected provider.

c. Select a PE from the PE Device Chooser.

d. Click **OK**.

If you're creating multiple service requests, you will receive the following **Set PE** informational prompt:

```
Changed the PE device for x service requests.
```

e. Click **OK**.

You return to the Service Request Editor, where the name of the selected PE is now displayed (see Figure 6-7).

Figure 6-7 PE Specified for Multiple Service Requests

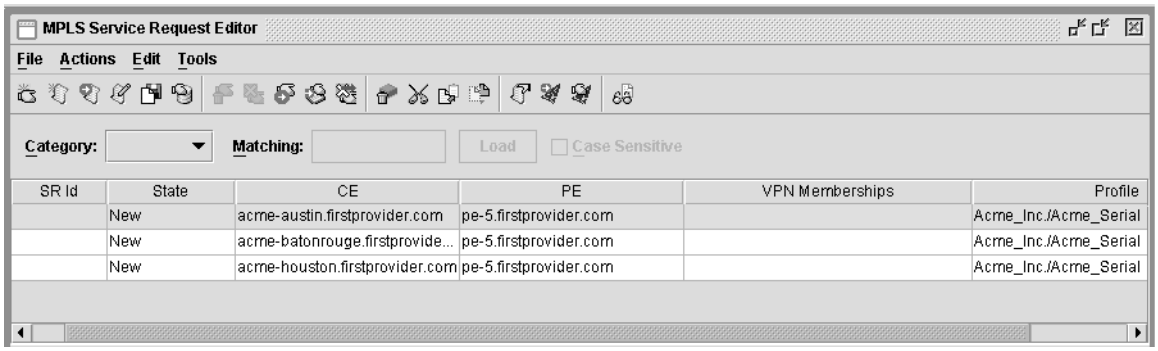

As you can see in Figure 6-7, a single PE (*PE-5.firstprovider.com*) is specified for the three CEs selected for these service requests.

About the VPN Membership Information

In this step, you specify the VPN that the selected PE-CE pairs belong to. Then specify the role that the selected CE (that is, the site) has in the VPN: whether the site should be able to communicate with all the other sites in the VPN, or whether you want the site to always communicate through the hub and not have access to all the other sites in the VPN.

The most common types of VPNs are *hub-and-spoke* and *full mesh*. These two basic types of VPNs—full mesh and hub and spoke—can be represented with a single CERC. This is controlled by the route target (RT) import and export statements in the Virtual Route Forwarding (VRF) definition.

- A *hub* receives routes from all members of the CERC, and its routes are distributed to all members of the CERC.

  A hub VRF imports and exports routes to the hub RT and imports routes from the spoke RT. This allows each hub to have routes from all the hubs and spokes in the CERC.

- A *spoke* receives routes from all hubs in the CERC, and its routes are sent only to hubs.

  A spoke VRF exports routes to the spoke RT and imports routes from the hub RT. This allows all hubs to receive the spoke's routes and allows the spoke to receive the hub routes, but the spoke does not receive other spoke routes.

- A hub-and-spoke CERC is one in which one or a few CEs act as hubs, and all spoke CEs talk only to or through the hubs, never directly to each other.

- A full mesh CERC is one in which every CE connects to every other CE.

Selecting the **Join as Spoke** option controls the VRF definition, as well as the RT import and export for the site being added.

For additional information on CE routing communities, see the "CE Routing Communities" section on page 1-18 and the "Defining CE Routing Communities" section on page 5-3.

Specifying the VPN Membership Information

Step 9 In the Editor pane of the MPLS Service Request Editor, select one or more of the PE-CE pairs you want to associate with a particular VPN.

Step 10 Choose **Actions > Set VPN Memberships**.

The VPN Memberships dialog box appears (see Figure 6-8).

Figure 6-8 Specifying the VPN Membership for the Devices

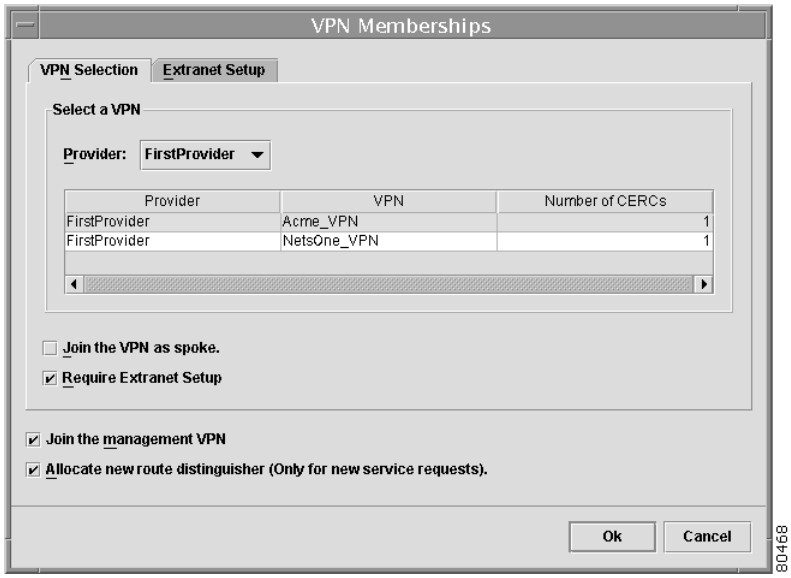

Specifying the VPN

Step 11 Select the VPN for this service request.

a. *Provider*: From the Provider drop-down list, select the name of the service provider.

b. From the list of VPNs displayed for the selected service provider, select the appropriate VPN.

Specifying a Hub-and-Spoke or Full Mesh VPN

Step 12 If you are building a VPN with a hub-and-spoke topology, enable the **Join as Spoke** option.

- If you want the CEs to *not* have access to all the other sites in the VPN, be sure to enable the **Join as spoke** option.
- If you want the CEs to have access to all the other sites in the VPN, do *not* enable the **Join as spoke** option.

Joining the Management VPN

Step 13 If you are adding a CE to the *management VPN*, enable the **Join the management VPN** option.

For more information on the management VPN, see Chapter 8, "The VPNSC Management Network."

When you use the VPN Solutions Center software to define a management VPN, the software automatically generates an *export route map* for the management VPN.

Set Up Load Balancing for a Multihomed CE

The **Allocate new route distinguisher** option allows the selected PEs to have unique route distinguishers (RDs) within the current VPN. Having unique RDs on each PE lets BGP load balance the traffic in the case where a Customer's network has a dual-homed CE with links to two PEs.

Note By default, the **Allocate new route distinguisher** option is not displayed. To make this option available, in the *csm.properties* file, change the following property to **true**:
netsys.mpls.svrc.OverrideVRFCreatedByPE.unix

Figure 6-9 shows a service provider BGP MPLS network that connects two networks (or subnets) to PE-1 and PE-2. Both of these PEs are configured for VPNv4 unicast iBGP peering. Network 2 is a multihomed subnet that is connected to both PE-1 and PE-2. Network 2 also has Extranet VPN services configured with Network 1. Both Network 1 and Network 2 are configured for eBGP peering with the PEs.

Figure 6-9 Unique RDs Assigned to PEs in the Same VPN

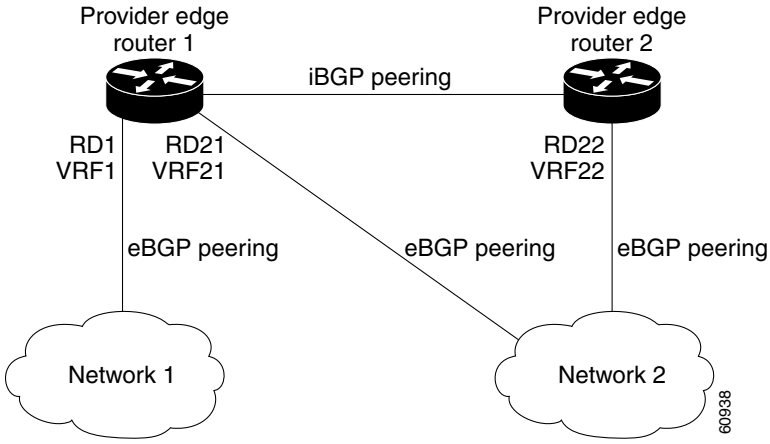

As shown in Figure 6-9, PE-1 and PE-2 have unique RDs (RD 21 and RD 22 respectively). The multipaths between Network 2 and PE-1 and PE-2 performs load balancing when the **Allocate new route distinguisher** option is enabled. Any prefix that is advertised from Network 2 will be received by RD 21 and RD 22. Thus, any traffic to Network 2 will be load balanced, with half the traffic going through PE-1 and half the traffic going through PE-2.

Step 14 To enable load balancing as described above, enable the **Allocate new route distinguisher** option.

If the VPN Has CEs in Other VPNs (Extranets)

Step 15 If you are building a VPN with CEs that are members of multiple VPNs (also referred to as *extranets*), enable the **Require Extranet Setup** option.

Extranet provisioning provides a way to create multiple VPN connectivity to a single VRF. You can add multiple CERCs to your VPN in any topology to form extranets. You can join an extranet in such a way that a CE can be a spoke in one VPN and a hub in another VPN.

- Enable the **Require Extranet Setup** option (by selecting the checkbox).

When you do so, the Extranet Setup tab is displayed (see Figure 6-10).

Figure 6-10 Extranet Setup

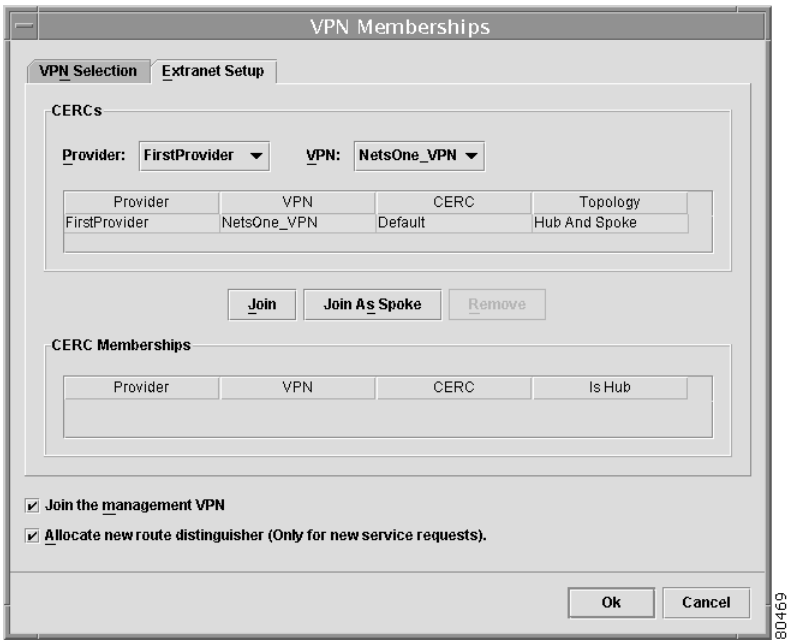

a. *Provider*: Specify the service provider name.

b. *VPN*: Specify the name of the other VPN that this CERC is a member of.

c. Specify whether the selected CERC is a hub or a spoke.

- If the selected CERC is a hub, click **Join**.

  The selected *hub CERC* is now displayed in the CERC Memberships panel. Note that the "Is Hub" checkbox is enabled (see Figure 6-11).

Figure 6-11 Hub CERC Added to Another VPN

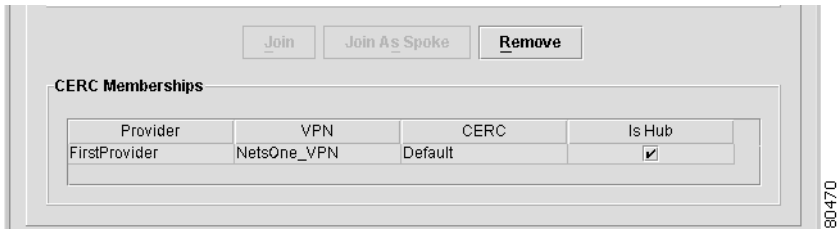

- If the selected CERC is a spoke, click **Join As Spoke**.

  The selected *spoke CERC* is now displayed in the CERC Memberships panel (see Figure 6-12)

Figure 6-12 Spoke CERC Added to Another VPN

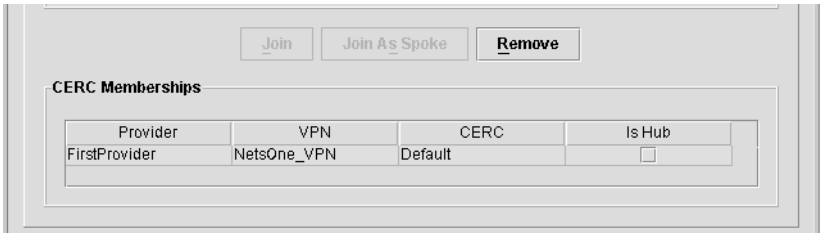

d. When satisfied with the Extranet settings, click **OK**.

If you're creating multiple service requests, you will receive the following **Set VPN** informational prompt:

```
Changed the VPN memberships for x service requests.
```

e. Click **OK**.

You return to the Service Request Editor, where all of the fields are filled in, indicating you are ready to deploy the service requests (see Figure 6-13).

Figure 6-13 Parameters Completed for Service Requests

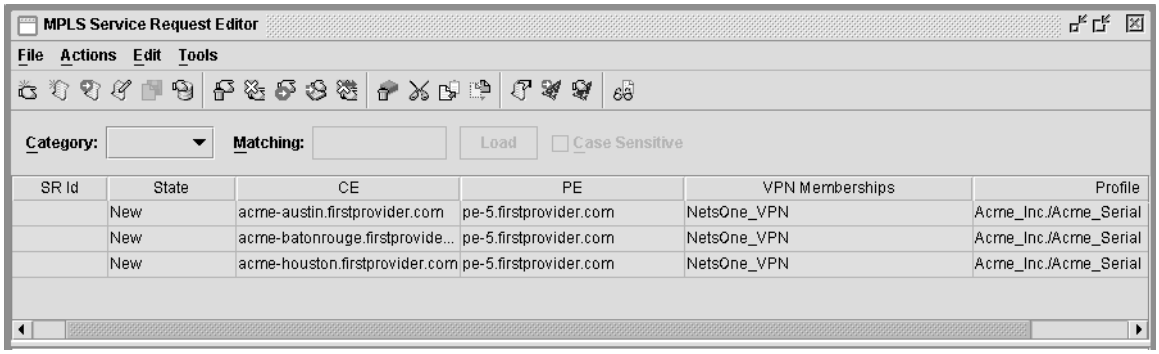

Tip At this point, you can customize the selected profile to accommodate the specific requirements of the service request (for example, to modify the PE or CE interfaces).
To do so, **double-click** the *Profile* cell and modify the service request profile as needed.

If You Wish to Delay Service Request Deployment

You are not required to immediately deploy the service requests you have set up. If you wish to delay service request deployment for any reason, you can commit to the Repository the existing service requests in their current state. You can either commit some or all of the existing service requests to the Repository.

To commit new service requests to the Repository:

1. If you wish to commit selected service requests only, from the MPLS Service Request Editor, select the service requests you want to commit.
2. Depending on whether you are committing some or all of the new service requests, do one of the following:
   - For selected service requests, choose **File > Commit to Repository**
   - For all the new service requests, choose **File > Commit All to Repository**.

The message bar at the bottom of the VPN Console displays the message:

```
Committed x service requests to the Repository.
```

VPN Solutions Center moves the selected service requests to the Requested state and assigns service request IDs to each (see Figure 6-14).

Figure 6-14 Service Requests Committed to the Repository

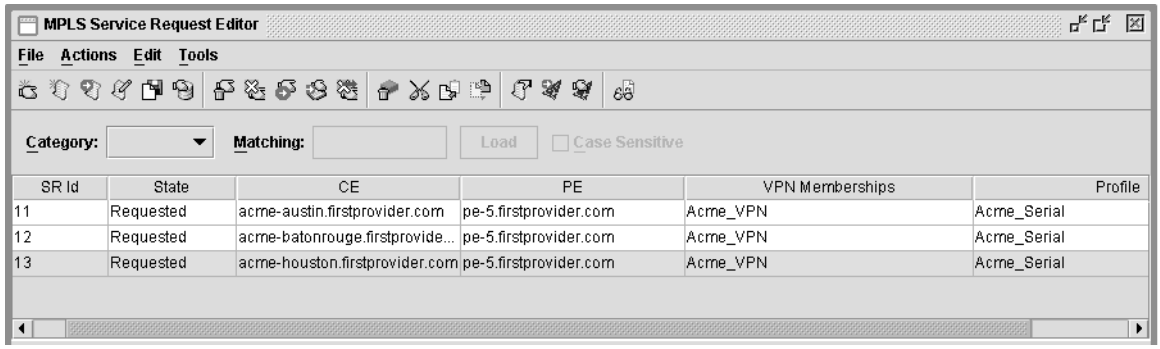

You can now deploy the committed service requests as your convenience.

Deploying Service Requests

When you have queued one or more service requests, you can then deploy them. This procedure automatically audits the new service requests. This audit passes the service request into an operational state.

VPN Solutions Center sets up a scheduled task that deploys service requests to the appropriate routers. This involves computing the configlets for each service request, downloading the configlets to the routers, and running audit reports to determine whether the service was successfully deployed.

You can choose to deploy the service requests immediately or schedule their deployment.

Step 1 From the VPN Console, choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor is displayed (see Figure 6-15).

Figure 6-15 MPLS Service Request Editor

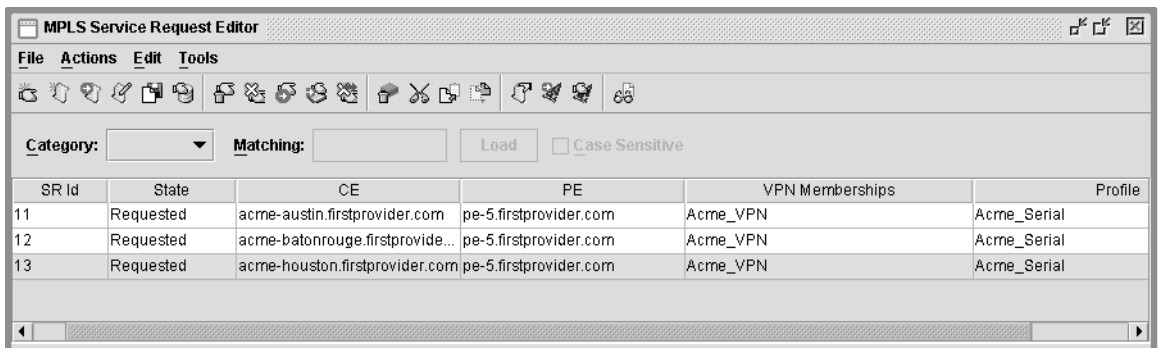

Step 2 Specify whether you want to deploy the service requests immediately or schedule their deployment:

- To deploy the new service requests immediately, choose **Actions > Deploy Service Requests > Deploy Now**.
- To schedule their deployment for another time, choose **Actions > Deploy Service Requests > Schedule Deploy**.

  When you choose **Schedule Deploy**, the Schedule VPN Service dialog box appears (see Figure 6-16).

Figure 6-16 Scheduling VPN Services

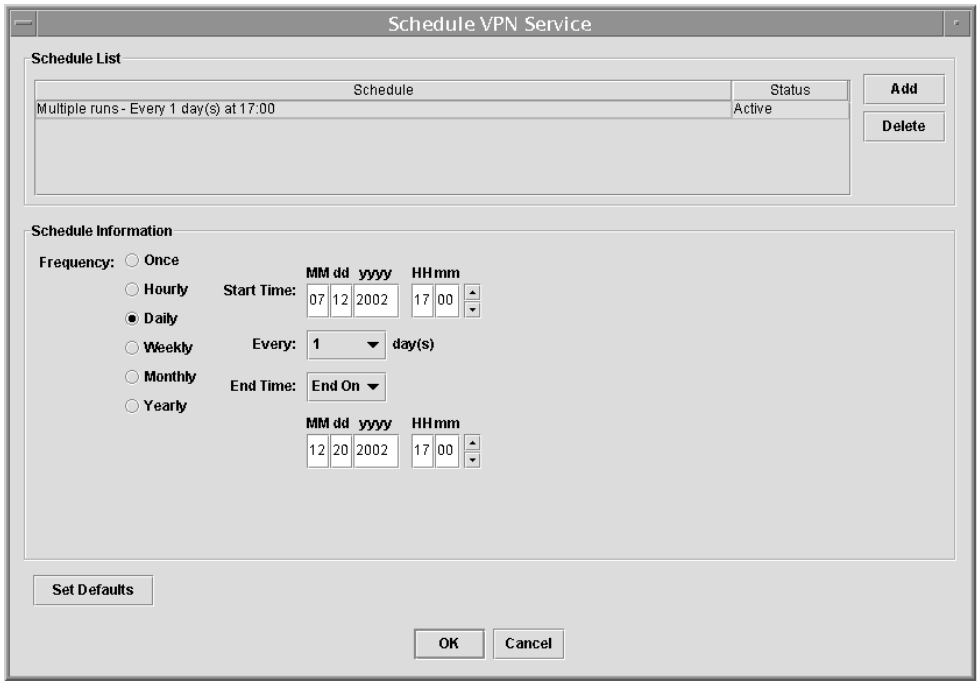

Step 3 Complete the fields in the Schedule VPN Service dialog box to schedule the service requests as needed.

a. *Frequency*: From the Frequency list, choose the desired frequency: **Once**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**.

   If you choose any other option except Once, new fields appear in the Schedule VPN Service dialog box.

b. *Start Time*: In the *Start Time* fields, specify the date and time to start the service.

c. *Every*: From the *Every* drop-down list, specify how often the service should run.

d. *End Time*: In the *End Time* drop-down list, choose either:

   - **No End** for a service with no termination time and date.
   - **End On** for a service that should end at a specific time and date.

e. If you select **End On**, specify the date and time to end the service.

Step 4 When you have scheduled the service requests to your satisfaction, click **Add**.

The service requests are added to the schedule list (displayed in the Schedule List panel).

You can delete a service request from the schedule list by selecting the pertinent line in the schedule list and clicking **Delete**. Then click **Yes** when prompted to confirm the deletion.

Step 5 When satisfied with the scheduling, click **OK**.

You receive the following confirmation prompt:

```
x service requests deployed.
```

Step 6 Click **OK**.

Overriding the Default VRF Name and Route Distinguisher Values

When you enable the VRF-RD Override property in the *csm.properties* file, the VPN Memberships dialog box presents options that allow you to override the default VRF name and Route Distinguisher (RD) values.

Caution Changing the default values for the VRF name and the Route Distinguisher (RD) value can alter or disable service requests that are currently running if not done correctly. Please make these changes with caution and only when absolutely necessary.

To override the default VRF name or the default RD values, follow these steps:

Step 1 On the VPN Solutions Center workstation, log in as **root** (**su**).

Step 2 Go to the */<installation_directory>/vpnadm/vpn/etc* directory.

Step 3 Open the *csm.properties* file with a text editor.

Step 4 Find the following section in the *csm.properties* file:

```
# Override VRF names and RD values.
# WARNING: This is an advanced feature. Overriding VRF names and RD values
# can potentially modify the intent of other service requests.
netsys.srvc.VRFRDOverride.unix=false
```

Step 5 Change the ***false*** value to ***true***, then save your changes and exit the file.

Step 6 If the Watch Dog is running, be sure to stop the Watch Dog, then start it again to enable this change.

Step 7 In the VPN Console, proceed through the Add Service to CE user interface as described in the previous sections.

When the VPN Memberships dialog box appears, it now displays fields for the VRF name and the RD value (see Figure 6-17).

Figure 6-17 VRF Name and RD Override Options

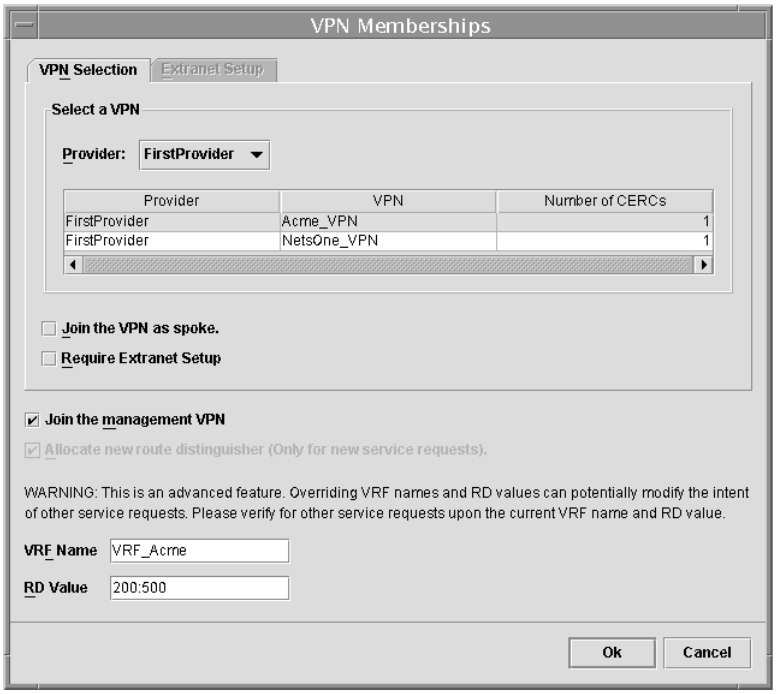

Step 8 *VRF Name*: To override the default VRF name, enter the new VRF name.

The maximum number of characters for the VRF name is 32.

Step 9 *RD Value*: To override the default Route Distinguisher value, enter the new RD value.

The RD number has the nomenclature **N1:N2**, where N1 is the BGP autonomous system number, and N2 is the RD number.

Step 10 When finished entering the necessary information, click **OK**.

Generating a Service Request Audit

When you initiate VPN Solutions Center service request audit, the audit tests all the configurations in the service request that were downloaded to the device(s). The service request is promoted to the Deployed state only when all the configurations pass the audit.

When a service request moves beyond the control of the Provisioning system, the Auditor for VPN Solutions Center takes control. The Auditor is a mechanism that monitors and reports the current state of a VPN service request over its lifetime. The lifetime of a VPN service request spans from the Requested state to the Closed state (see the "Service Request Summary" section on page 6-1). The Auditor also provides the reasons why the service request is in its current state. The Auditor saves the state transition (if any) into the VPN Inventory Repository.

To audit the service requests for an MPLS VPN, follow these steps:

Step 1 From the VPN Console hierarchy pane, expand the VPN Inventory hierarchy to display the **VPNs** folder (see Figure 6-18).

Figure 6-18 Selecting the VPN for a Service Request Audit

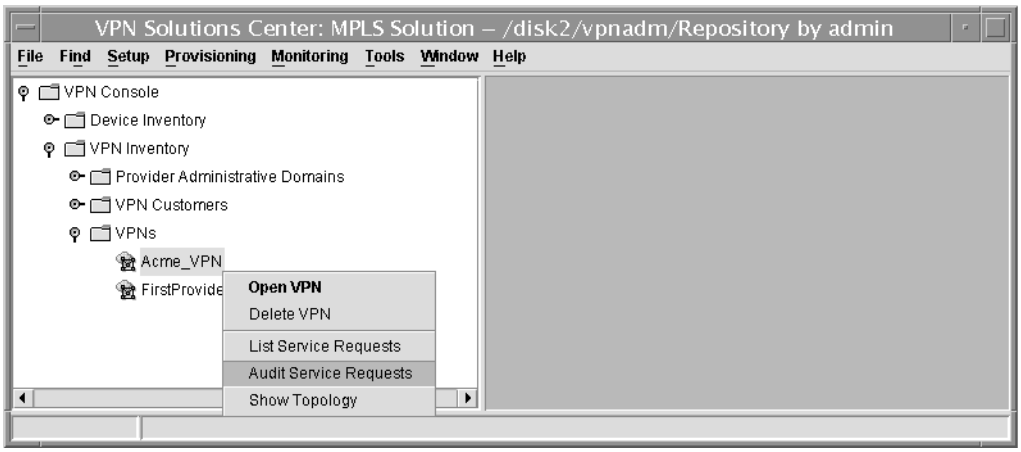

Step 2 Expand the VPNs folder to see the list of VPNs.

Step 3 Select the name of the VPN you want to audit, then **right-click**.

Step 4 From the menu, choose **Audit Service Requests**.

The Audit Customer dialog box shown in Figure 6-19 appears.

Figure 6-19 Choosing Which Service Requests to Audit

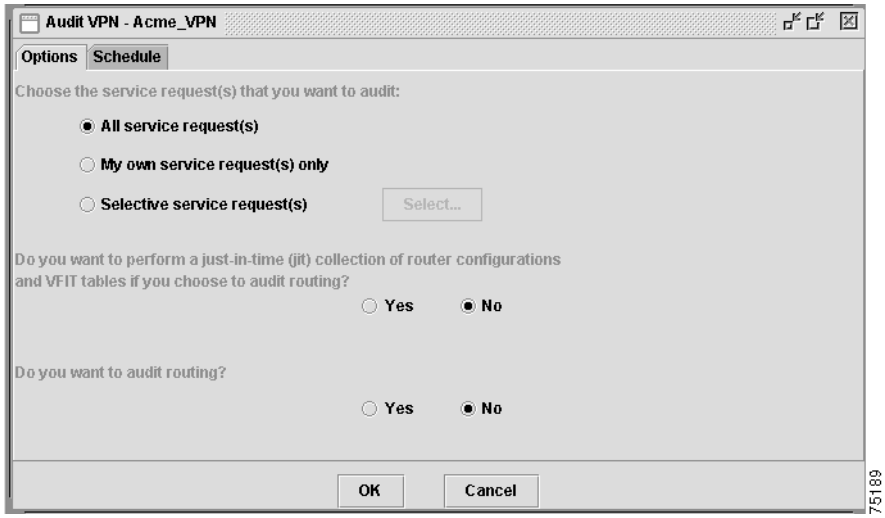

Step 5 Choose which service requests you want to audit:

- All service requests for the selected VPN
- Your own service requests only
- Selected service requests

Step 6 If you want to audit only selected service requests, select the **Selective service requests** option, then click **Select**.

A table showing the service requests from which you can select is displayed (see Figure 6-20).

Figure 6-20 List of Service Requests to Choose From

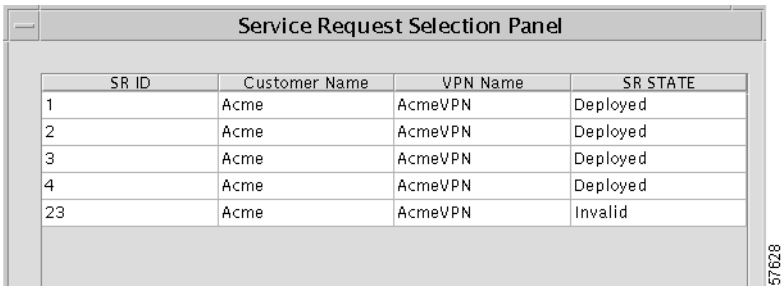

a. Select one or more service requests from the list.

b. Click **OK**.

You return to the Audit Options dialog box shown in Figure 6-19.

Step 7 *Just-in-time (jit) configuration collection*: If you want to collect the latest router configuration files from the routers effected by the selected service request before VPN Solutions Center runs the audit, select the **Yes** option.

If you do not want to collect the latest router configuration files before the audit begins, accept **No**, the default *Just-in-time (jit) configuration collection* option.

Step 8 When the audit options are set to your satisfaction, click the **Schedule** tab.

The Audit Schedule dialog box appears (see Figure 6-21).

Figure 6-21 Scheduling the Audit

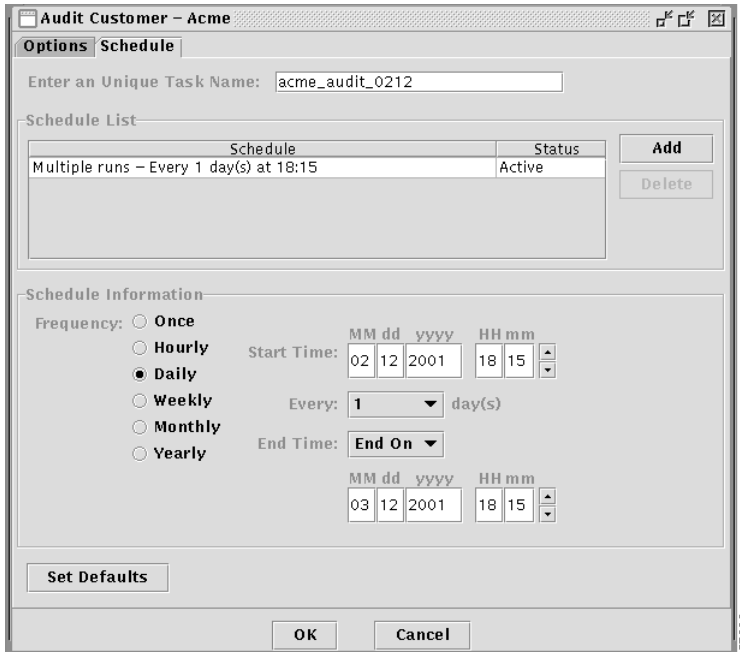

Step 9 Complete the fields in the dialog box to schedule the audit as needed.

a. *Task name*: Enter a unique audit task name.

b. *Frequency*: From the *Frequency* list, choose the desired frequency: **Once**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**.

c. *Start Time*: Set the *Start Time* to **Now** or **Later**.

If you chose any other option except **Once**, new fields appear in the Schedule dialog box.

d. *Later*: If you chose **Later**, in the *Start Time* fields, specify the date and time to start the service.

e. *Every*: from the *Every* drop-down list, specify how often the service should run.

f. *End Time*: In the *End Time* drop-down list, choose either:

- **No End** for a service with no termination time and date.
- **End On** for a service that should end at a specific time and date.

g. If you selected **End On**, specify the date and time to end the service.

Step 10 When you have scheduled the audit to your satisfaction, click **Add**.

The audit is added to the Schedule List, displayed in the upper area of the dialog box (as shown in Figure 6-21).

Step 11 Click **OK**.

You return to the VPN Console.

Viewing Audit Reports

Before you view audit reports, you must first generate an audit as described in the previous section. To view audit reports, follow these steps:

Step 1 From the VPN Console menu, choose **Provisioning > List All Service Requests**.

The All VPN Service Requests Report appears (see Figure 6-22).

Figure 6-22 All VPN Service Requests Report

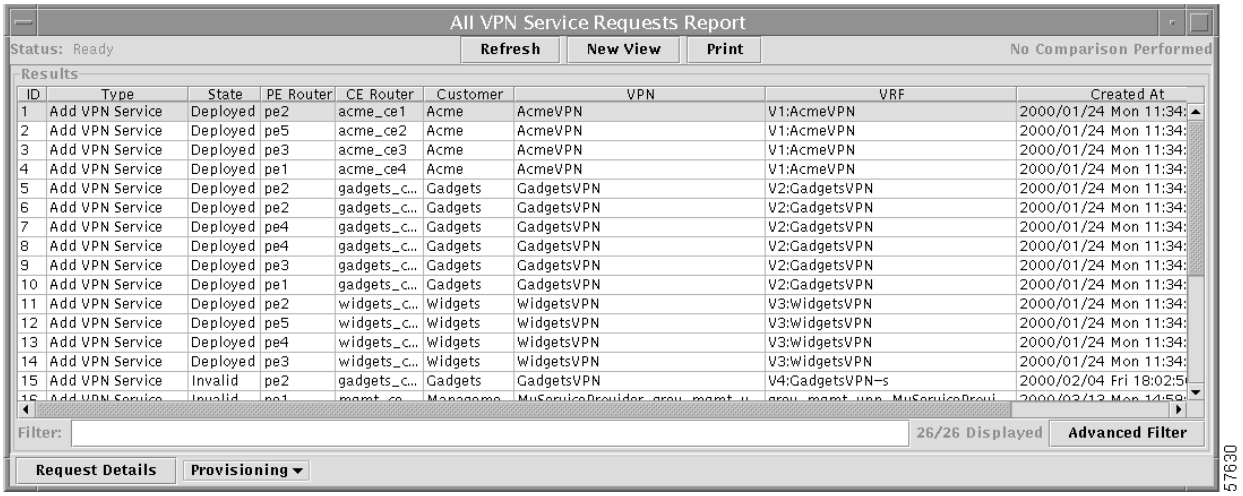

Step 2 Select (highlight) the desired service request.

Step 3 Click **Request Details**.

The Service Request Details Report appears (see Figure 6-25 on page 6-24).

Step 4 From the Service Request Details Report, click **Audit Detail**.

The Service Request Audit Report appears (see Figure 6-23).

Figure 6-23 Audit Details Report

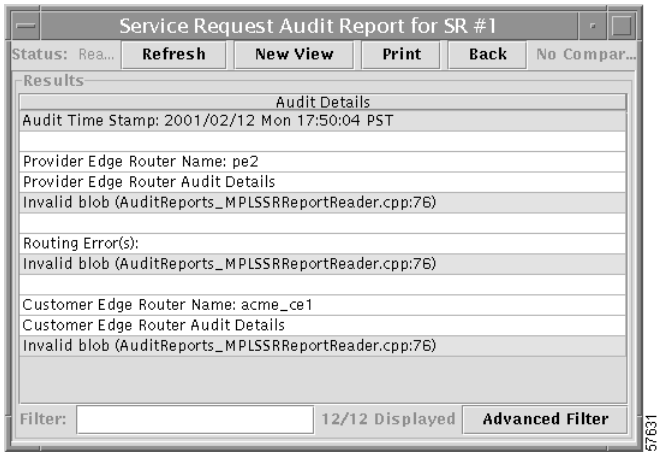

Those items that the audit discovered problems with are highlighted in yellow.

Checking Service Request Deployment Details

Once you have created and queued a service request, you can discover the details about its deployment. You can view the configlet generated for the service request. If the service request failed, you can discover why it failed by using the Service Request Audit report.

Step 1 To check service request details, choose **Provisioning>List All Service Requests**.

The All VPN Service Requests Report appears (see Figure 6-24).

Figure 6-24 All VPN Service Requests Report

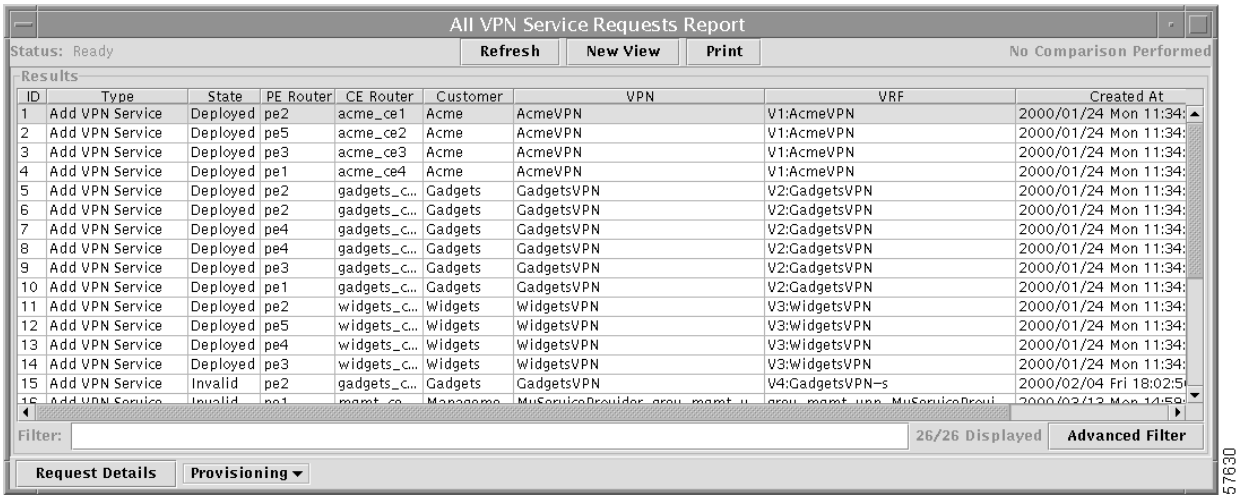

This report provides the following information:

- Service request ID number
- Type of request
- Current state

  If the current state is either Deployed or Functional, the service request is deployed.
- Names of the PE and CE router the service is for
- Customer name
- VPN name
- VRF name
- Time and date the service request was created
- Time and date when the service was last changed

Step 2 Select the service request you want detailed information on.

Step 3 Click **Request Details**.

The Service Request Details Report appears (see Figure 6-25).

Figure 6-25 Service Request Details Report

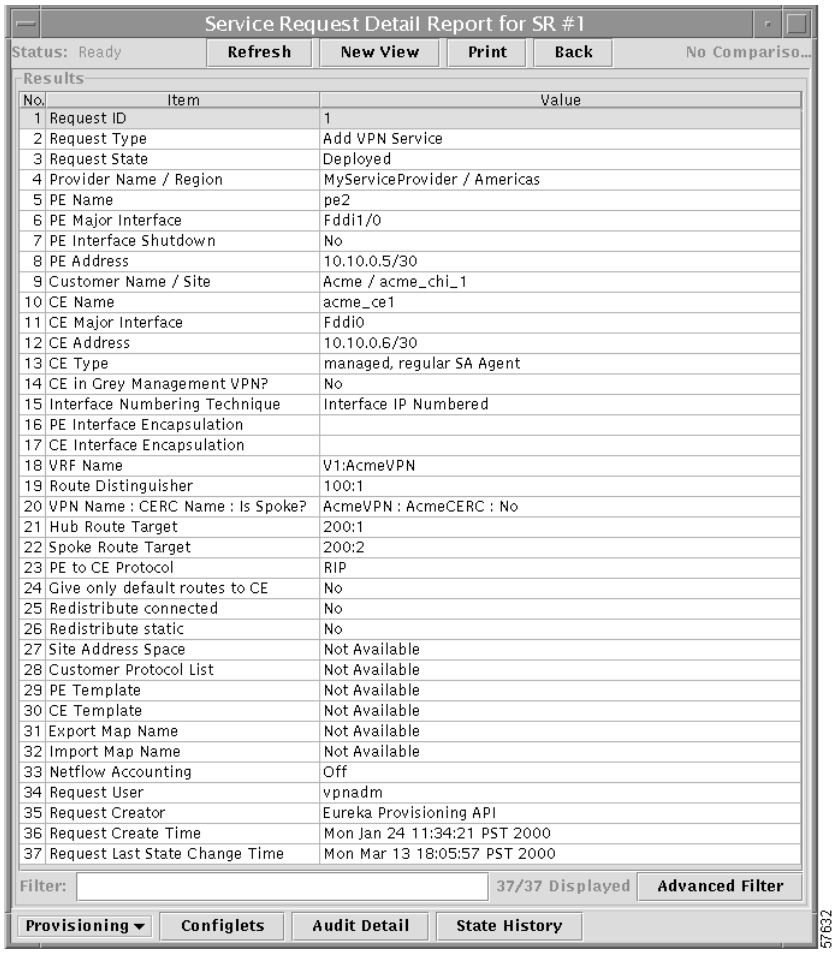

No.	Item	Value
1	Request ID	1
2	Request Type	Add VPN Service
3	Request State	Deployed
4	Provider Name / Region	MyServiceProvider / Americas
5	PE Name	pe2
6	PE Major Interface	Fddi1/0
7	PE Interface Shutdown	No
8	PE Address	10.10.0.5/30
9	Customer Name / Site	Acme / acme_chi_1
10	CE Name	acme_ce1
11	CE Major Interface	Fddi0
12	CE Address	10.10.0.6/30
13	CE Type	managed, regular SA Agent
14	CE in Grey Management VPN?	No
15	Interface Numbering Technique	Interface IP Numbered
16	PE Interface Encapsulation	
17	CE Interface Encapsulation	
18	VRF Name	V1:AcmeVPN
19	Route Distinguisher	100:1
20	VPN Name : CERC Name : Is Spoke?	AcmeVPN : AcmeCERC : No
21	Hub Route Target	200:1
22	Spoke Route Target	200:2
23	PE to CE Protocol	RIP
24	Give only default routes to CE	No
25	Redistribute connected	No
26	Redistribute static	No
27	Site Address Space	Not Available
28	Customer Protocol List	Not Available
29	PE Template	Not Available
30	CE Template	Not Available
31	Export Map Name	Not Available
32	Import Map Name	Not Available
33	Netflow Accounting	Off
34	Request User	vpnadm
35	Request Creator	Eureka Provisioning API
36	Request Create Time	Mon Jan 24 11:34:21 PST 2000
37	Request Last State Change Time	Mon Mar 13 18:05:57 PST 2000

Step 4 To view the configlets generated for the selected service request, click **Configlets**.

The report shown in Figure 6-26 appears.

Figure 6-26 Service Request Configlets Report

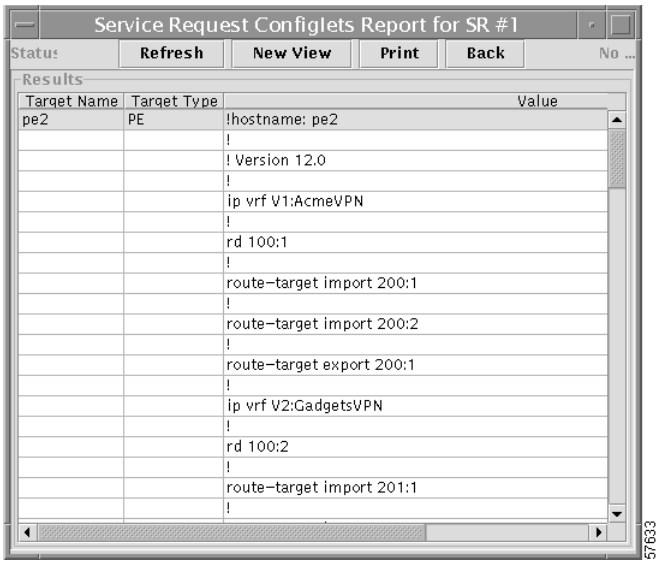

To return to the Service Request Detail Report, click **Back**.

Modifying an Existing Service Request

A service request is an instance of service contract between a CE and a PE. You can modify this service by creating a new service request. When you do so, VPN Solutions Center creates a new service request with a new ID. (The service request ID is displayed in the *SR ID* column in the MPLS Service Request Editor.) The new service request subsumes the earlier one and becomes the current service request.

When you modify a service request, you can modify the settings for the PE-CE link, except for the CE and the PE themselves. This procedure takes through the same user interface as described in the "Adding a Service for a PE-CE Link" section on page 6-6, except that the settings are based on the service request's current values.

To modify a service, follow these steps:

Step 1 Choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor appears (see Figure 6-27).

Figure 6-27 Modifying a Service Request

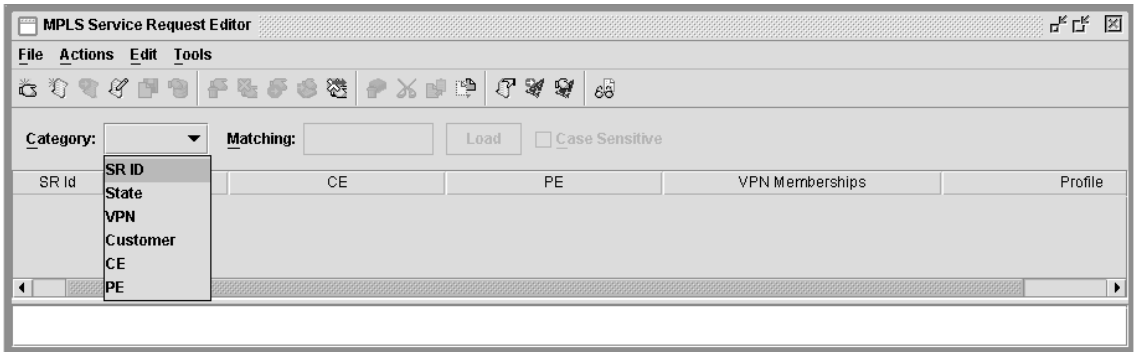

Step 2 From the list of service requests, select the service request you need to modify.

You can search for service requests by the following categories:

- Service request ID
- Service request state
- VPN name
- Customer name
- CE name
- PE name

To locate a particular service request:

a. *Category*: Select the category by which you want VPNSC to search.

b. *Matching*: Enter the string to search for.

c. Click **Load**.

The settings for the selected service request are displayed in the Service Request Editor. You can modify the VPN Memberships or the profile associated with the service.

Modifying the VPN Membership Information for the Service

To modify the VPN membership information for the selected service request:

a. Select the *VPN Memberships* cell.

b. Choose **Actions > Set VPN Memberships**.

The VPN Memberships dialog box appears.

c. Modify the VPN membership information as necessary (for details, see the "Specifying the VPN Membership Information" section on page 6-10).

d. Click **OK**.

You receive the following message:

```
Changed the VPN Memberships for 1 service request.
```

e. Click **OK**.

Modifying the Profile for an Existing Service

To modify the profile for an existing service request:

a. Select the *Profile* cell for the specific service request, then **double-click**.

   The Profile Editor displays the profile associated with this service request.

b. Edit the MPLS attributes you need to change. For details, see the "Specifying the MPLS Attributes for a Service Request Profile" section on page 5-20.

c. When satisfied with the changes, click **Apply**.

   You receive the message:

   ```
   Created a schedule to deploy n service requests.
   ```

   The edited service is moved to the *Closed* state and new service request is created and placed in the **Requested** state.

Changing the Profile Associated with a Service Request

To change the profile associated with a service request:

a. Select the *Profile* cell.

b. Choose **Actions > Set Profile**.

   The Open Profile dialog box appears.

c. Select the appropriate service request profile for this service.

d. Click **Open**.

   You receive the message:

   ```
   Changed the profile for 1 service request.
   ```

e. Click **OK**.

Step 3 Choose **File > Commit to Repository**.

This message is displayed:

```
Committed 1 service request to the Repository.
```

Step 4 Choose **Actions > Deploy Service Requests**, then choose **Deploy Now** or **Schedule Deploy**.

You have now queued the modified service request. It is entered into the Repository and placed in the "Requested" state.

Decommissioning a VPN Service

Decommissioning a service request from VPN Solutions Center is a four-task process:

1. *Create a service request to initiate the decommissioning of the selected VPN service.*

   When you decommission a VPN service, VPN Solutions Center replaces the old service request with a new one whose purpose is to remove the pertinent commands from the PE and CE router configuration files.

2. *Deploy the remove VPN service request.*

   The new remove VPN service request will be in Requested state, and you should deploy it as you do any other service request.

Deploying a "Remove VPN Service" request deletes individual commands from the PE and CE configuration files, which were put there by the original provisioning request, and are not in use by any other service or feature in the router configuration.

To ensure that the service removal is safe requires that not all commands that were provisioned are removed. In cases where VPN Solutions Center cannot know whether a provisioned command is being used for some other purpose, the command is not removed. Examples of router commands not removed for a "Remove VPN Service" request include routing protocols created during service provisioning, such as BGP or RIP. These are not removed from the router's configuration file, although some of their subcommands are removed when they support only the original service request.

3. *Audit the remove VPN service request.*

   The auditing process uploads the revised configuration file into VPN Solutions Center and checks to confirm that the appropriate commands have been removed from the file. If the audit is successful, the remove VPN service is moved to the Closed state. Though this step is not required, it is highly recommended.

4. *Purge the remove VPN service request.*

   To delete the remove VPN service request from the Repository, you must purge the service request.

Removing a VPN Service Request

To remove a specified VPN service:

Step 1 Choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor appears (see Figure 6-28).

Figure 6-28 Removing a VPN Service

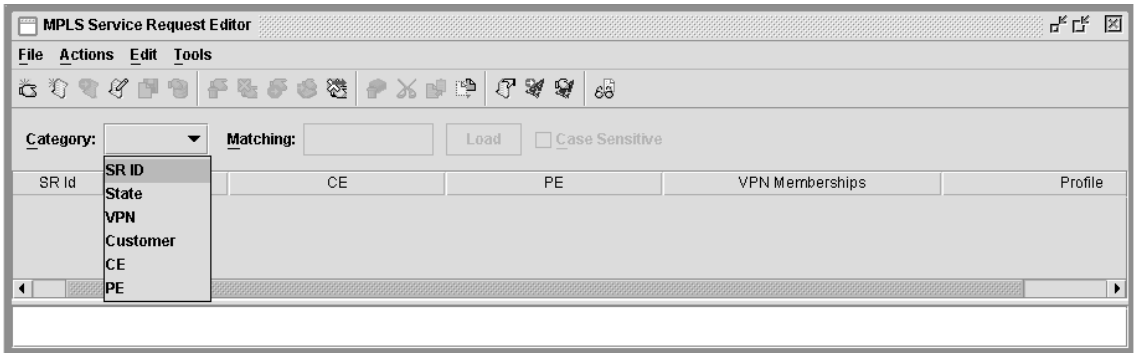

Step 2 From the list of service requests, select the service requests you want to remove.

You can search for service requests by the following categories:

- Service request ID
- Service request state
- VPN name
- Customer name
- CE name

- PE name

To locate a particular service request:

a. *Category*: Select the category by which you want VPNSC to search.

b. *Matching*: Enter the string to search for.

c. Click **Load**.

The specified service requests are displayed in the Service Request Editor.

Step 3 Choose **Actions > Remove Service Request(s)**.

You receive this warning message:

This will submit a new service request to remove the VPN service between the PE and CE. New configlets will be generated with the appropriate "no" commands to remove the VPN service. Service Request n to Add VPN Service will no longer be active. Do you want to continue?

Step 4 Click **Yes** to proceed, or **No** to cancel the Remove VPN Service operation.

If you click **Yes**, you receive the following message:

A new service request has been submitted to remove the VPN service specified in service request n.

Step 5 Click **OK**.

The new remove VPN service request is placed in the Requested state.

Step 6 Deploy the service request as described in the "Deploying Service Requests" section on page 6-15.

Auditing the Remove VPN Service Request to Close It

The auditing process uploads the revised configuration file into VPN Solutions Center and checks to confirm that the appropriate commands have been removed from the file. If the audit is successful, VPNSC sets the remove VPN service to the Closed state.

Note Though this step is not required, it is highly recommended. If you choose not to audit the remove VPN service, you can close the service manually as described in the "Closing Service Requests Manually" section on page 6-31.

To audit the remove VPN service, follow these steps:

Step 1 From the VPN Console hierarchy pane, expand the VPN Console hierarchy to display the **VPN Customers** folder.

Step 2 Expand the VPN Customers folder to see the list of VPNs.

Step 3 Select the name of the VPN you want to audit, then **right-click**.

Step 4 From the menu, choose **Audit Service Requests**.

The Audit Customer dialog box appears.

Step 5 From the Audit Customer dialog box, choose the **Selective service requests** option, then click **Select**.

A table showing the service requests from which you can select is displayed.

Step 6 Select the remove VPN service request from the list, then click **OK**.

You return to the Audit Options dialog box.

Step 7 If you want to collect the latest router configuration files from the routers effected by the selected service request before VPN Solutions Center runs the audit, accept the **Yes** option.

If you do not want to collect the latest router configuration files before the audit begins, click **No**.

Step 8 When the audit options are set to your satisfaction, click the **Schedule** tab.

The Audit Schedule dialog box appears.

Step 9 Complete the fields in the dialog box to schedule the audit as needed.

a. *Task Name*: Enter a unique task name for the audit operation.

b. *Frequency*: From the *Frequency* list, choose the desired frequency: **Once**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**.

c. *Start Time*: Set the *Start Time* to **Now** or **Later**.

d. *Later*: If you choose **Later**, specify the date and time to start and end the audit.

e. If you choose anything other than **Once**, specify how often the audit should run from the **Every** drop-down list.

Step 10 When you have scheduled the audit to your satisfaction, click **Add**.

The audit is added to the Schedule List, displayed in the upper area of the dialog box.

Step 11 Click **OK**.

You return to the VPN Console.

Purging a Closed Service from the Repository

The final task in removing a VPN service is to purge the service from the Repository. VPN Solutions Center software does not automatically remove closed service requests from the Repository (in case you need them for your records). But keeping closed service requests can be a waste of disk space, therefore, the VPN Solutions Center software provides a way to purge obsolete request data from the Repository.

To purge closed service requests from the Repository, follow these steps:

Step 1 If you have not already done so, close the service request(s) you want to remove as described in the previous section, "Auditing the Remove VPN Service Request to Close It."

If the audit operation does not close the target service request(s), you can close a service request manually as described in the "Enabling Manual Closure of Service Requests" section on page 6-31.

Step 2 From the VPN Console, choose **Provisioning > Purge Closed Requests from Database**.

You receive the following Purge Confirmation message (see Figure 6-29):

Figure 6-29 The Purged Closed Service Requests Message Dialog Box

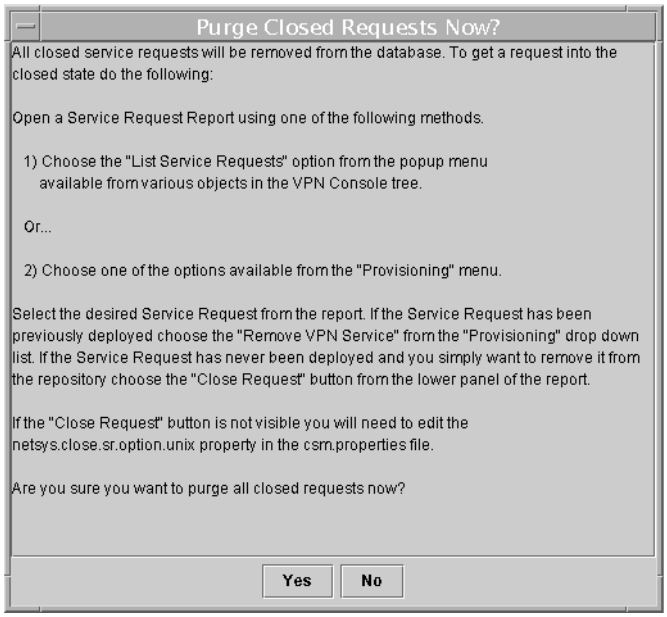

Step 3 If you wish to proceed with the service request removal operation, click **Yes**.

Closing Service Requests Manually

When you manually close a service request, VPN Solutions Center changes the state of the service to *Closed* in the Repository. VPN Solutions Center does not make any modifications to the router's configuration file when you close a service request. You cannot remove a service request from the Repository until it is closed.

Enabling Manual Closure of Service Requests

Before you can manually close a service request, you must enable a certain property in the *csm.properties* file. Changing this value in the *csm.properties* file provides a new option in the VPN Console that allows you to remove service requests in any state.

Step 1 If VPN Solutions Center is running, shut it down.

Step 2 On the VPN Solutions Center workstation, log in as the *vpnadm* administrative user.

Step 3 Go to the */<installation_directory>/vpnadm/vpn/etc* directory.

Step 4 Open the *csm.properties* file with a text editor.

Step 5 Find the following property in the *csm.properties* file:

```
netsys.close.sr.option.unix = Off
```

Step 6 Change the `Off` value to **On** as follows:

```
netsys.close.sr.option.unix = On
```

Step 7 Save your changes and exit from the file.

Step 8 If the Watch Dog is running, be sure to stop the Watch Dog, then start it again to enable this change.

Step 9 Restart VPN Solutions Center.

Closing a Service Request

To manually close a service request, follow these steps:

Step 1 From the VPN Console menu bar, choose **Provisioning > List All Service Requests**.

The All VPN Service Requests Report appears.

A new option for closing service requests—**Close Request**—is displayed on the menu bar at the bottom of the All VPN Service Requests Report (see Figure 6-30).

Figure 6-30 Close Service Request Option Enabled

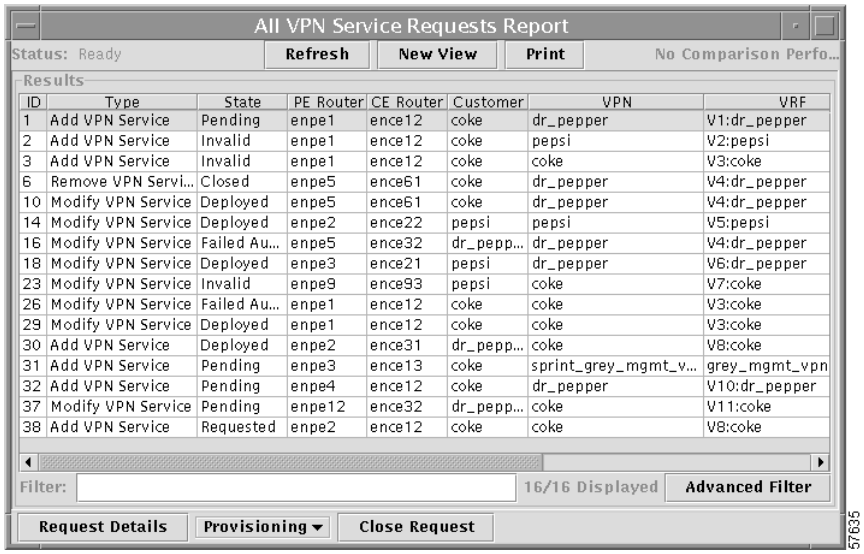

Step 2 Select the service request you want to close.

Step 3 From the menu bar at the bottom of the All VPN Service Requests Report, click **Close Request**.

You receive the following confirmation prompt:

Service request n is about to be forced into the Closed state. Do you want to continue?

Step 4 To close the selected service requests, click **Yes**.

To cancel the close operation, click **No**.

VPN Solutions Center changes the state of the selected services to Closed in the Repository.

Performing a Customized Service Request Deployment

The procedure to perform a customized service request deployment deploys the service request immediately. This customized deployment does not perform an audit, nor does it allow you to schedule the audit.

Step 1 From the VPN Console, choose **Provisioning > List All Service Requests**.

The All VPN Service Requests Report appears (see Figure 6-31).

Figure 6-31 All VPN Service Requests Report

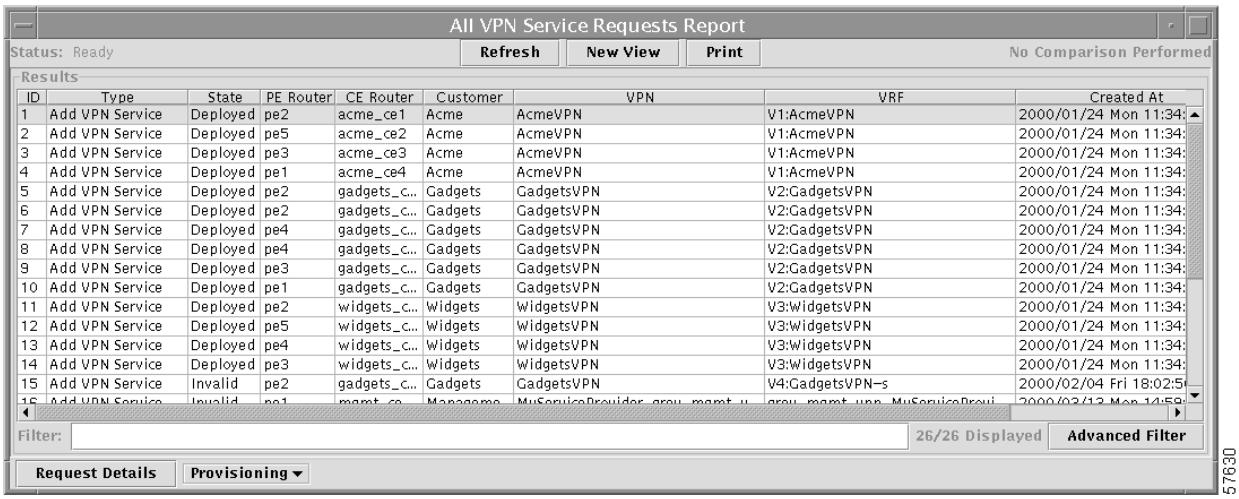

Step 2 Select the service request you want to deploy.

Step 3 From the Provisioning menu at the bottom of the window, click **Provisioning**.

The Service Request Provisioning drop-down menu appears (see Figure 6-32).

Figure 6-32 Service Request Provisioning Menu

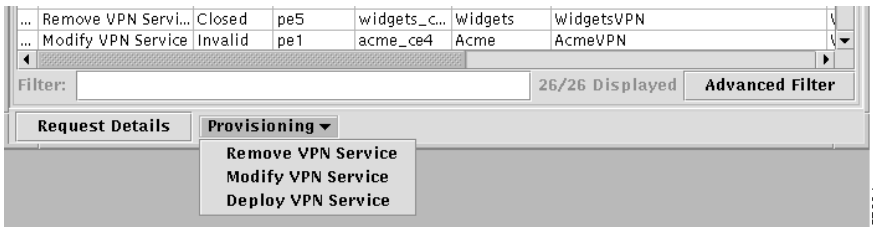

Step 4 From the drop-down menu, choose **Deploy VPN Service**.

The following message is displayed:

This will deploy the selected VPN service request now. Do you want to continue?

Step 5 Click **Yes**.

The selected service request is Deployed and placed in the Pending state.

Using the Task Manager

VPN Solutions Center provides a Task Manager that allows you to view pertinent information about both current and expired provisioning tasks, as well as create and schedule tasks, delete specified tasks, and delete the expired tasks.

To bring up the Task Manager, choose **Tools > Tasks** from the VPN Console menu.

The Task Manager window appears (see Figure 6-33).

Figure 6-33 The Task Manager Window

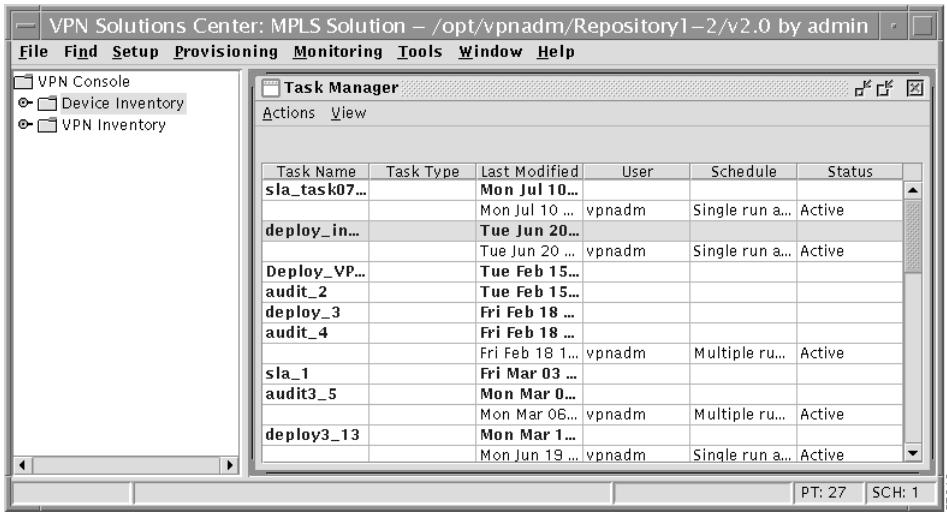

The Task Manager window provides information on each task by name, including the task type, the date when the task was last modified, the VPN Solutions Center username, schedule summary information, and its current status—expired or active.

From the Actions menu (shown in Figure 6-34), you can execute all the necessary task-related functions:

Figure 6-34 The Task Manager Actions Menu

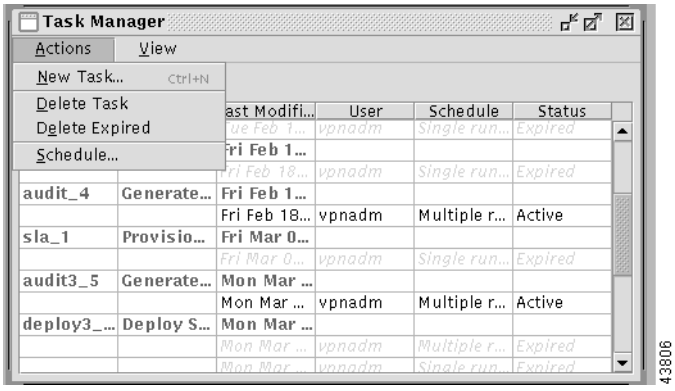

Creating a New Task

To create a new task:

Step 1 From the VPN Console, choose **Tools > Tasks**.

The Task Manager window appears.

Step 2 To create a new task, choose **Actions > New Task** from the Task Manager menu.

The Task Chooser appears (see Figure 6-35).

Figure 6-35 The Task Chooser

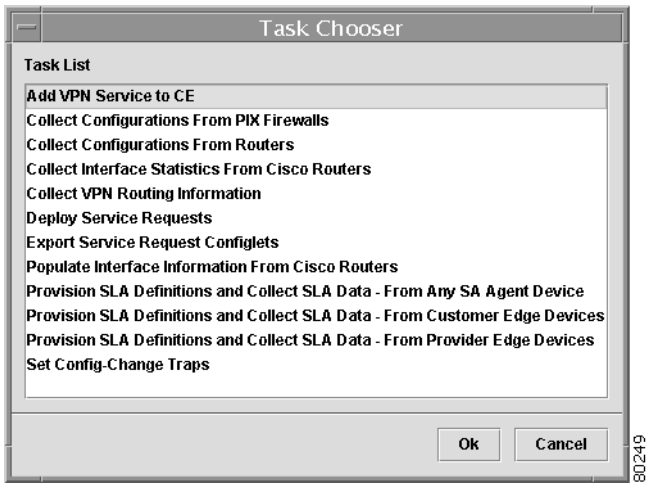

Step 3 From the Task Chooser's Task List, select the task you want to execute and press **OK**.

Step 4 Complete the task user interface as required.

Deleting a Task

When you delete a task through the Task Manager, you delete both the persistent and the scheduled tasks. VPN Solutions Center removes the task from the Task Repository and updates the task logs (see also the "Deleting Task Logs" section on page 6-42).

To delete one or more tasks, do the following:

Step 1 From the VPN Console, choose **Tools > Tasks**. The Task Manager window appears.

Step 2 In the Task Manager window, select one or more tasks to delete.

Step 3 Choose **Actions > Delete Task**.

You are asked to confirm the deletion request:

You have selected to delete n task(s) from the Repository. Do you want to continue?

Step 4 To delete the selected tasks, click **Continue**.

You can also cancel the operation at this point by clicking **Cancel**.

VPN Solutions Center deletes the selected tasks and redisplays the current list of tasks in the Task Manager window.

Deleting Expired Tasks

To delete tasks that have expired, follow these steps:

Step 1 From the VPN Console, choose **Tools > Tasks**. The Task Manager window appears.

Step 2 From the Task Manager, choose **Actions > Delete Expired**.

You are asked to confirm the deletion request:

You have selected to delete the expired tasks from the Repository. Do you want to continue?

Step 3 To delete the expired tasks, click **Continue**.

You can also cancel the operation at this point by clicking **Cancel**.

VPN Solutions Center deletes the expired tasks and redisplays the current list of tasks in the Task Manager window.

Scheduling a Task

The VPN Solutions Center Task Manager allows you to schedule a selected task. To schedule a task, follow these steps:

Step 1 From the VPN Console, choose **Tools > Tasks**. The Task Manager window appears.

Step 2 From the Task Manager window, select the task you want to schedule.

Step 3 From the Task Manager, choose **Actions > Schedule**.

The Scheduler dialog box appears (see Figure 6-36).

Figure 6-36 The Scheduler Dialog Box

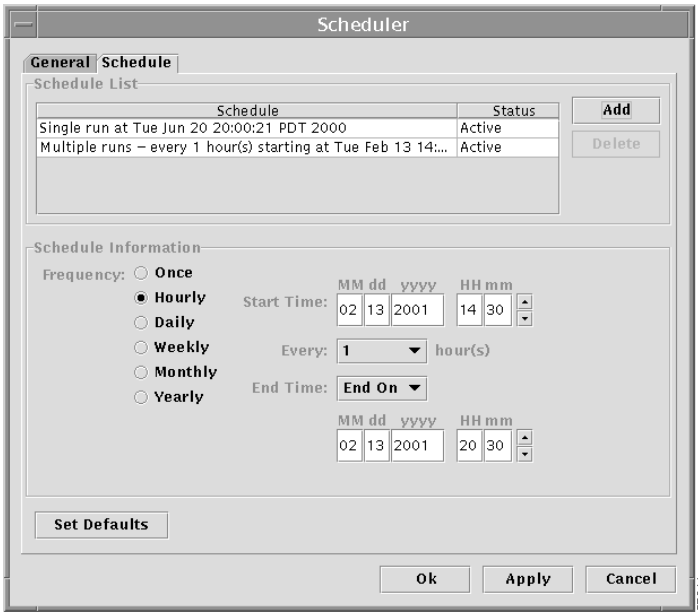

Step 4 Complete the fields in the dialog box to schedule the task as needed.

a. *Frequency*: From the Frequency list, choose the desired frequency: **Once**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**.

b. *Start Time*: Set the Start Time: **Now** or **Later**.

c. If you choose **Later**, specify the date and time to start and end the service.

d. If you choose anything other than **Once**, specify how often the service should run from the **Every** drop-down list.

Step 5 When you have scheduled the task to your satisfaction, click **Add**.

The task is added to the Schedule List, displayed in the upper area of the dialog box (as shown in Figure 6-36).

Step 6 Click **Apply** to save your changes and remain in the Scheduler; click **OK** to save and exit the Scheduler.

When you click **OK**, you return to the VPN Console. The task is added to the VPN Solutions Center task queue; it will begin executing on the date and time specified.

Using the Task Logs

To access the VPN Solutions Center task logs, do the following:

Step 1 From the VPN Console, choose **Tools > Task Logs**.

VPN Solutions Center starts the browser and displays the Task Logs window.

If Netscape has not already started, the Netscape Password dialog box appears (see Figure 6-37).

Figure 6-37 Logging in to the VPN Solutions Center Browser

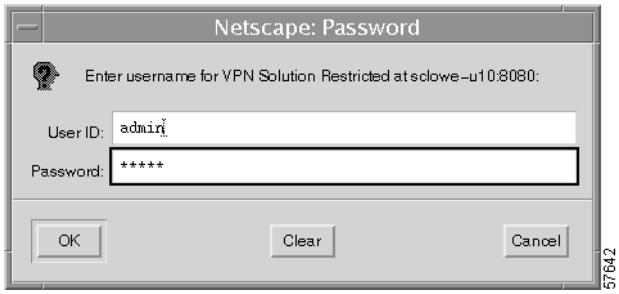

Step 2 In the Netscape Password dialog box, enter the VPN Solutions Center administrative username and password, then click **OK**.

The default administrative username and password is **admin** (for both).

The Task Logs page is displayed (see Figure 6-38).

The tasks are listed in order of the task start time; the task with the latest start time is listed at the top of list, and the task with the earliest start time is listed at the bottom of the list.

Notice the **Logs** column in Figure 6-38, which provides a **Log** link for every task listed.

Figure 6-38 Viewing the List of Tasks

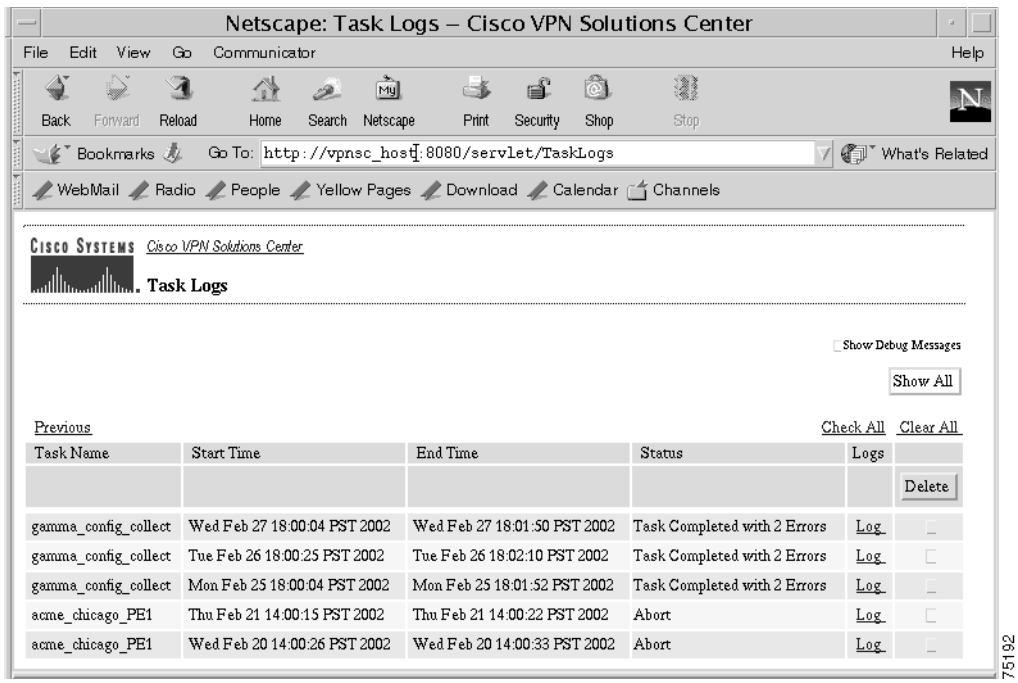

The task logs are displayed in sets of 5 logs per page in descending order, with the latest task on top of the list.

Jumping to the Next and Previous Pages of Task Logs

To jump to the previous page of task logs, click the **Previous** link (in the upper left corner of the task logs table).

When there are task log pages above and below the current position, the task logs page provides both **Next** and **Previous** links so you can navigate efficiently through multiple pages.

Viewing Debug Messages

Not all tasks include debug messages. If the specific task in question does have debug messages implemented, you can view in the task logs any debug messages that were generated by checking the **Show Debug Messages** checkbox.

Step 3 Scroll to the name of the task whose log you want to view, then click the corresponding **Log** link on that row.

The status report for the selected task appears in the lower left pane, as shown in Figure 6-39.

Figure 6-39 Displaying the Task Status

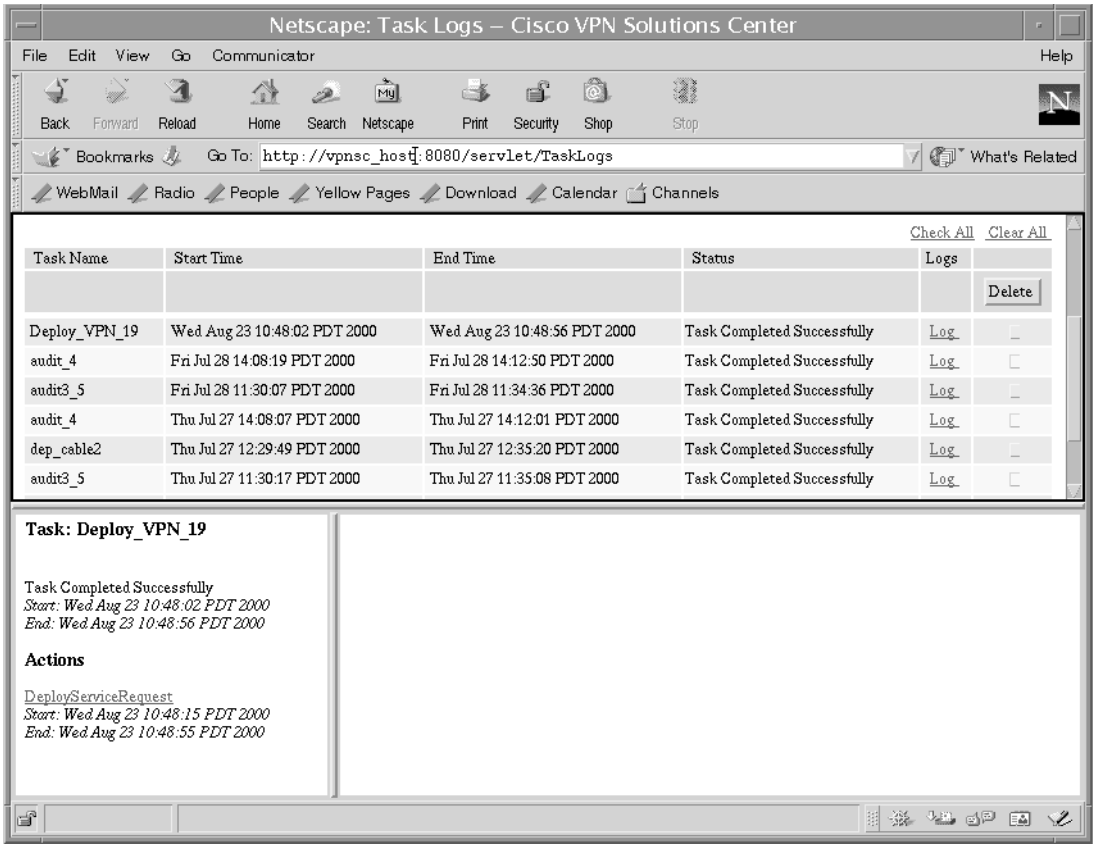

The status pane shows the following information for the selected task:

- Task status summary

  The possible task status summary states can be any one of the following:

  - **Task Completed Successfully**
  - **Running**, **Task Terminated**
  - **Task Completed with *n* Errors**
  - **Task Completed with *n* Warnings**
  - **Device Connection Error Occurred**

- Start time
- End time
- All the actions under the currently selected task
- Start time and end time for all the actions

Viewing a Task Action Log

Step 4 To see an Action log for the selected action, click the link displayed under the Actions heading.

The Action log appears in the lower right pane (see Figure 6-40).

For example, for the task shown, *Deploy_VPN_19*, there is only one action—"Deploy Service Request." (Some tasks have more than one action listed.)

To view the action report for the action "Deploy Service Request" for task "Deploy_VPN_19," click the **DeployServiceRequest** link.

Figure 6-40 The Task Action Report

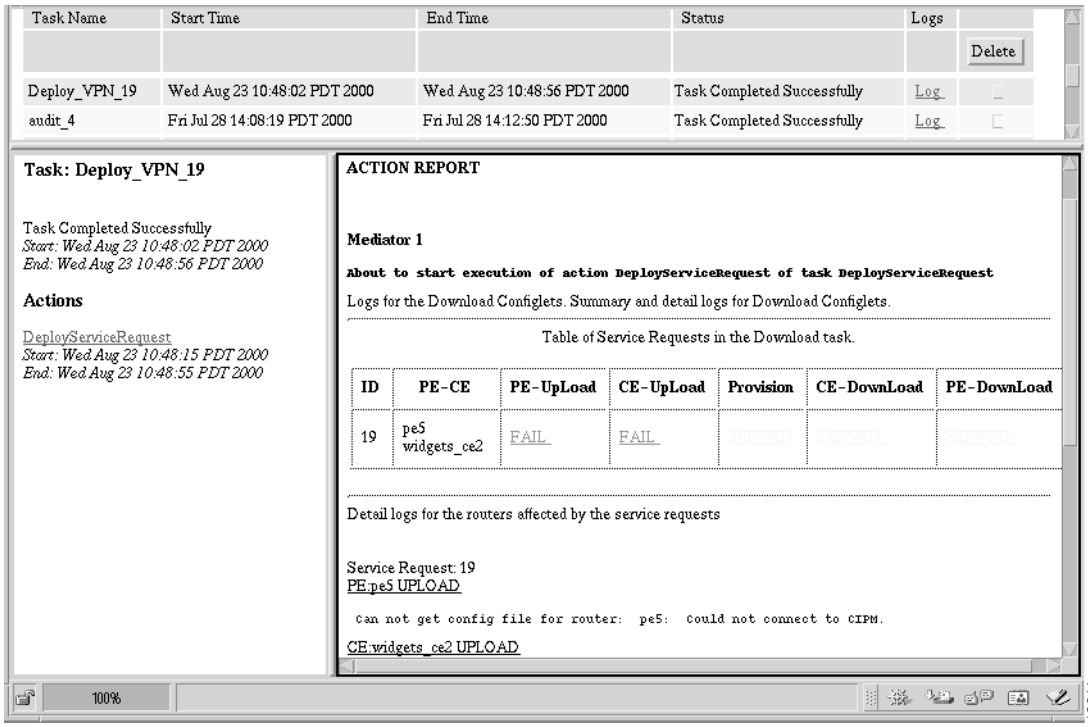

Viewing the Standard Output/Standard Error Log

Step 5 To view the standard output and error logs for the selected task, click the **Stdout/Stderr** link.

The Standard Output and Error log appears in the lower right pane (see Figure 6-41).

Figure 6-41 Standard Output and Error Logs

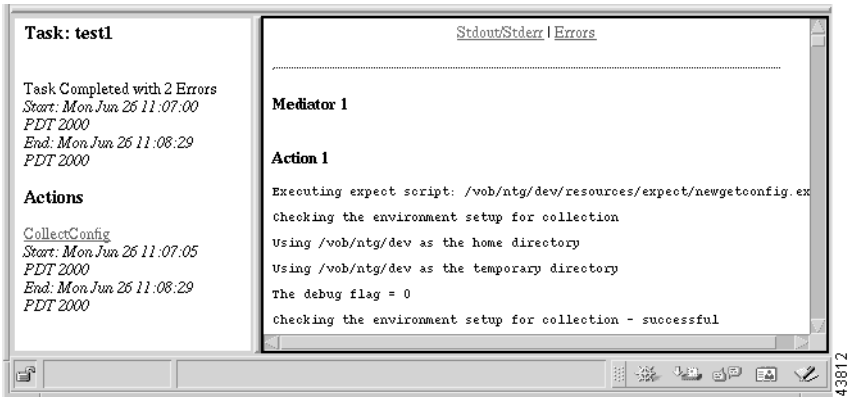

Viewing the Task Log Error Report

Step 6 To view the Error Report for the selected task log, click the **Errors** link.

The Task Log Error Report appears in the lower right pane (see Figure 6-42).

Figure 6-42 Task Log Error Report

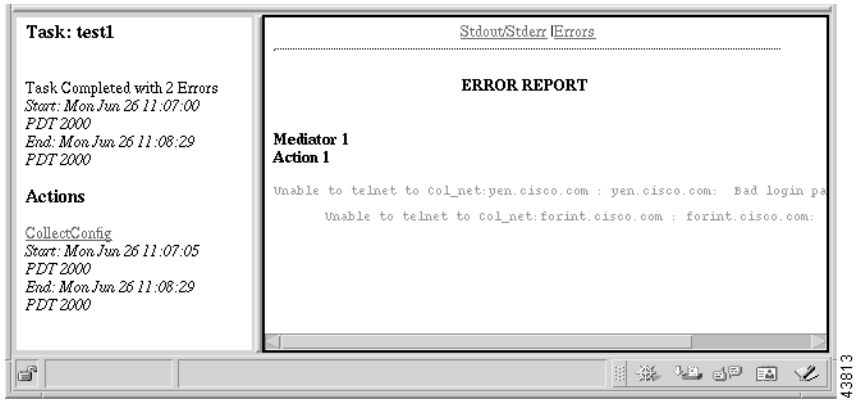

Deleting Task Logs

To save disk space, you can delete task logs when they become obsolete. When you delete a task from the task logs, you delete the run-time task and the logs associated with the task.

Tip We recommend that you delete no more than 10 task logs at a time.
When you have a large number of tasks and would like to delete the oldest logs, you may find it more convenient to click the **Show All** button. **Show All** displays all the tasks in one page, so if you have a large number of tasks in the Repository, the browser often takes a long time to display all of the tasks. But the advantage to this procedure is that you can delete the oldest logs (which would be the last set of tasks in the list), rather than having to click **Next** many times to reach the last page of logs.

To delete task logs, follow these steps:

Step 1 From the VPN Console, choose **Tools > Task Logs**.

VPN Solutions Center starts the browser and displays the Task Logs window. Notice the **Delete** column (the rightmost column in the Task Logs window) as shown in Figure 6-43.

The tasks are listed in order of the task start time; the task with the latest start time is listed at the top of list, and the task with the earliest start time is listed at the bottom of the list.

Figure 6-43 The Task Logs Window

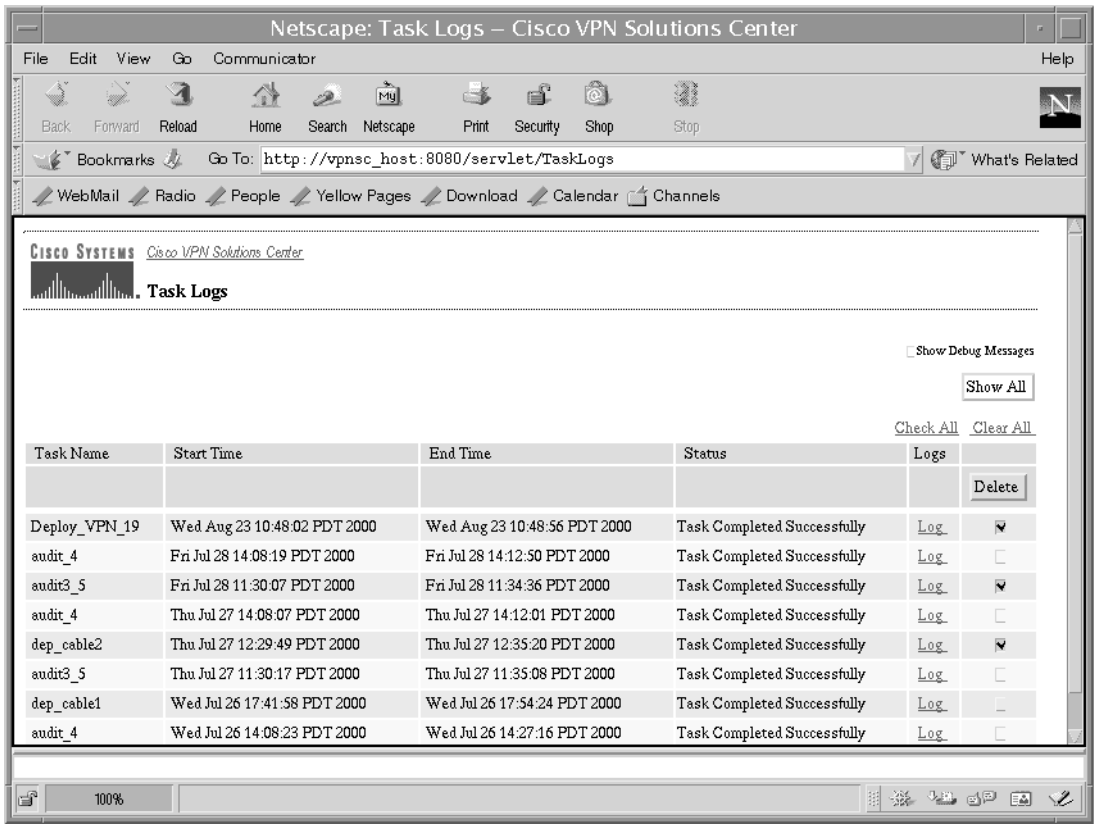

Step 2 Check the **Delete** check box for each task log you want to delete.

a. To mark all the tasks in the current page for deletion, click **Check All**.

When you click **Check All**, all the tasks on the current page are checked—not all the tasks in the Repository. Thus, when you click **Check All**, then click **Delete**, the application deletes only the tasks in the current page.

b. To clear all the check boxes in the current page, click **Clear All**.

Step 3 When you have indicated which task logs are to be deleted, click **Delete**.

The VPN Solutions Center software deletes the selected task logs and redisplays the Task Logs window.

Monitoring MPLS VPN Performance

This chapter provides an overview of performance monitoring and data collection tasks. It contains the following sections:

Updating Configuration Information on Routers in the Network

To update configuration information on routers in the network, you must collect configurations from Cisco routers, as described in this section:

Step 1 From the VPN Console, choose **Monitoring > Collect Configurations From Routers**.

The Introductory screen tells you that "This wizard sets up a scheduled task that collects Cisco router configuration files directly from the selected routers."

Click **Next**.

The Get Devices dialog box appears (see Figure 7-1).

Select Routers for Collection Task

Figure 7-1 Specifying Routers for the Collection Task

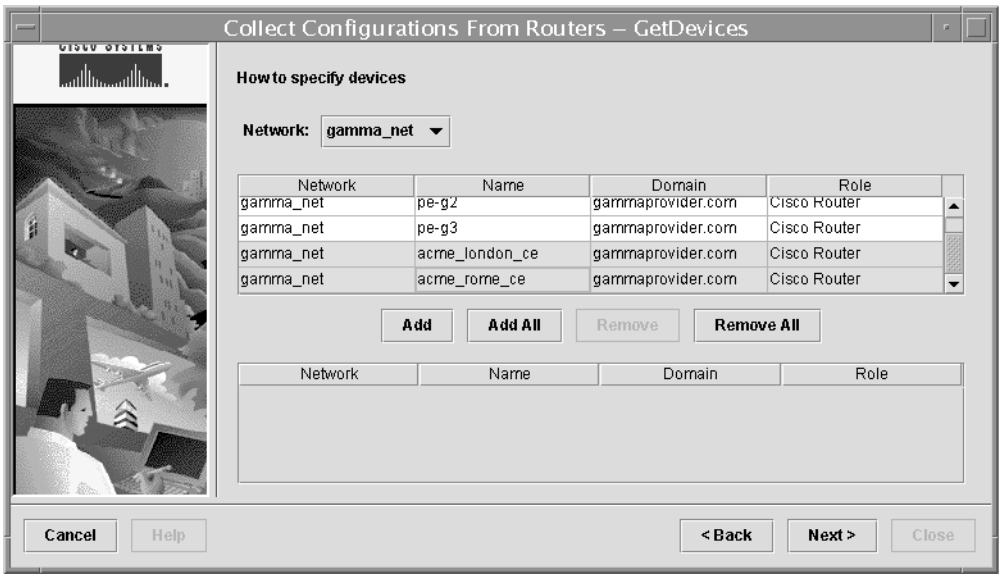

Step 2 Specify the routers that you want to collect configurations from as follows:

a. *Network*: From the Network drop-down list, choose the name of the network the routers are in.

b. From the list of routers, select one or more routers.

c. Click **Add**.

The selected routers are added to the collection list displayed in the lower pane.

If you wish to add all the routers in the selected network to the list of routers from which configurations are to be collected, click **Add All**.

Likewise, you can remove selected routers from the collection list by selecting the routers in the collection list and clicking **Remove**; or remove all the selected routers from the collection list by clicking **Remove All**.

d. When satisfied with your selections, click **Next**.

Masking Passwords in the Collected Configurations

The next dialog box lets you indicate whether you want to mask the passwords in the collected configuration files (see Figure 7-2).

Figure 7-2 Masking Passwords Option

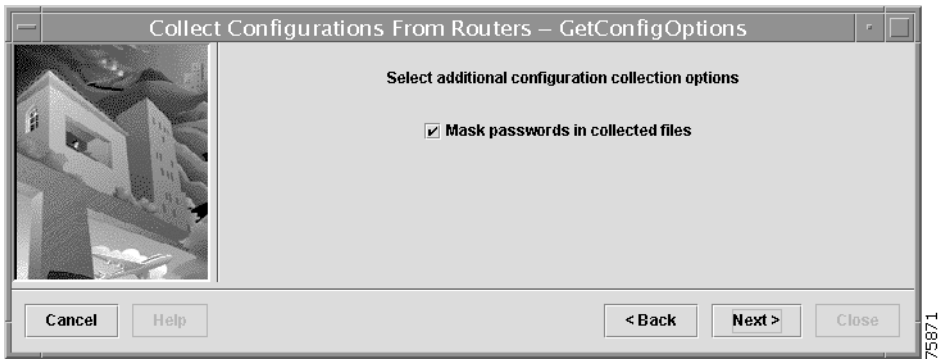

By default, the option to mask passwords is enabled. When you enable this option, VPNSC places *x* marks in the router's password field to mask the actual characters that are typed in the field.

Step 3 Indicate whether you want to mask the passwords in the collected configuration files.

- If you want to mask the passwords, click **Next**.
- If you do not choose to mask the passwords in the configurations, deselect the check in the checkbox, then click **Next**.

Provide a Task Name

The next dialog box asks you enter a task name for this collection task (see Figure 7-3).

Figure 7-3 Entering the Collection Task Name

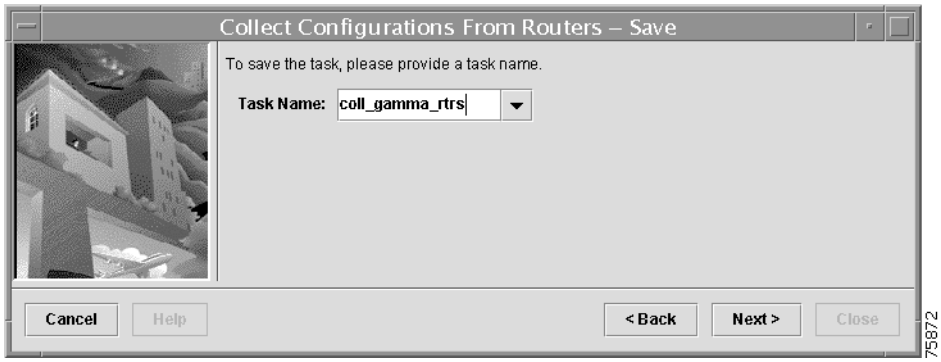

Step 4 Enter a unique name for the configuration collection task.

You can view a list of up to 30 existing task names for the appropriate task type by clicking the down arrow.

Schedule the Task

The next screen asks if you want to create a schedule for the task **Now**, in the **Future**, or **No**.

- If you choose **Now**, the task will be deployed immediately.

- If you choose **No**, the Task Manager saves the task, but the service is not scheduled for deployment.
- If you choose **Future**, the Schedule dialog box appears (see Figure 7-4).

Figure 7-4 The Schedule Dialog Box

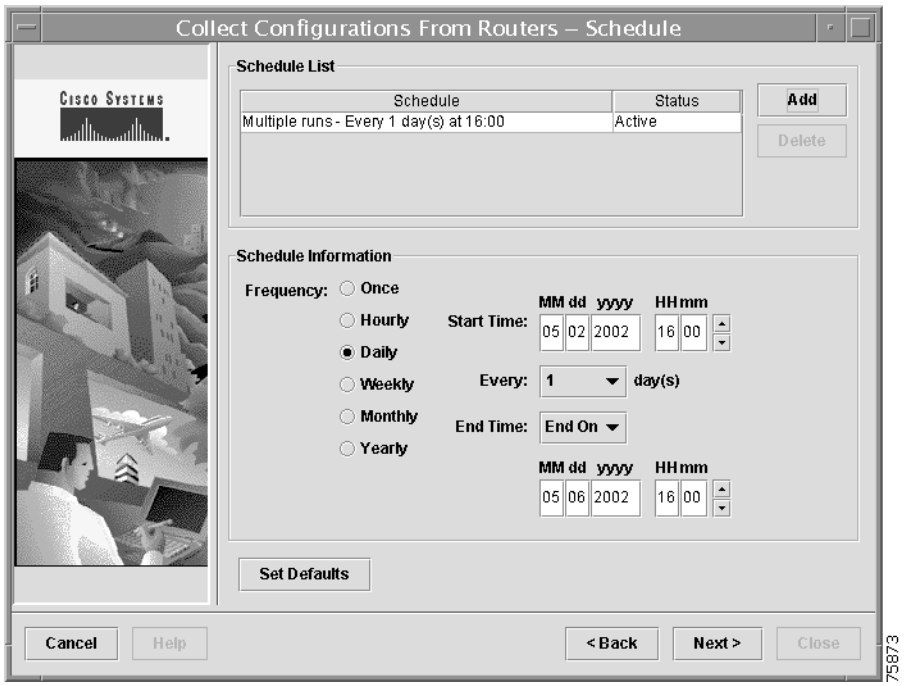

Step 5 Complete the fields in the Schedule dialog box to schedule the task as needed.

a. From the *Frequency* list, choose the desired frequency: **Once**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**.

b. Set the *Start Time*: **Now** or **Later**.

If you chose any other option except **Once**, new fields appear in the Schedule dialog box.

c. *Start Time*: If you chose **Later**, in the *Start Time* fields, specify the date and time to start the service.

d. *Every*: from the *Every* drop-down list, specify how often the service should run.

e. *End Time*: In the *End Time* drop-down list, choose either:

- **No End** for a service with no termination time and date.
- **End On** for a service that should end at a specific time and date.

f. If you selected **End On**, specify the date and time to end the service.

Step 6 When you have scheduled the configuration collection task to your satisfaction, click **Add**.

The configuration collection task is added to the schedule list (displayed in the Schedule List panel).

You can delete a configuration collection task from the schedule list by selecting the pertinent line in the schedule list and clicking **Delete**. Then click **Yes** when prompted to confirm the deletion.

Step 7 Click **Next** twice, then click **Close**.

Before You Create SLAs in VPN Solutions Center Software

Before you create SLAs in VPN Solutions Center software, you must enter some configuration changes on each CE and PE from which you want to collect performance data. PEs and CEs in the Customer's VPN must be able to communicate with the HTTP server in the service provider network.

Setting Up the Edge Devices for Gathering SLA Data

To set up edge devices for gathering performance data with SLAs, complete the following tasks:

Step 1 Enable SNMP and set the SNMP read-only and read-write community strings on all the PEs and CEs in the service provider's network. For instructions, see the following:

- "Setting Up SNMPv1 and SNMPv2 on the Routers in the Service Provider Network" section on page 2-4
- "Setting the SNMPv3 Parameters on the Routers in the Service Provider Network" section on page 2-5.

Step 2 Enable SA Agent on the CEs configured as SA Agent CEs.

SA Agent is automatically enabled when VPN Solutions Center software provisions a CE that is running SA Agent.

Setting Up VPN Solutions Center for Gathering SLA Data

To set up VPN Solutions Center for gathering performance data with SLAs, complete the following tasks:

Step 1 Verify that the edge device targets have been properly imported and defined in VPNSC.

Step 2 Make sure that when you add a CE to VPN Customer that the CE is configured as a *managed CE* with either *Regular SA Agent* status or *Shadow SA Agent* status enabled.

For the procedures to do so, see the "Specifying the Management Status for CE Routers" section on page 4-48.

Step 3 Verify the IP addresses are populated into the Repository for each target that is a source or destination for an SLA probe.

To check to see if a device's IP addresses are populated into the Repository, follow these steps:

a. Bring up the VPN Console in MPLS mode.

b. In the hierarchy pane under the Device Inventory, **double-click** the name of the desired network listed in the Networks folder.

The Network window is displayed, as shown in Figure 7-5.

Figure 7-5 The Network Window

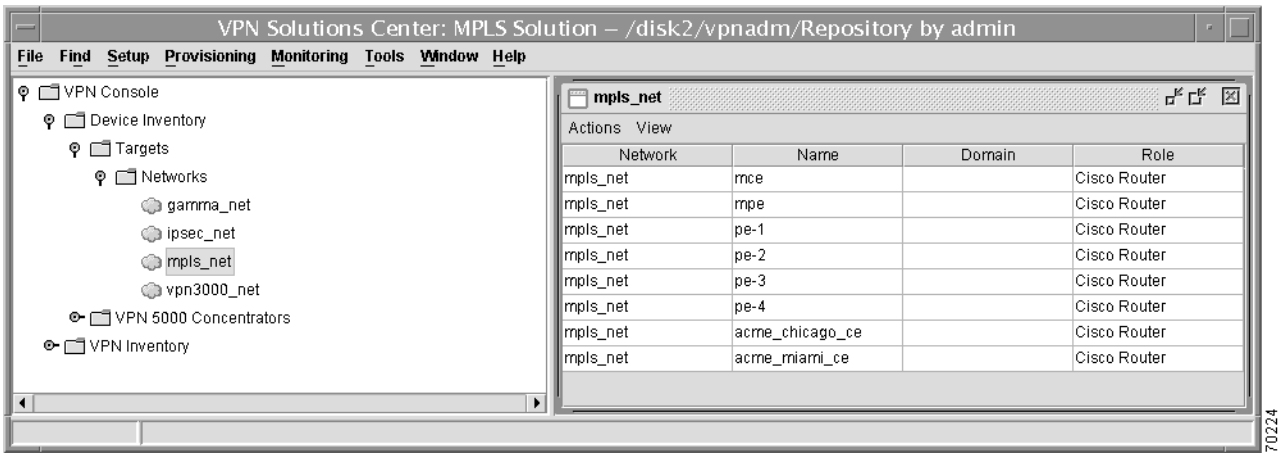

c. Select a row that lists the target router.

d. From the Network window, choose **Actions > Edit Target**.

e. Choose the **IP Addresses** tab.

A complete list of all currently populated IP addresses for the selected target (router) is displayed (see Figure 7-6).

Figure 7-6 IP Addresses Displayed for the Selected Device

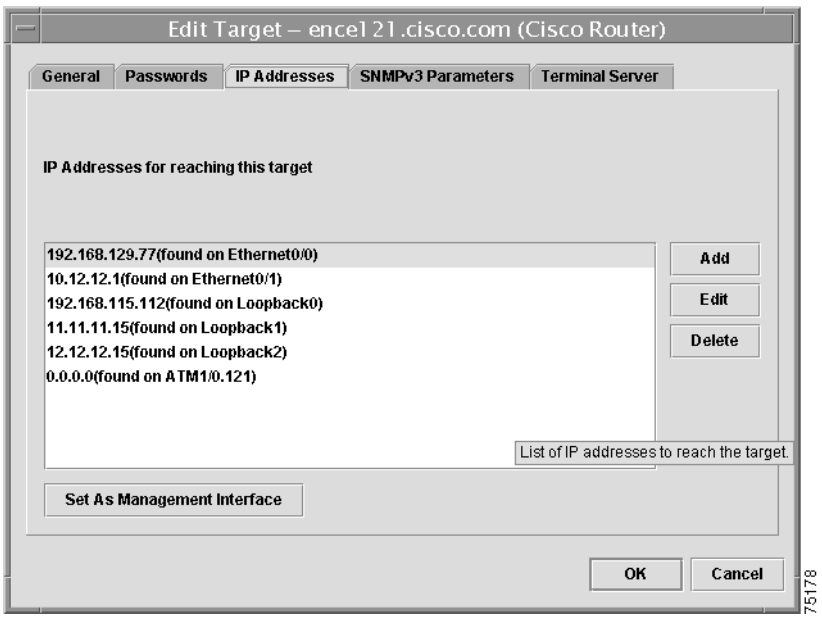

If all of the significant IP addresses are not listed, you must populate the IP addresses to the Repository. For the procedure, see the "Populating Router Interface Information to the Repository" section on page 7-31.

Provisioning Service Level Agreements

Each Service Level Agreement (SLA) is associated with a customer, the source and destination addresses on the target CEs, the protocol used for the SA Agent probe, and the threshold for delay.

In MPLS mode, VPN Solutions Center provides for SLA data collection for the following types of devices:

- Customer Edge (CE) devices (see the "Provisioning SLAs for Customer Edge Devices" section on page 7-8).
- Provider Edge (PE) devices

  Shadow CEs (a CE connected directly to a PE via Ethernet) are also considered as PEs in the context of SLA reports and APIs (see the "Provisioning VRF-Aware SLAs on PEs" section on page 7-14).
- Any SA Agent device

  This option allows you to provision SLA definitions and collect SLA data from routers that are not part of the same VPN. You can select routers from any network that has been properly defined in VPN Solutions Center (see the "Provisioning SLAs for Routers Outside a VPN" section on page 7-21).

You can create an SLA from a PE to any other PE or CE in the same VPN. When you create an SLA on a PE, you must specify a VRF name (which indicates the VPN) and an interface on the device. A PE can have the same VRF name associated with more than one interface. When provisioning a PE-to-PE SLA, the interface of the destination PE must be the interface associated with the same VPN.

When you create an SLA, VPN Solutions Center software creates an SA Agent probe on the source device.

After you provision an SLA, you must collect SLA data (see "Collecting SA Agent Data to Monitor SLAs" section on page 7-29).

When you collect data for SLA monitoring, VPN Solutions Center software downloads SLA statistics collected over the last hour from one or more specified routers. The specified routers must have the SA Agent probes configured on them. For information on defining a CE as a router running SA Agent, see the "Before You Create SLAs in VPN Solutions Center Software" section on page 7-5.

Note When you initially create an SLA, you must wait at least sixty minutes before attempting to view SLA data. If you try to view SLA data before sixty minutes elapses, the data will not yet be available and the SLA reports will be empty.

Provisioning SLAs for Customer Edge Devices

To create an SLA for Customer Edge devices, follow these steps:

Step 1 From the VPN Console, choose **Monitoring > Provision SLA Definitions and Collect SLA Data > From Customer Edge Devices** (see Figure 7-7).

Figure 7-7 Provision SLA Menu

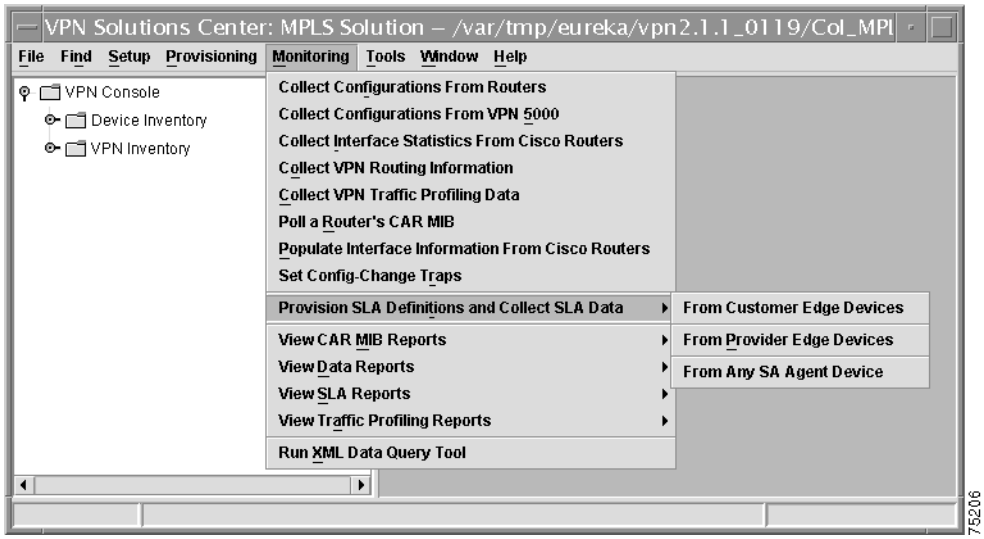

The first wizard window is informational. Click **Next** to continue.

The Specify SLA Operation dialog box appears (see Figure 7-8).

Figure 7-8 Specifying the SLA Operation

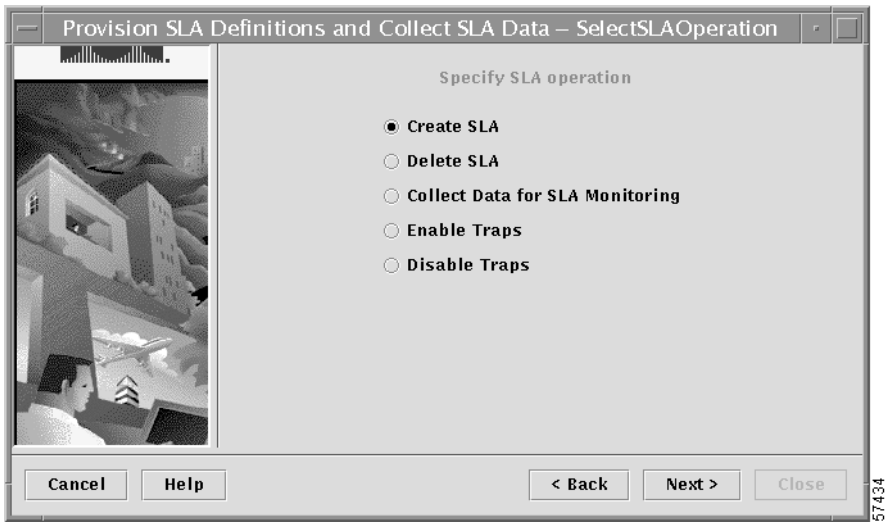

As shown in Figure 7-8, you can create an SLA, delete an SLA, or collect data for SLA monitoring.

You can also enable or disable traps. For more information, see the "Enabling Traps for SLA Data" section on page 7-38 and the "Disabling Traps" section on page 7-42.

For information on creating, selecting, and deleting SLAs for APIs by using the command line interface, see the *Cisco VPN Solutions Center: MPLS Solution API Programmer Guide*.

Step 2 To create an SLA from a CE router, choose **Create SLA**, then click **Next.**

The dialog box shown in Figure 7-9 directs you to select the source CE (or CEs)—that is, the CE you select here sends the SLA probe.

Figure 7-9 Selecting the Source CE(s) for the SLA Probe

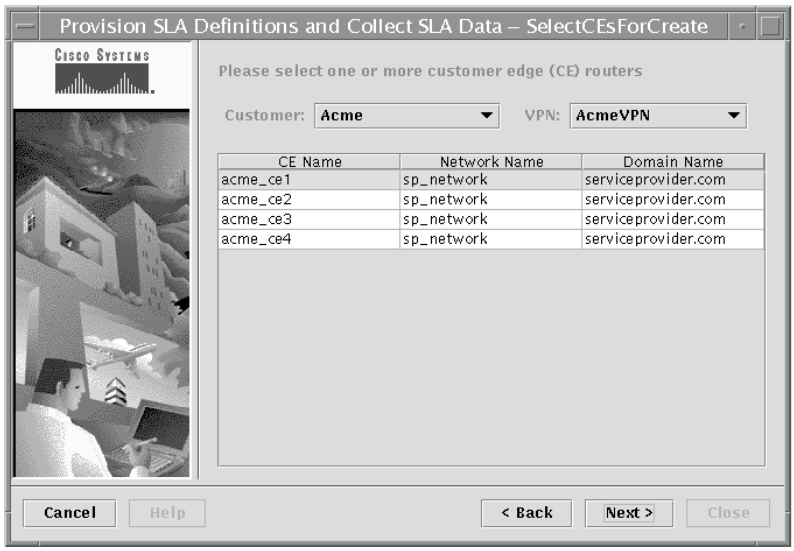

Step 3 Select one or more source CEs for the SLA probe.

a. *Customer*: From the Customer drop-down list, choose the name of the customer.

b. *VPN*: From the VPN drop-down list, choose the name of the VPN.

c. Select one or more source CEs for the SLA probe, then click **Next**.

To select multiple CEs from the list, hold down the **Ctrl** key, then click the additional router names.

As shown in Figure 7-10, the next dialog box directs you to indicate the source IP address for the source CE.

Figure 7-10 Select Source IP Address for SLA Probe

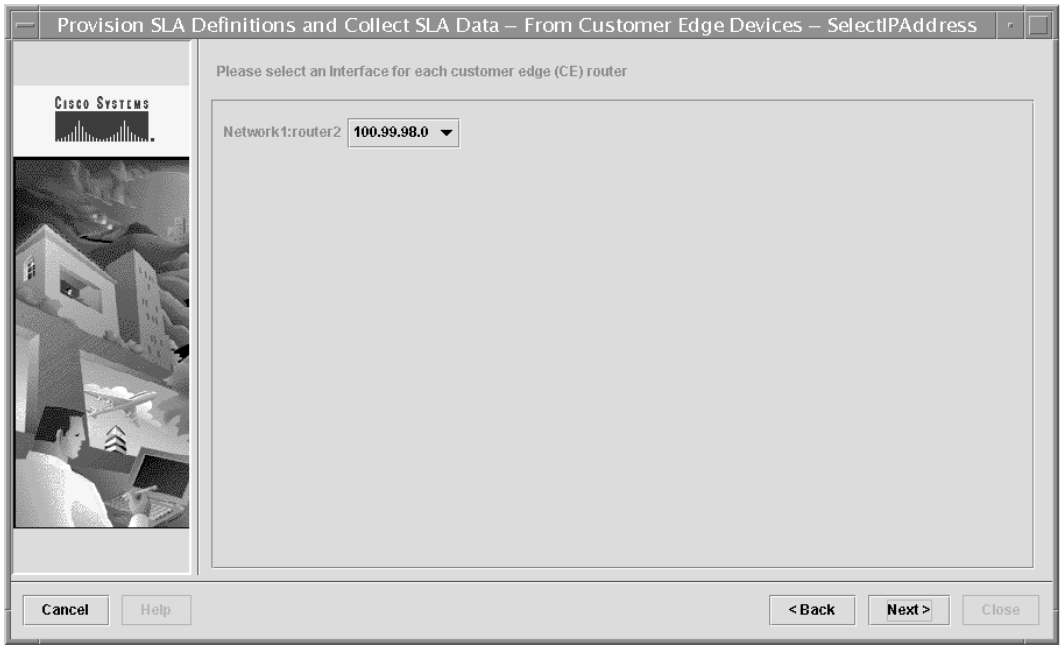

Step 4 From the drop-down list, choose the IP address for the appropriate interface on the source CE. The name of the selected CE is displayed to the left of the IP address.

When finished, click **Next**.

The next dialog box (see Figure 7-11) directs you to specify the common parameters for the SLA.

Figure 7-11 Specify SLA Common Parameters

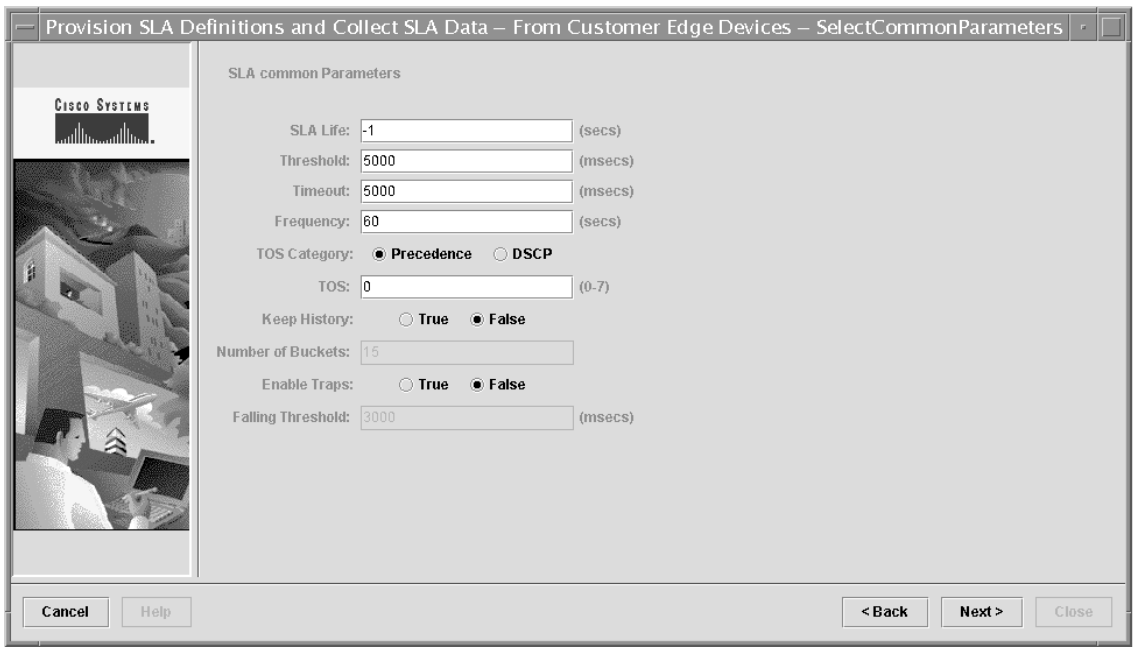

Step 5 Enter the appropriate values for the SLA parameters common to each of the SLA protocols, then click **Next**.

The fields in the SLA Common Parameters dialog box are as follows:

a. *SLA Life* is the number of seconds that the probe will be active (with the maximum value of a 32-bit integer in seconds). If the value is set to **-1**, the typical value, the probe is active indefinitely. The default value is **-1**.

b. *Threshold* is an integer that defines the threshold limit in milliseconds. The maximum value is the maximum value of a 32-bit integer. If the SA Agent operation time exceeds this limit, the threshold violation is recorded by the SA Agent. The default value is **5000**.

c. *Timeout* is the duration in milliseconds to wait for an SA Agent operation completion. The value for *Timeout* must be less than the value for *Frequency*. The default value is **5000**.

d. *Frequency* is the duration in seconds between initiating each SA Agent operation. The default value is **60**.

e. *TOS Category*:

- **Precedence**: Specifies the importance or priority of the traffic. The Precedence designation is to be used within a network only.
- **DSCP**: The Differentiated Service Code Point (DSCP) octet, in the IP header, classifies the packet service level. The DSCP maps to a particular observable forwarding behavior called a Per Hop Behavior (PHB). The DSCP replaces the ToS octet in the IPv4 header, and the Class octet in the IPv6 header. Currently, only the first six bits are used, allowing up to 64 different classifications for service levels. The DSCP is unstructured, but it does reserve some values to maintain limited backward compatibility with the precedence bits in the ToS octet.

  For details, see the "Differentiated Service Code Point (DSCP)" section on page 7-28.

f. *TOS:*

Precedence: When the ToS Category is set to **Precedence**, the valid values are an integer ranging from **0** to **7**. These values represent the type of service precedence bits in an IP header. The default value is **0**. Table 7-1 defines the *TOS* precedence values.

Table 7-1 ToS Precedence Values in SLA Parameters

ToS Value	Binary Value	Meaning
7	111	In contract, best class
6	110	In contract, second best class
5	101	In contract, third best class
4	100	In contract, worst class
3	011	Out of contract, best class
2	010	Out of contract, second best class
1	001	Out of contract, third best class
0	000	Out of contract, worst class

DSCP: When the ToS Category is set to **DSCP**, the valid values are an integer ranging from **0** to **63**. These values represent the type of service DSCP bits in an IP header. The default value is **0**.

For details, see the "Differentiated Service Code Point (DSCP)" section on page 7-28.

Note Type of Service does not apply to the DNS and DHCP types of SLA probes. VPNSC ignores any ToS value set for these two types of SLA probes. For example, if you first choose a ToS value of 5, then choose both the DHCP and ICMP protocols for an SLA probe, VPNSC will apply the selected ToS value to the ICMP probe only.

Step 6 Set the next set of SLA parameters as necessary:

a. *Keep History*

The VPN Solutions Center history table records the round trip time (that is, the delay) of operations in milliseconds. The history table does not apply to the jitter and http SLA probes.

The statistics table, which is unrelated to the history table, records the sum of the round trip times, calculates averages, and records the minimum and maximum delay values.

When you set the *Keep History* parameter to **True**, it configures the SLA probe to keep both the history table and statistics table.

b. *Numbered Buckets*

The *Numbered Buckets* parameter determines the number of samples saved for each operation. This parameter indicates the number of history delay values retained in the history table.

c. *Enable Traps*

When you set *Enable Traps* for a new SLA probe, the traps are set before the SLA operation activates. VPN Solutions Center sends a trap in the event of a timeout, a connection loss, or threshold violation (see also "Enabling Traps for SLA Data" section on page 7-38).

d. *Falling Threshold*

If you enable traps for an SLA, you must specify the *Falling Threshold* value, which triggers a threshold resolution trap. The default is 3000 milliseconds.

The next dialog box (see Figure 7-12) directs you to specify the type of SLA protocol and set its corresponding parameters.

Figure 7-12 Select SLA Protocol Data Parameters

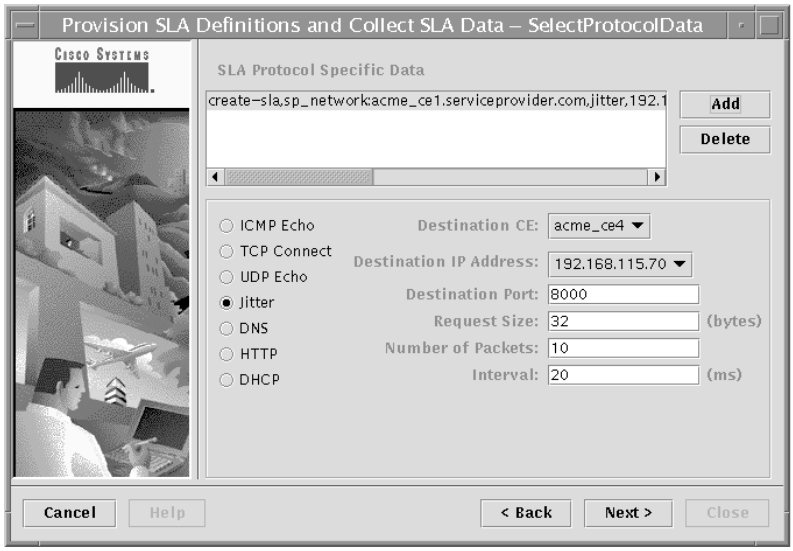

Step 7 Select an SLA protocol.

- Internet Control Message Protocol Echo (ICMP Echo)
- Transmission Control Protocol Connect (TCP Connect)
- User Datagram Protocol Echo (UDP Echo).
- Jitter (voice jitter)

  If you use the Jitter protocol, you must manually enable SA Agent on the target devices. For instructions, see the "Enabling SA Agent on Edge Routers for SLA Jitter Probes" section on page 2-7.
- Dynamic Host Configuration Protocol (DHCP)
- Hyper text Transfer Protocol (HTTP)
- Domain Name System (DNS)

a. For each SLA protocol selected, enter the desired values for the fields associated with them.

You can add additional SLA protocols as necessary.

b. When finished specifying the SLA protocol probe parameters, click **Add**.

The SLA protocol probe parameters are displayed in the SLA Protocol Specific Data pane at the top of the dialog box.

c. When satisfied with the SLA protocol data settings, click **Next**.

For details on the parameters and values for each SLA protocol listed here, refer to "Provision SLA Definitions and Collect SLA Data" in Chapter 9 of the *Cisco VPN Solutions Center: MPLS Solution User Reference*.

The dialog box shown in Figure 7-13 directs you to select the SNMP security level for the SLA.

Figure 7-13 Specifying the SNMP Security Level

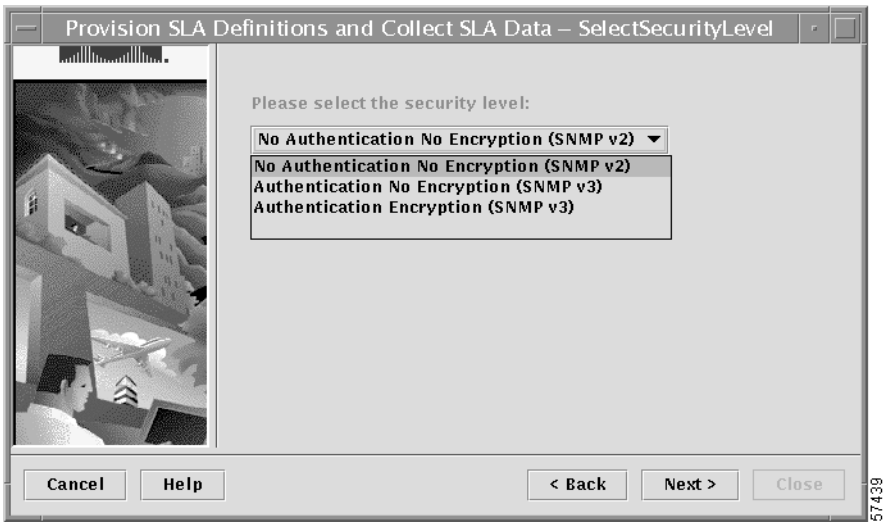

Step 8 From the drop-down list, choose the appropriate SNMP security level:

- *No Authentication, No Encryption (SNMPv2)*
- *Authentication, No Encryption (SNMPv3)*
- *Authentication, Encryption (SNMPv3)*

When you have selected the SNMP security level for this SLA, click **Next**.

Step 9 Enter a unique task name, then click **Next**.

To help you specify a unique task name, the Task Name drop-down list shows the list of existing task names.

The dialog box shown in Figure 7-14 asks if and when you want to schedule the task.

Figure 7-14 Specifying When to Run the Task

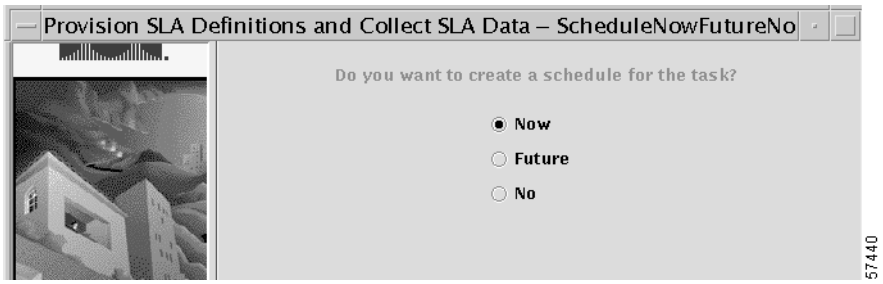

You have three options:

- **Now**. The task is scheduled to run immediately.
- **Future**. The Schedule dialog box appears.
- **No**. The SLA task is canceled.

Step 10 To run the task now, choose **Now**; to schedule the task, choose **Future**, then click **Next**.

If you choose to schedule the task for some time in the future, the Schedule dialog box appears.

Step 11 Set all the pertinent scheduling information in the Schedule dialog box, then click **Add**.

The SLA is added to the Schedule List (and displayed in the upper pane).

Step 12 Click **Next** twice, then click **Close**.

When you have collected data for SLAs, you can view the data (see the "Viewing SLA Reports" section on page 7-44).

Provisioning VRF-Aware SLAs on PEs

With Cisco IOS 12.2.1(T) and later, it is possible to configure Provider Edge routers with SLAs that monitor the VPN routes. A CE that is directly connected to the PE via Ethernet (called a *shadow CE*) monitors SLAs for the attached PE. The SLAs configured on the shadow CE monitor the VPN routes.

VPN Solutions Center 2.2 supports provisioning SLAs on shadow CEs, as well as on PEs. Shadow CEs and PEs are both considered as PEs with regard to SLA reports and APIs. VPNSC supports three types of SLA probes on PEs:

- Internet Control Message Protocol Echo (ICMP Echo)
- User Datagram Protocol Echo (UDP Echo).

- Jitter (voice jitter)

  If you use the Jitter protocol, you must manually enable SA Agent on the target devices. For instructions, see the "Enabling SA Agent on Edge Routers for SLA Jitter Probes" section on page 2-7.

It is the service provider's responsibility to make sure the selected PE for the SLA task meets these conditions:

- The router must have the correct Cisco IOS image installed—12.2.1(T) and later.
- The selected PEs must be SA Agent-enabled devices.
- The SNMP parameters must be configured appropriately on the PEs (see "Setting the SNMPv3 Parameters on the Routers in the Service Provider Network" section on page 2-5).
- The SNMP parameters must be configured appropriately in VPN Solutions Center software (see the "Specifying the Default SNMPv3 Attributes for PEs" section on page 4-22 and the "Specifying the SNMPv3 Attributes for CEs" section on page 4-58.

If any of these conditions are not met, the SLA provisioning task fails.

You can create SLAs from a PE to any other PE or CE in the same VPN. The procedure for creating an SLA on a PE includes selecting a VRF name and PE interface. When you specify the VRF name, the VPN is selected as well (since the VRF name includes the VPN name).

When you provision a PE-to-PE SLA, make sure that the interface of the destination PE is associated with the same VPN as the source PE.

To provision VRF-aware SLAs on PEs, follow these steps:

Step 1 From the VPN Console, choose **Monitoring > Provision SLA Definitions and Collect SLA Data > From Provider Edge Devices** (see Figure 7-15).

Figure 7-15 Provision SLA Menu

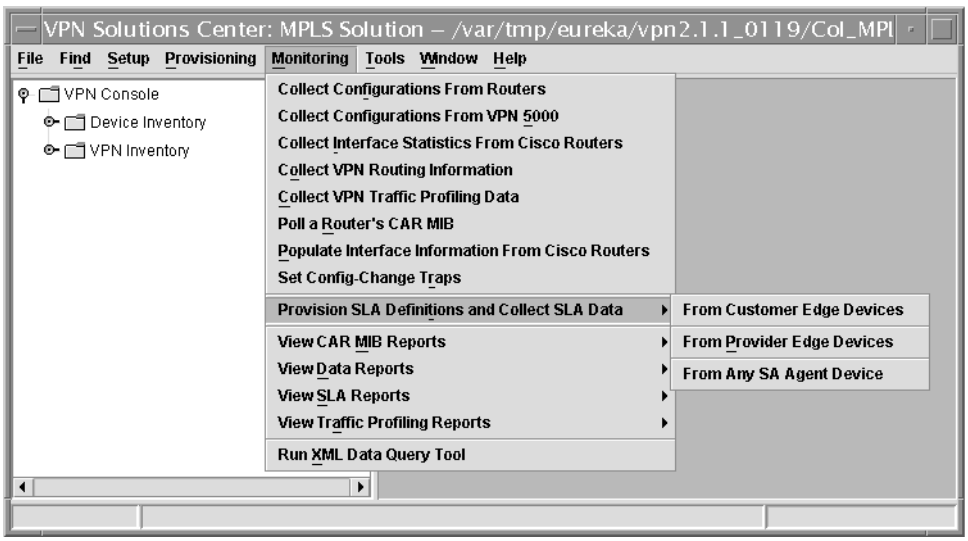

The first wizard window is informational. Click **Next** to continue.

The Specify SLA Operation dialog box appears (see Figure 7-16).

Figure 7-16 Specifying the SLA Operation

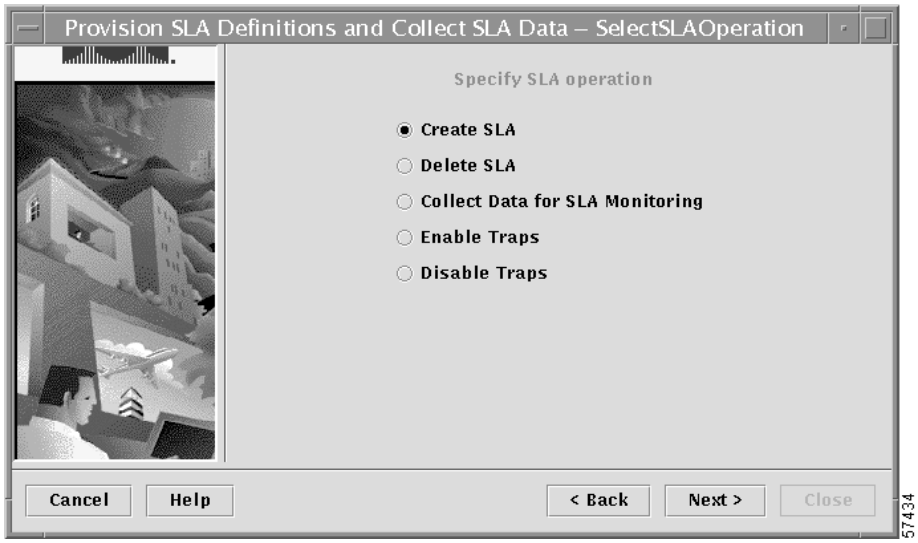

As shown in Figure 7-16, you can create an SLA, delete an SLA, or collect data for SLA monitoring.

Step 2 To create an SLA from a CE router, choose **Create SLA**, then click **Next.**

The dialog box shown in Figure 7-17 directs you to select the source PE (or PEs)—that is, the PE you select here sends the SLA probe.

Figure 7-17 Selecting the Source PE(s) for the SLA Probe

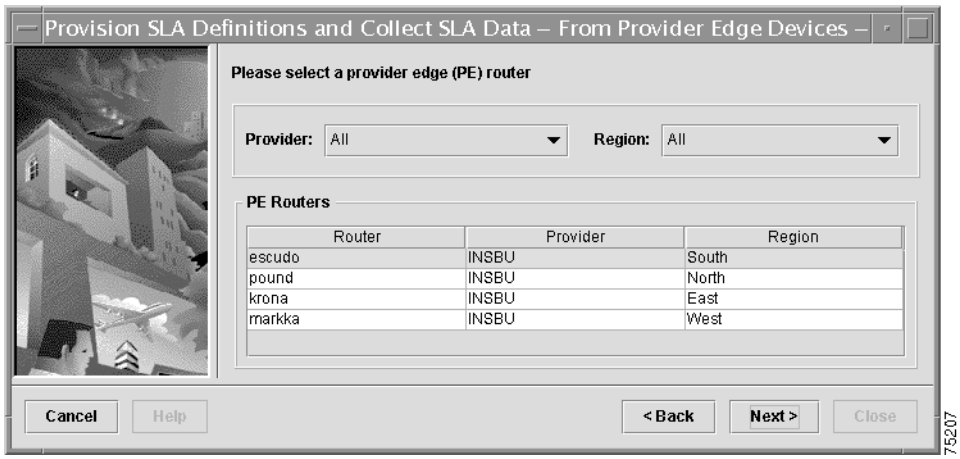

Step 3 Select one or more source PEs for the SLA probe:

a. *Provider*: From the Provider drop-down list, choose the name of the service provider.

b. *Region*: From the Region drop-down list, choose the name of the Region.

c. *PE Routers*: Select one or more source PEs for the SLA probe, then click **Next**.

To select multiple PEs from the list, hold down the **Ctrl** key, then click the additional router names.

As shown in Figure 7-18, the next dialog box directs you to indicate the VRF name—and therefore the VPN—for the source PE.

Figure 7-18 Specifying the VRF Name (VPN)

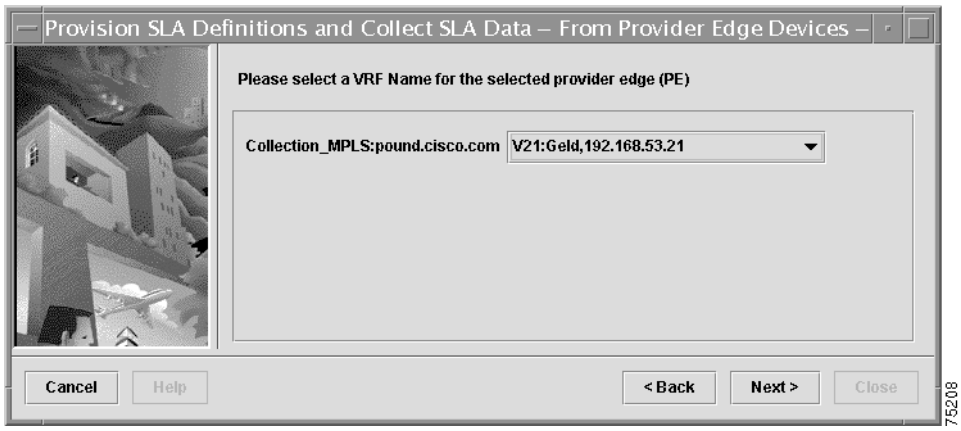

Step 4 *VRF Name (VPN)*: Specify the VPN for this PE probe by selecting the appropriate VRF name from the drop-down list, t hen click **Next**.

The next dialog box (see Figure 7-19) directs you to specify the common parameters for the SLA.

Figure 7-19 Specify SLA Common Parameters

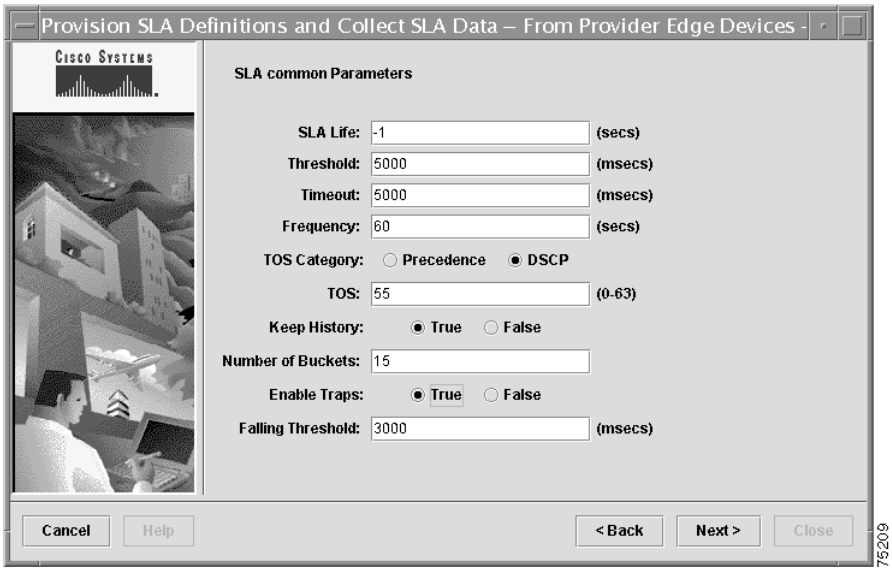

Step 5 Enter the appropriate values for the SLA parameters common to each of the SLA protocols, then click **Next**.

The fields in the SLA Common Parameters dialog box are as follows:

a. *SLA Life* is the number of seconds that the probe will be active (with the maximum value of a 32-bit integer in seconds). If the value is set to **-1**, the typical value, the probe is active indefinitely. The default value is **-1**.

b. *Threshold* is an integer that defines the threshold limit in milliseconds. The maximum value is the maximum value of a 32-bit integer. If the SA Agent operation time exceeds this limit, the threshold violation is recorded by the SA Agent. The default value is **5000**.

c. *Timeout* is the duration in milliseconds to wait for an SA Agent operation completion. The value for *Timeout* must be less than the value for *Frequency*. The default value is **5000**.

d. *Frequency* is the duration in seconds between initiating each SA Agent operation. The default value is **60**.

e. *TOS Category*:

- **Precedence**: Specifies the importance or priority of the traffic. The Precedence designation is to be used within a network only.
- **DSCP**: The Differentiated Service Code Point (DSCP) octet, in the IP header, classifies the packet service level. The DSCP maps to a particular observable forwarding behavior called a Per Hop Behavior (PHB). The DSCP replaces the ToS octet in the IPv4 header, and the Class octet in the IPv6 header. Currently, only the first six bits are used, allowing up to 64 different classifications for service levels. The DSCP is unstructured, but it does reserve some values to maintain limited backward compatibility with the precedence bits in the ToS octet.

  For details, see the "Differentiated Service Code Point (DSCP)" section on page 7-28.

f. *TOS:*

Precedence: When the ToS Category is set to **Precedence**, the valid values are an integer ranging from **0** to **7**. These values represent the type of service precedence bits in an IP header. The default value is **0**. See Table 7-1 on page 7-11 for a description of the *TOS* precedence values.

DSCP: When the ToS Category is set to **DSCP**, the valid values are an integer ranging from **0** to **63**. These values represent the type of service DSCP bits in an IP header. The default value is **0**.

Note Type of Service does not apply to the DNS and DHCP types of SLA probes. VPNSC ignores any ToS value set for these two types of SLA probes. For example, if you first choose a ToS value of 5, then choose both the DHCP and ICMP protocols for an SLA probe, VPNSC will apply the selected ToS value to the ICMP probe only.

Step 6 Set the next set of SLA parameters as necessary:

a. *Keep History*

The VPN Solutions Center history table records the round trip time (that is, the delay) of operations in milliseconds. The history table does not apply to the jitter and http SLA probes.

The statistics table, which is unrelated to the history table, records the sum of the round trip times, calculates averages, and records the minimum and maximum delay values.

When you set the *Keep History* parameter to **True**, it configures the SLA probe to keep both the history table and statistics table.

b. *Numbered Buckets*

The *Numbered Buckets* parameter determines the number of samples saved for each operation. This parameter indicates the number of history delay values retained in the history table.

c. *Enable Traps*

When you set *Enable Traps* for a new SLA probe, the traps are set before the SLA operation activates. VPN Solutions Center sends a trap in the event of a timeout, a connection loss, or threshold violation (see also "Enabling Traps for SLA Data" section on page 7-38).

d. *Falling Threshold*

If you enable traps for an SLA, you must specify the *Falling Threshold* value, which triggers a threshold resolution trap. The default is 3000 milliseconds.

The next dialog box (see Figure 7-12) directs you to specify the type of SLA protocol and set its corresponding parameters.

Figure 7-20 Specify SLA Common Parameters

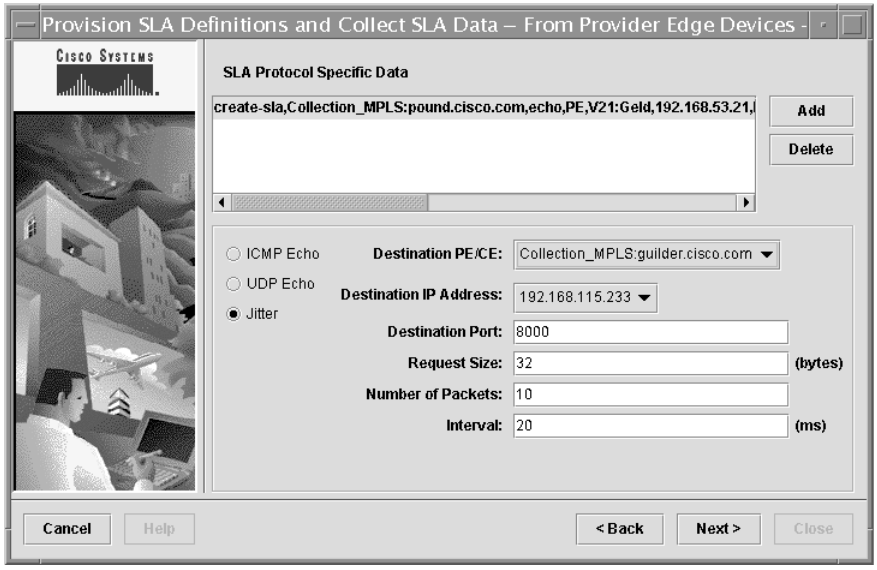

Step 7 Select the appropriate SLA protocol(s).

- Internet Control Message Protocol Echo (ICMP Echo)
- User Datagram Protocol Echo (UDP Echo).
- Jitter (voice jitter)

  If you use the Jitter protocol, you must manually enable SA Agent on the target devices. For instructions, see the "Enabling SA Agent on Edge Routers for SLA Jitter Probes" section on page 2-7.

a. For each SLA protocol selected, enter the desired values for the fields associated with them.

You can add additional SLA protocols as necessary.

b. When finished specifying the SLA protocol probe parameters, click **Add**.

The SLA protocol probe parameters are displayed in the SLA Protocol Specific Data pane at the top of the dialog box.

c. When satisfied with the SLA protocol data settings, click **Next**.

The dialog box shown in Figure 7-13 directs you to select the SNMP security level for the SLA.

Figure 7-21 Specifying the SNMP Security Level

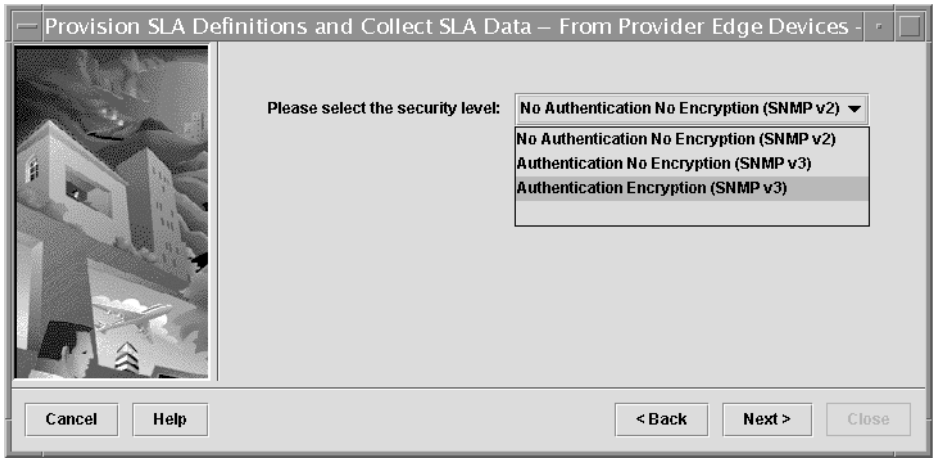

Step 8 From the drop-down list, choose the appropriate SNMP security level:

- *No Authentication, No Encryption (SNMPv2)*
- *Authentication, No Encryption (SNMPv3)*
- *Authentication, Encryption (SNMPv3)*

When you have selected the SNMP security level for this SLA, click **Next**.

Step 9 Enter a unique task name, then click **Next**.

To help you specify a unique task name, the Task Name drop-down list shows the list of existing task names.

The dialog box shown in Figure 7-14 asks if and when you want to schedule the task.

Figure 7-22 Specifying When to Run the Task

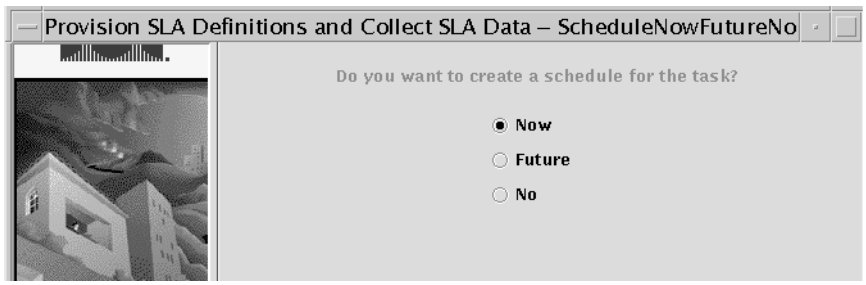

You have three options:

- **Now.** The task is scheduled to run immediately.
- **Future**. The Schedule dialog box appears.
- **No**. The SLA task is canceled.

Step 10 To run the task now, choose **Now**; to schedule the task, choose **Future**, then click **Next**.

If you choose to schedule the task for some time in the future, the Schedule dialog box appears.

Step 11 Set all the pertinent scheduling information in the Schedule dialog box, then click **Add**.

The SLA is added to the Schedule List (and displayed in the upper pane).

Step 12 Click **Next** twice, then click **Close**.

When you have collected data for SLAs, you can view the data (see the "Viewing SLA Reports" section on page 7-44).

Provisioning SLAs for Routers Outside a VPN

VPN Solutions Center provides a method to provision SLA definitions and collect SLA data from routers that are not part of a VPN. Because routers that are not part of a VPN are not necessarily associated with a customer or a VPN, you can select any Cisco router defined as a target in VPN Solutions Center software—that is, you can select routers from different networks.

If the non-VPN SLAs are have a probe type of DNS, HTTP, or DHCP, you do not have to specify a destination router. However, you must specify a destination router if the SLA is of any other probe type: ICMP, TCP, UDP, or Jitter.

To provision SLA definitions and collect SLA data for routers outside a VPN, follow these steps:

Step 1 From the VPN Console, choose **Monitoring > Provision SLA Definitions and Collect SLA Data > From Any SA Agent Device** (see Figure 7-23).

Figure 7-23 Provision SLA Menu

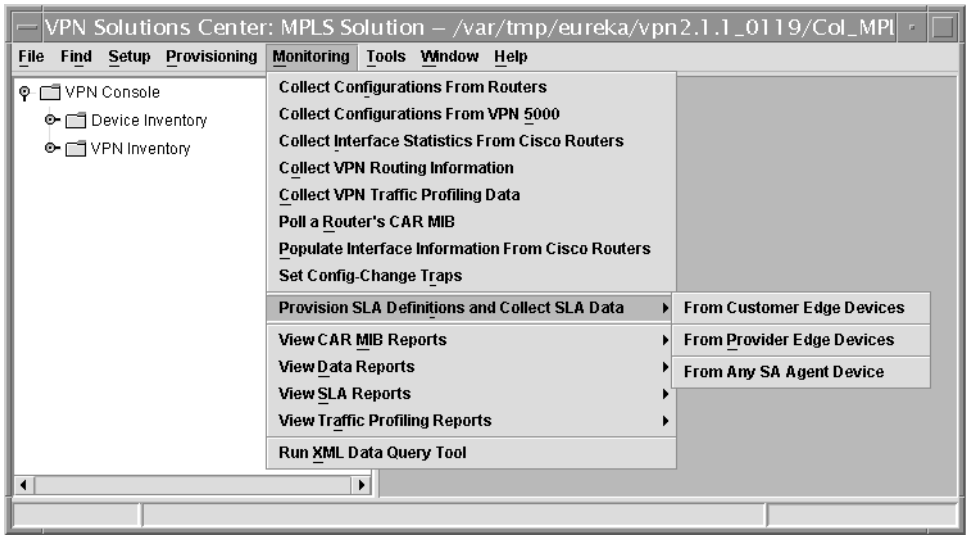

The SLA task wizard is launched and presents the introductory screen. Click **Next**.

The following dialog box appears (see Figure 7-24).

Figure 7-24 Selecting Source Routers for SLA

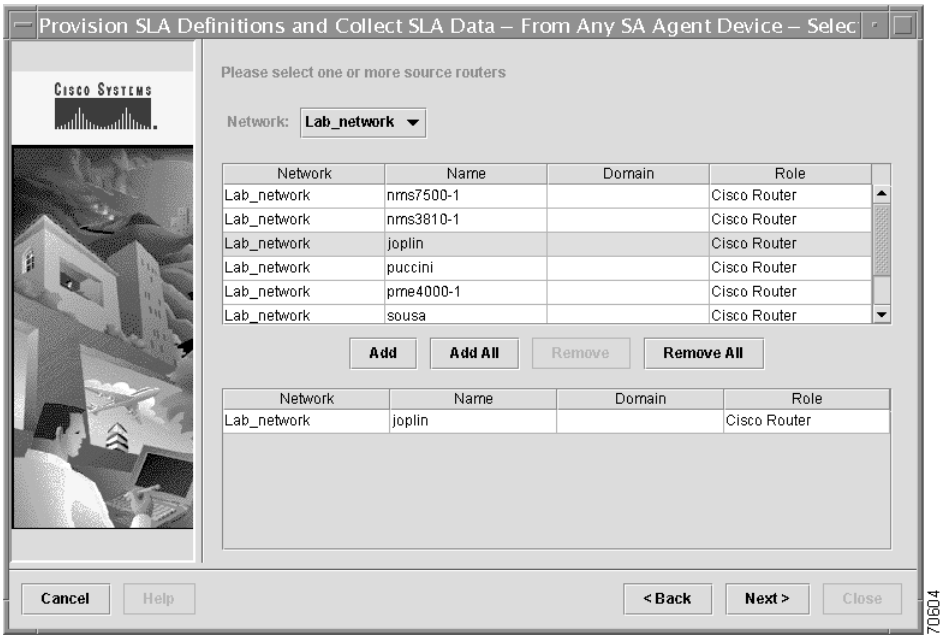

Step 2 Select one or more source routers for the SLA probe.

a. *Network*: From the drop-down list, select the network for the source router(s).

b. From the list of devices in the selected network, select one or more routers that will be the source devices for the SLA probe.

c. Click **Add**.

The selected routers are added to the list of source routers displayed in the lower pane.

If you wish to add all the routers in the selected network to the list of source routers, click **Add All**.

Likewise, you can remove selected routers from the list by selecting the routers in the list and clicking **Remove**; or remove all the selected routers from the list by clicking **Remove All**.

d. When satisfied with your selections, click **Next**.

The following dialog box appears (see Figure 7-25).

Figure 7-25 Specifying the IP Addresses for Each Source Router

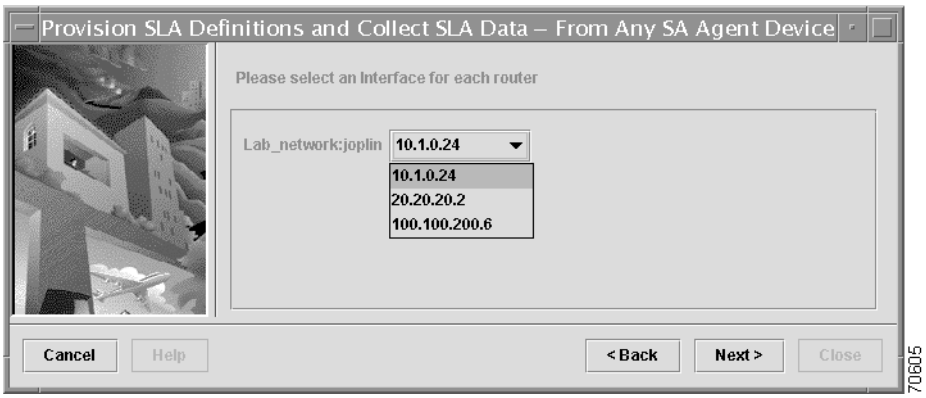

In the example shown in Figure 7-25, only one device IP address list appears in the dialog box; when you select multiple source devices, this dialog box displays a list of IP addresses for each selected source device.

Step 3 From the drop-down list, choose the appropriate IP addresses on each selected source router, then click **Next**.

The following dialog box appears (see Figure 7-26).

Figure 7-26 Selecting the Destination Routers for the SLA

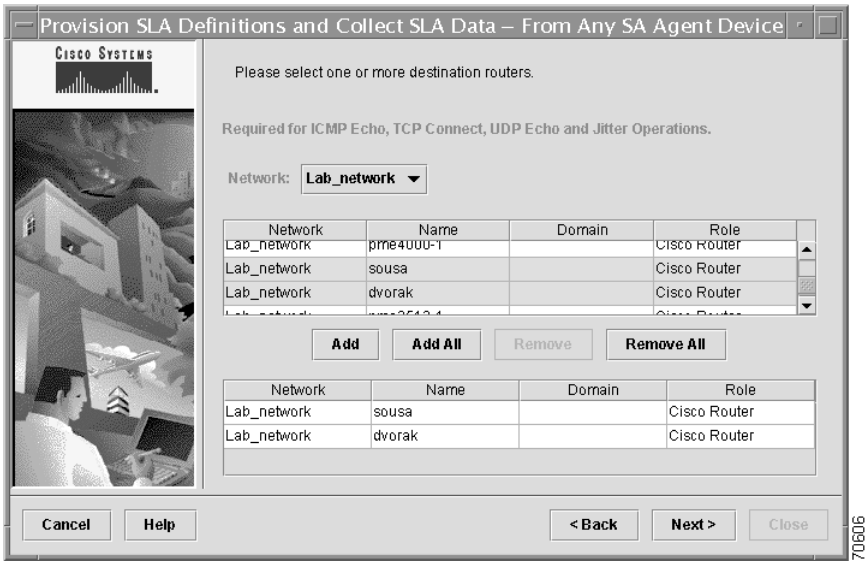

Step 4 Select the destination routers for the SLA probe.

a. *Network*: From the drop-down list, select the network for the destination router(s).

b. From the list of devices in the selected network, select one or more routers that will be the destination devices for the SLA probe.

c. Click **Add**.

The selected routers are added to the list of destination routers displayed in the lower pane.

If you wish to add all the routers in the selected network to the list of destination routers, click **Add All**.

Likewise, you can remove selected routers from the list by selecting the routers in the list and clicking **Remove**; or remove all the selected routers from the list by clicking **Remove All**.

d. When satisfied with your selections, click **Next**.

The following dialog box appears (see Figure 7-27).

Figure 7-27 Specifying the IP Addresses for Each Destination Router

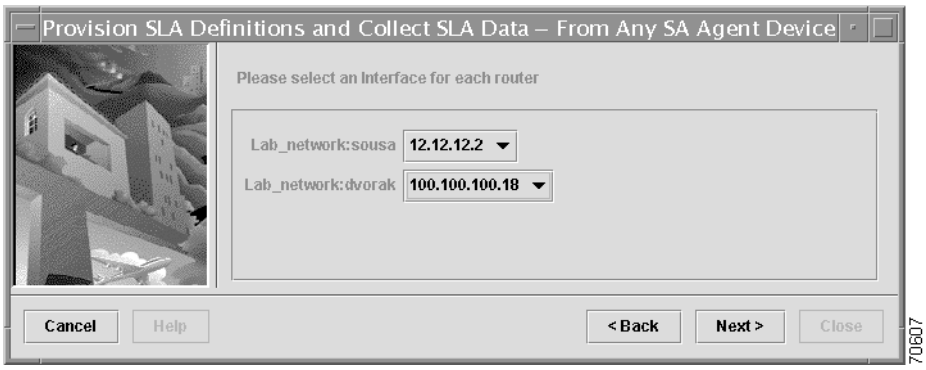

Step 5 Specify the appropriate IP addresses for each of the destination routers selected for the SLA probe, then click **Next**.

The following dialog box appears (see Figure 7-28).

Figure 7-28 Specifying the SLA Protocol Common Parameters

Step 6 Enter the appropriate values for the SLA parameters common to each of the SLA protocols being used in the network, then click **Next**.

a. *SLA Life* is the number of seconds that the probe will be active (with the maximum value of a 32-bit integer in seconds). If the value is set to **-1**, the typical value, the probe is active indefinitely. The default value is **-1**.

b. *Threshold* is an integer that defines the threshold limit in milliseconds. The maximum value is the maximum value of a 32-bit integer. If the SA Agent operation time exceeds this limit, the threshold violation is recorded by the SA Agent. The default value is **5000**.

c. *Timeout* is the duration in milliseconds to wait for an SA Agent operation completion. The value for *Timeout* must be less than the value for *Frequency*. The default value is **5000**.

d. *TOS Category*:

- **Precedence**: Specifies the importance or priority of the traffic. The Precedence designation is to be used within a network only.
- **DSCP**: The Differentiated Service Code Point (DSCP) octet, in the IP header, classifies the packet service level. The DSCP maps to a particular observable forwarding behavior called a Per Hop Behavior (PHB). The DSCP replaces the ToS octet in the IPv4 header, and the Class octet in the IPv6 header. Currently, only the first six bits are used, allowing up to 64 different classifications for service levels. The DSCP is unstructured, but it does reserve some values to maintain limited backward compatibility with the precedence bits in the ToS octet.

  For details, see the "Differentiated Service Code Point (DSCP)" section on page 7-28.

e. *Frequency* is the duration in seconds between initiating each SA Agent operation. The default value is **60**.

f. *TOS:*

Precedence: When the ToS Category is set to **Precedence**, the valid values are an integer ranging from **0** to **7**. These values represent the type of service precedence bits in an IP header. The default value is **0**. Table 7-2 defines the *TOS* precedence values.

Table 7-2 ToS Precedence Values in SLA Parameters

ToS Value	Binary Value	Meaning
7	111	In contract, best class
6	110	In contract, second best class
5	101	In contract, third best class
4	100	In contract, worst class
3	011	Out of contract, best class
2	010	Out of contract, second best class
1	001	Out of contract, third best class
0	000	Out of contract, worst class

DSCP: When the ToS Category is set to **DSCP**, the valid values are an integer ranging from **0** to **63**. These values represent the type of service DSCP bits in an IP header. The default value is **0**.

Note Type of Service does *not* apply to the DNS and DHCP types of SLA probes. VPNSC ignores any ToS value set for these two types of SLA probes. For example, if you first choose a ToS value of 5, then choose both the DHCP and ICMP protocols for an SLA probe, VPNSC will apply the selected ToS value to the ICMP probe only.

Step 7 Set the next set of SLA parameters as necessary:

a. *Keep History*

The VPN Solutions Center history table records the round trip time (that is, the delay) of operations in milliseconds. The history table does not apply to the jitter and http SLA probes.

The statistics table, which is unrelated to the history table, records the sum of the round trip times, calculates averages, and records the minimum and maximum delay values.

When you set the *Keep History* parameter to **True**, it configures the SLA probe to keep both the history table and statistics table.

b. *Numbered Buckets*

The *Numbered Buckets* parameter determines the number of samples saved for each operation. This parameter indicates the number of history delay values retained in the history table.

c. *Enable Traps*

When you set *Enable Traps* for a new SLA probe, the traps are set before the SLA operation activates. VPN Solutions Center sends a trap in the event of a timeout, a connection loss, or threshold violation (see also "Enabling Traps for SLA Data" section on page 7-38).

d. *Falling Threshold*

If you enable traps for the SLA, you must specify the *Falling Threshold* value, which triggers a threshold resolution trap. The default is **3000 milliseconds**.

The following dialog box appears (see Figure 7-29).

Figure 7-29 Specifying the Selected SLA Protocol Data Parameters

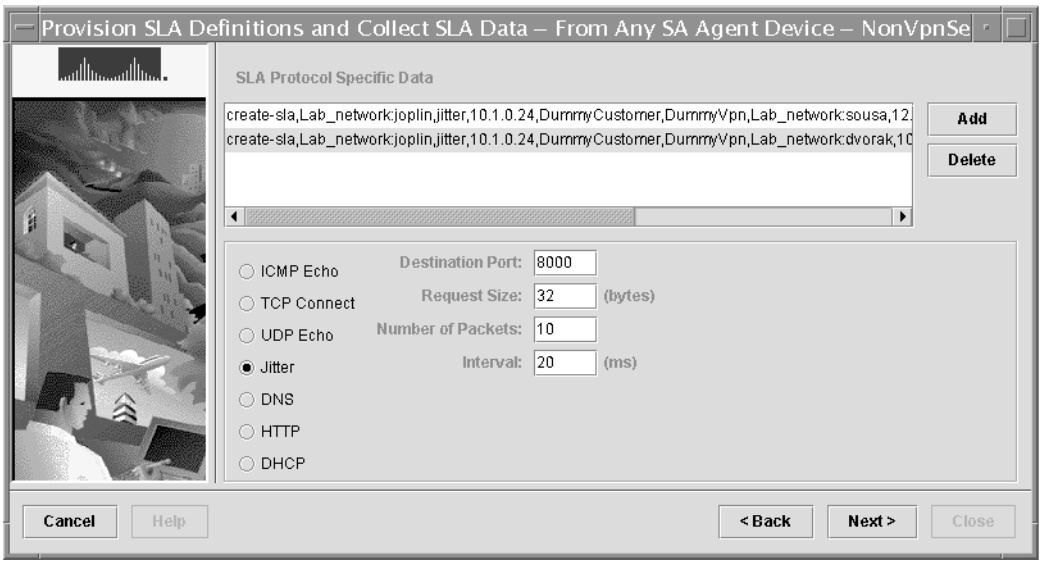

Step 8 Select an SLA protocol.

- Internet Control Message Protocol Echo (ICMP Echo)
- Transmission Control Protocol Connect (TCP Connect)
- User Datagram Protocol Echo (UDP Echo).
- Jitter (voice jitter)

  If you use the Jitter protocol, you must manually enable SA Agent on the target devices. For instructions, see the "Enabling SA Agent on Edge Routers for SLA Jitter Probes" section on page 2-7.
- Dynamic Host Configuration Protocol (DHCP)
- Hyper text Transfer Protocol (HTTP)

- Domain Name System (DNS)

a. For each SLA protocol selected, enter the desired values for the fields associated with them.

You can add additional SLA protocols as necessary.

b. When finished specifying the SLA protocol probe parameters, click **Add**.

The SLA protocol probe parameters are displayed in the SLA Protocol Specific Data pane at the top of the dialog box.

When you click **Add**, VPN Solutions Center creates SLA probes between all the possible permutations of the specified source routers and destination routers. If a router is specified as both a source router and destination router, VPNSC does not (and cannot) create a probe to and from itself.

c. When satisfied with the SLA protocol data settings, click **Next**.

The following dialog box appears (see Figure 7-30).

Figure 7-30 Specifying the SNMP Security Level

Step 9 From the Security Level drop-down list, choose the appropriate SNMP security level:

- *No Authentication, No Encryption (SNMPv2)*
- *Authentication, No Encryption (SNMPv3)*
- *Authentication, Encryption (SNMPv3)*

When you have selected the SNMP security level for this SLA, click **Next**.

Step 10 Enter a unique task name, then click **Next**.

To help you specify a unique task name, the Task Name drop-down list shows the list of last thirty existing task names.

The dialog box shown in Figure 7-31 asks if and when you want to schedule the task.

Figure 7-31 Specifying When to Run the Task

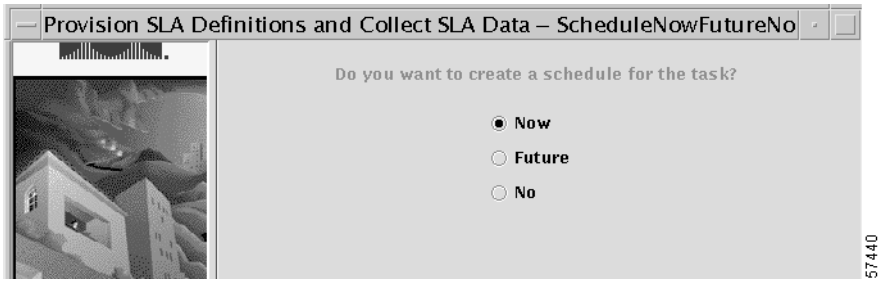

Step 11 To run the task now, choose **Now**; to schedule the task, choose **Future**, then click **Next**.

If you choose to schedule the task for some time in the future, the Schedule dialog box appears.

Step 12 Set all the pertinent scheduling information in the Schedule dialog box, then click **Add**.

The SLA is added to the Schedule List (and displayed in the upper pane).

Step 13 Click **Next** twice, then click **Close**.

When you have collected data for SLAs, you can view the data (see the "Viewing SLA Reports" section on page 7-44).

Differentiated Service Code Point (DSCP)

In the IP header, DSCP classifies the packet service level. The DSCP maps to a particular observable forwarding behavior called a Per Hop Behavior (PHB). The DSCP replaces the ToS octet in the IPv4 header, and the Class octet in the IPv6 header. Currently, only the first six bits are used, allowing up to 64 different classifications for service levels. The DSCP is unstructured, but it does reserve some values to maintain limited backward compatibility with the precedence bits in the ToS octet.

Bits 3 and 4 of the ToS field are the DSCP bits. They allow further priority granularity through the specification of a packet drop probability for any of the defined classes. Collectively, Class 1 through 4 are referred to as Assured Forwarding (AF).

Table 7-3 illustrates the DSCP coding for specifying the priority level (or class) and the drop percentage.

- Bits 0, 1, and 2 define the class.
- Bits 3 and 4 specify the drop percentage.
- Bit 5 is always 0.

Table 7-3 DSCP Coding for Specifying Class and Drop Percentage

	Class 1	Class 2	Class 3	Class 4
Low Drop Percentage	001010	010010	011010	100010
Medium Drop Percentage	001100	010100	011100	100100
High Drop Percentage	001110	010110	011110	100110

Using this system, a device would first prioritize traffic by class, then differentiate and prioritize same-class traffic by considering the drop percentage. Please note that this standard has not specified a precise definition or values for low, medium, and high drop percentages. Additionally, not all devices will recognize the DSCP settings for bits 3 and 4. The DSCP proposal allows a finer granularity of priority setting for the applications and devices that can make use of it, but it does not specify the action to be taken.

Collecting SA Agent Data to Monitor SLAs

After you provision SLAs, you must collect the SA Agent data for the SLAs.

Note When you initially create an SLA, you must wait at least sixty minutes before attempting to view SLA data. If you try to view SLA data before sixty minutes elapses, the data will not yet be available and the SLA reports will be empty.

To collect SA Agent data for SLAs, follow these steps:

Step 1 From the VPN Console, choose **Monitoring > Provision SLA Definitions and Collect SLA Data**, then choose one of the three options from the menu:

- **From Customer Edge Routers**
- **From Provider Edge Routers**
- **From Any SA Agent Device**

Step 2 The first wizard window is informational. Click **Next** to continue.

The Specify SLA Operation dialog box is displayed.

Step 3 Choose **Collect Data for SLA Monitoring**, then click **Next**.

The dialog box shown in Figure 7-32 appears.

Figure 7-32 Select Source CE for SLA Probe

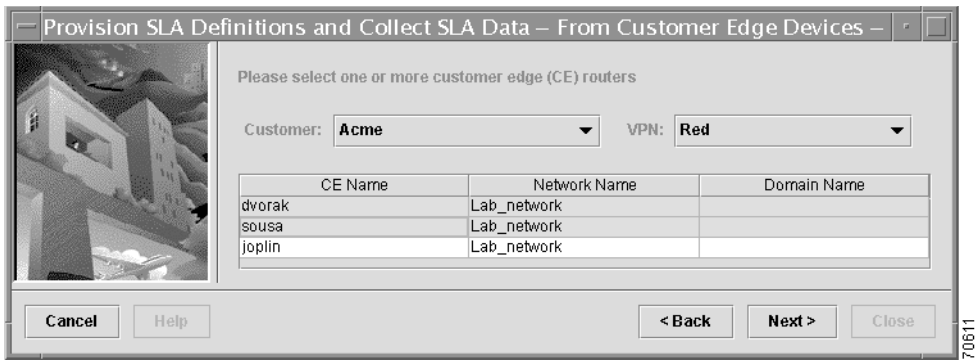

Step 4 Select the source device (or devices) for the SLA probe.

The device you select here sends the SLA probe to the routers that have SA Agent enabled.

a. *Customer*: Choose the appropriate Customer from the Customer drop-down list.

b. *VPN*: Choose the appropriate VPN from the VPN drop-down list.

c. *Devices*: Select one or more devices from which you want to collect SLA data, then click **Next**.

The following dialog box appears (see Figure 7-33).

Figure 7-33 Specifying the SNMP Security Level

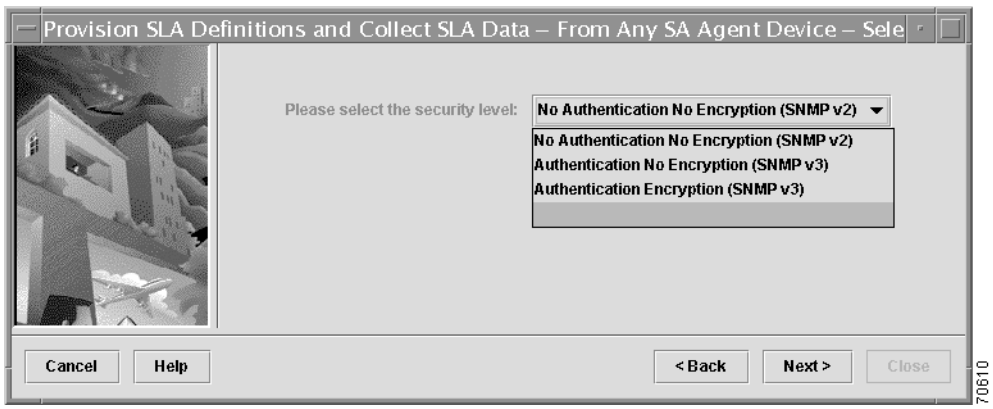

Step 5 From the Security Level drop-down list, choose the appropriate SNMP security level:

- *No Authentication, No Encryption (SNMPv2)*
- *Authentication, No Encryption (SNMPv3)*
- *Authentication, Encryption (SNMPv3)*

When you have selected the SNMP security level for this SLA, click **Next**.

Step 6 In the next dialog box, provide a unique task name, then click **Next.**

To help you specify a unique task name, the Task Name drop-down list shows the list of existing task names.

The next dialog box asks if and when you want to schedule the task. You have three options:

- *Now*. The task is scheduled to run immediately.
- *Future*. The Schedule dialog box appears.
- *No*. The SLA task is canceled.

Step 7 To run the task now, choose **Now**; to schedule the task, choose **Future**, then click **Next**.

If you choose to schedule the task for some time in the future, the Schedule dialog box appears.

Step 8 Set all the pertinent scheduling information in the Schedule dialog box, then click **Add**.

The SLA is added to the Schedule List (and displayed in the upper pane).

Step 9 To save the SA Agent collection task, click **Next**.

If you chose to schedule the SA Agent collection task, that will also occur.

You are informed that all steps are done.

Step 10 Click **Close** to close the wizard.

Collecting Changed Configuration Files Only

Router configuration files are usually collected at regular intervals and then examined for changes that affect the way the routers function. While the routers whose configuration files have changed are the only ones that need to be collected, the normal collection process does not separate the routers whose configuration files have changed from the routers whose configuration files have not.

About Smart Collector

Smart Collector finds the routers whose configuration files have changed and organizes them into a group to have their configuration files collected.

With Smart Collector, VPN Solutions Center creates a task and schedules it to be run once. When the task executes, all the targeted routers are instructed to advise the VPN Solutions Center software that uses the Simple Network Management Protocol (SNMP) of any change to their configuration files. VPN Solutions Center notes these traps and keeps track of the routers whose configuration files have changed, and thus need to be collected. The purpose of configuring traps (through Smart Collector) is to efficiently collect router configuration files from a set of routers that can belong to more than one network.

An example of the potential substantial savings this feature affords is a scenario in which a network has 200 routers, but the configuration files for only 20 of the routers have changed. In this example, Smart Collector collects only the configuration files for the 20 that have changed rather than for all 200 routers. If only 10 percent of the routers have their configuration files changed between scheduled collections, each Smart Collection takes only 10 percent of the resources of a full collection.

Note that periodically the scheduler ignores the reduced target list and collects from all routers in the original target list. Thus even those routers whose traps failed to reach VPN Solutions Center are collected periodically.

Configure SNMP on PEs and CEs

The Simple Network Management Protocol (SNMP) must be configured on each PE router and CE router in the service provider network.

To determine whether SNMP is enabled and the SNMP community strings are set on a router, see the "Setting Up SNMPv1 and SNMPv2 on the Routers in the Service Provider Network" section on page 2-4 and the "Setting the SNMPv3 Parameters on the Routers in the Service Provider Network" section on page 2-5.

Populating Router Interface Information to the Repository

Prior to registering the configuration file change traps, you must populate the router interface information into the Repository, as described in this section. This information is used to create the various reports and to map the "config-change" traps to the appropriate routers.

To populate the IP address information to the Repository, follow these steps:

Step 1 From the VPN Console, choose **Monitoring > Populate Interface Information From Cisco Routers**.

The initial screen presents the following information:

This procedure sets up a scheduled task that polls for information about router interfaces. It extracts the interface name, index number, and IP address, and subnet mask for each interface. The collected interface information is stored with each router definition.

Click **Next**.

The following dialog box appears (see Figure 7-34).

Figure 7-34 Selecting Devices For Interface Information

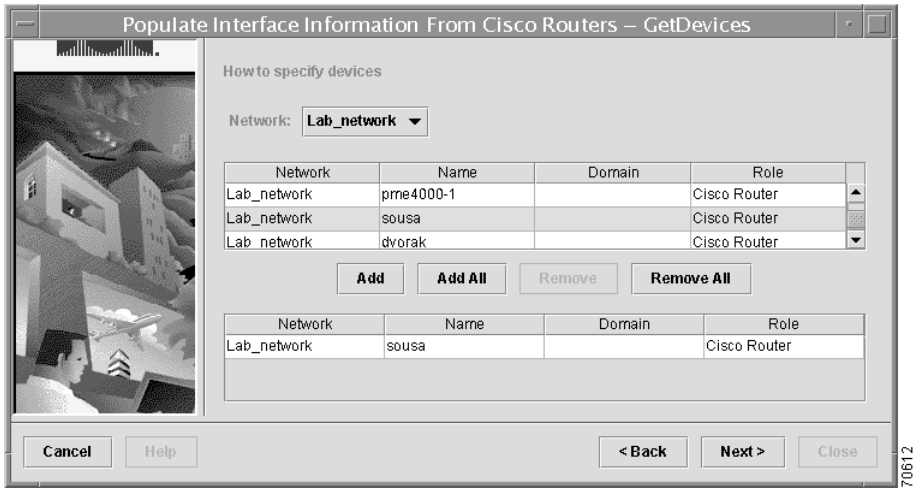

Step 2 Select one or more source routers from which you want to extract their interface information.

a. *Network*: From the drop-down list, select the network for the source router(s).

b. From the list of devices in the selected network, select one or more routers.

c. Click **Add**.

The selected routers are added to the list of source routers displayed in the lower pane.

If you wish to add all the routers in the selected network to the list of routers, click **Add All**.

Likewise, you can remove selected routers from the list by selecting the routers in the list and clicking **Remove**; or remove all the selected routers from the list by clicking **Remove All**.

d. When satisfied with your selections, click **Next**.

The Task Name dialog box appears.

Step 3 Enter a unique task name, then click **Next**.

The next dialog box asks if and when you want to schedule the task.

You have three options:

- *Now*. The interface population task is scheduled to run immediately.
- *Future*. The Schedule dialog box appears.
- *No*. The interface population task is canceled.

Step 4 To run the task now, choose **Now**; to schedule the task, choose **Future**, then click **Next**.

If you choose to schedule the task for some time in the future, the Schedule dialog box appears.

Step 5 Set all the pertinent scheduling information in the Schedule dialog box, then click **Add**.

The task is added to the Schedule List (and displayed in the upper pane).

Step 6 Click **Next** twice, then click **Close**.

Setting Traps for Changed Configuration Files

This section explains how to set traps for changed configuration files. You specify the routers for which configurations will be collected only if the routers have changed.

Step 1 From the VPN Console, choose **Monitoring > Set Config-Change Traps**.

The introductory screen appears with the following information:

This wizard allows you:

1. To configure selected Cisco routers to send configuration traps to this workstation. This enables SmartCollection, in which router configuration files are collected only from those routers that have configuration files that have actually changed.
2. To configure selected routers to stop sending configuration traps to this workstation, the routers will no longer be part of SmartCollector, the process in which configuration files are collected only if a configuration change has been detected.

Click **Next**. The following dialog box appears (see Figure 7-35).

Figure 7-35 Registering for Configuration Traps

Step 2 Choose **Register for Config Traps** (the default selection), then click **Next**.

The following dialog box appears (see Figure 7-36).

Figure 7-36 Specifying the Devices for Configuration-Change Traps

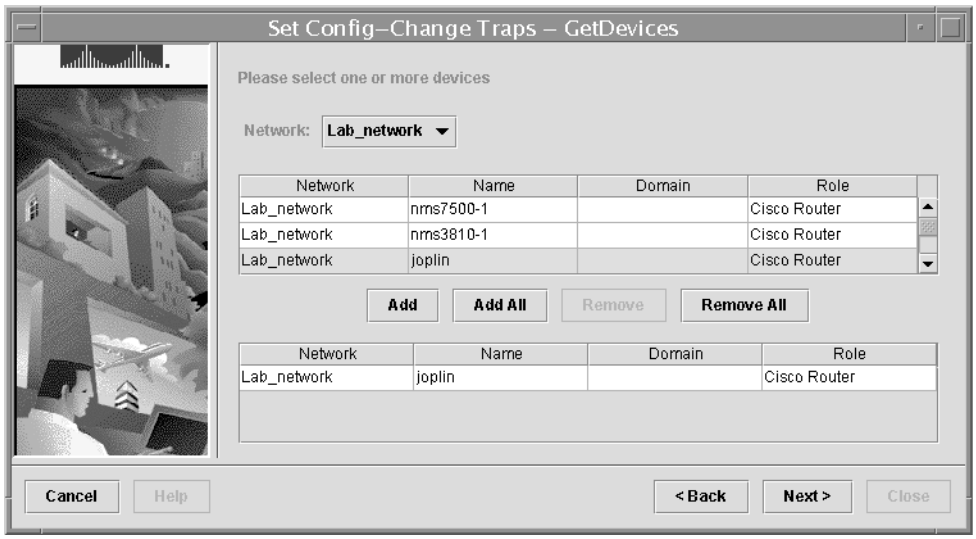

Step 3 Specify the target devices for configuration-change traps.

Note Provider Edge routers *must* run Cisco IOS version 12.x or higher to return traps.

a. *Network*: From the drop-down list, select the network where the target devices reside.

b. From the list of devices in the selected network, select one or more routers that you wish to be watched for configuration changes.

c. Click **Add**.

The selected routers are added to the list of configuration-change routers displayed in the lower pane.

If you wish to add all the routers in the selected network to the list of destination routers, click **Add All**.

Likewise, you can remove selected routers from the list by selecting the routers in the list and clicking **Remove**; or remove all the selected routers from the list by clicking **Remove All**.

d. When satisfied with your selections, click **Next**.

Step 4 Enter a unique task name, then click **Next**.

To help you specify a unique task name, the Task Name drop-down list shows the list of last thirty existing task names.

The next dialog box asks if and when you want to schedule the task.

You have three options:

- *Now*. The task is scheduled to run immediately.
- *Future*. The Schedule dialog box appears.
- *No*. The SLA task is canceled.

Step 5 To run the task now, choose **Now**; to schedule the task, choose **Future**, then click **Next**.

If you choose to schedule the task for some time in the future, the Schedule dialog box appears.

Step 6 Set all the pertinent scheduling information in the Schedule dialog box, then click **Add**.

The task is added to the Schedule List (and displayed in the upper pane).

Step 7 Click **Next** twice, then click **Close**.

Deregistering Traps for Changed Configuration Files

This section explains how to deregister the traps for that have been configured for changed configuration files. This procedure configures selected Cisco routers to stop sending "config-change" traps to the current VPN Solutions Center workstation. The selected routers will no longer be part of SmartCollection, through which configuration files are collected only from those routers whose configuration files have changed.

Step 1 From the VPN Console, choose **Monitoring > Set Config-Change Traps**.

The introductory screen appears.

Click **Next**. The following dialog box appears (see Figure 7-37).

Figure 7-37 Registering for Configuration Traps

Step 2 Choose **Deregister for Config Traps**, then click **Next**.

The following dialog box appears (see Figure 7-38).

Figure 7-38 Specifying the Devices for Configuration-Change Traps

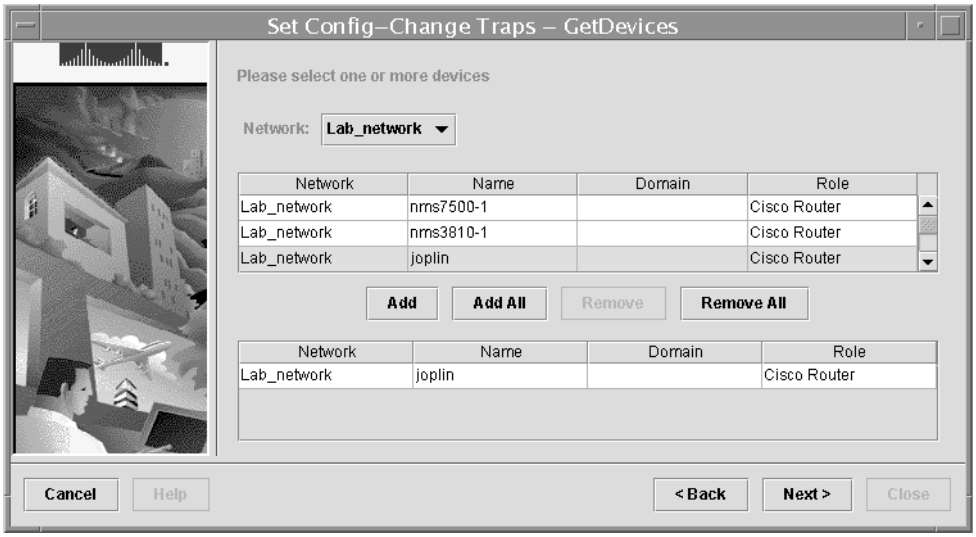

Step 3 Specify the target devices for configuration-change traps.

a. *Network*: From the drop-down list, select the network where the target devices reside.

b. From the list of devices in the selected network, select one or more routers that you wish to deregister for configuration changes.

c. Click **Add**.

The selected routers are added to the list of deregistered routers displayed in the lower pane.

If you wish to add all the routers in the selected network to the list of destination routers, click **Add All**.

Likewise, you can remove selected routers from the list by selecting the routers in the list and clicking **Remove**; or remove all the selected routers from the list by clicking **Remove All**.

d. When satisfied with your selections, click **Next**.

Step 4 Enter a unique task name, then click **Next**.

The next dialog box asks if and when you want to schedule the task.

You have three options:

- *Now*. The task is scheduled to run immediately.
- *Future*. The Schedule dialog box appears.
- *No*. The SLA task is canceled.

Step 5 To run the task now, choose **Now**; to schedule the task, choose **Future**, then click **Next**.

If you choose to schedule the task for some time in the future, the Schedule dialog box appears.

Step 6 Set all the pertinent scheduling information in the Schedule dialog box, then click **Add**.

The task is added to the Schedule List (and displayed in the upper pane).

Step 7 Click **Next** twice, then click **Close**.

Deleting an SLA

Deleting an SLA from VPN Solutions Center deletes an SA Agent probe from the source CE router.

Step 1 From the VPN Console, choose **Monitoring > Provision SLA Definitions and Collect SLA Data**, then choose one of the three options from the menu:

- **From Customer Edge Routers**
- **From Provider Edge Routers**
- **From Any SA Agent Device**

Step 2 The first wizard window is informational. Click **Next** to continue.

The Specify SLA Operation dialog box is displayed (as shown in Figure 7-8).

Step 3 To delete an SLA in VPN Solutions Center, choose **Delete SLA**, then click **Next**.

The dialog box shown in Figure 7-39 appears.

Figure 7-39 Select Source CE of SLA

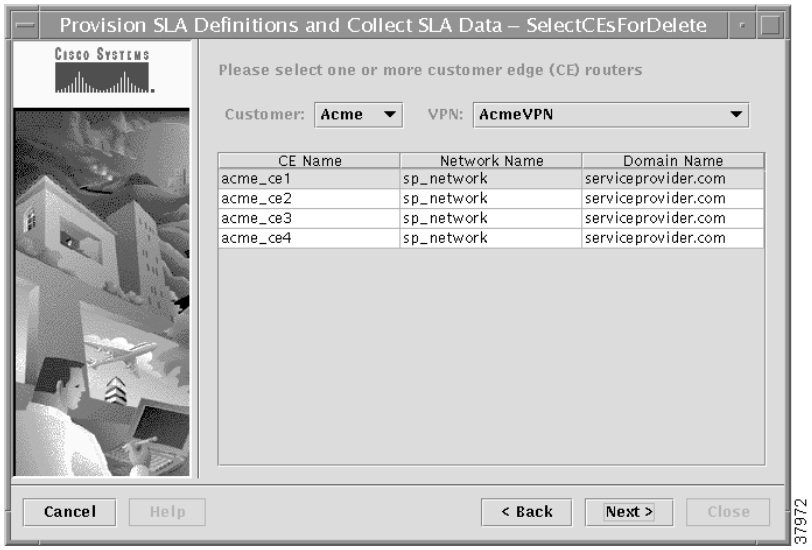

Step 4 In the **Customer** and **VPN** drop-down lists, select the pertinent Customer name and VPN name.

The CE pane lists all the CEs in the selected VPN that are running SA Agent.

Step 5 Select the name of the source CE for the SLA probe you want to delete, then click **Next**.

The next dialog box directs you to select the SLA you want to delete.

Figure 7-40 Select SLAs to Delete

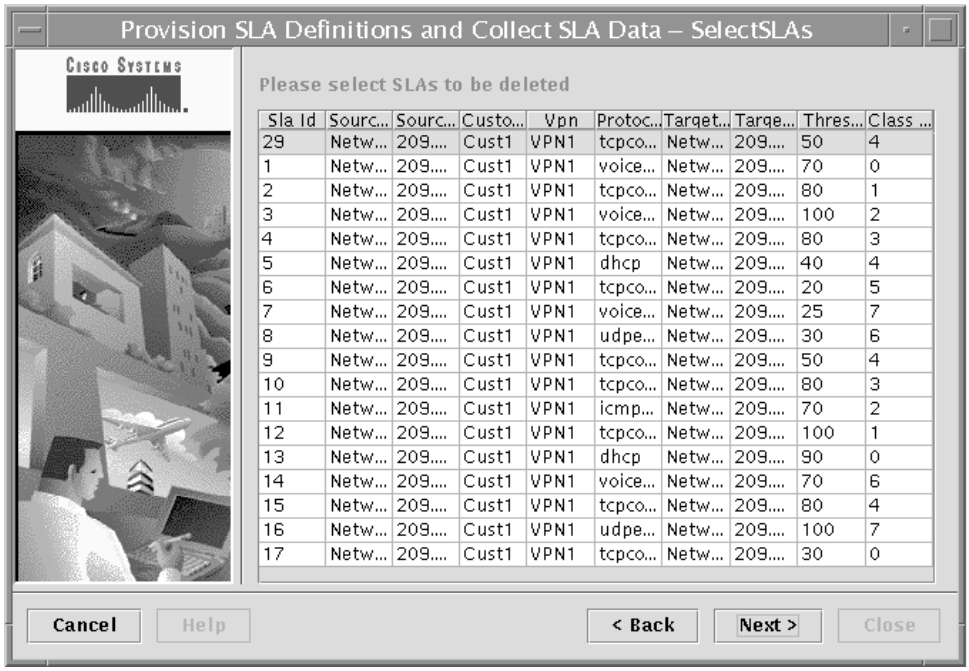

Step 6 Click the appropriate lines in the list to select the SLAs you want to delete, then click **Next**.

To select multiple items, hold down the **Ctrl** key and click each item you want to add.

Step 7 Enter a unique task name, then click **Next**.

Step 8 Choose the default (**Yes**) to proceed to schedule the task, then click **Next**.

Step 9 From the Schedule dialog box, set all the pertinent scheduling information, then click **Add**.

The SLA deletion request is added to the Schedule List (and displayed in the upper pane).

Step 10 Click **Next** twice, then click **Close**.

Enabling Traps for SLA Data

In VPN Solutions Center, you can set traps per SLA either when you create an SLA or on an actively running SLA probe. When you configure traps during SLA creation, the traps are set before the SLA operation activates. In this case, VPNSC sends a trap in the event of a connection loss, a timeout, or threshold violation. When you configure a trap on an SLA probe that is already running, VPNSC does not send a trap after the first operation that triggers the trap until it sends the resolution trap.

An indication as to whether traps are sent on each SLA is recorded in the Repository. When a router reboots, VPNSC recreates the SLA and configures the traps according the data in the Repository.

In VPN Solutions Center (MPLS mode), you can enable or disable traps on three types of network devices:

- Customer Edge devices (CEs)
- Provider Edge devices (PEs)
- Any SA Agent-enabled device in the provider networks

Because the process for enabling traps is almost identical for all three device types, this section describes the procedure for CEs only. The other procedures vary only in the specific devices selected.

To enable traps for SLA data:

Step 1 From the VPN Console, choose **Monitoring > Provision SLA Definitions and Collect SLA Data**.

Step 2 Choose one of the device options:

- **From Customer Edge Devices**
- **From Provider Edge Devices**
- **From Any SA Agent Device**

The introductory screen is displayed. Click **Next**.

The following screen is displayed (see Figure 7-41).

Figure 7-41 Selecting the Enable Traps Option

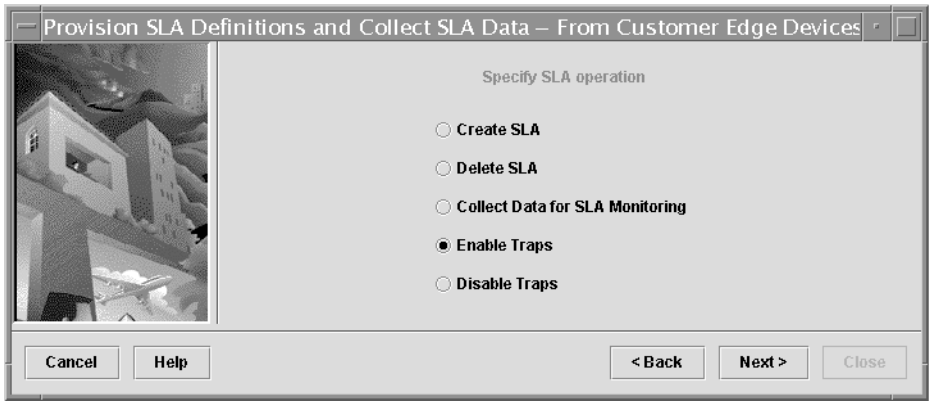

Step 3 Choose **Enable Traps**, then click **Next**.

The following screen is displayed (see Figure 7-42).

Figure 7-42 Selecting the Devices for Enabled Traps

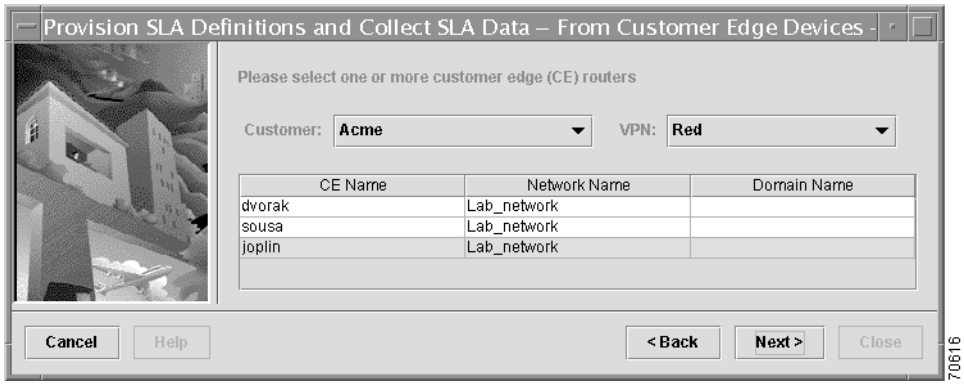

Step 4 Select one or more devices.

a. *Customer*: From the drop-down list, select the pertinent Customer.

b. *VPN*: From the drop-down list, select the pertinent VPN.

c. From the list of devices for the selected Customer and VPN, select one or more devices.

d. Click **Next**.

The following screen is displayed (see Figure 7-43).

Figure 7-43 Specifying the Falling Threshold

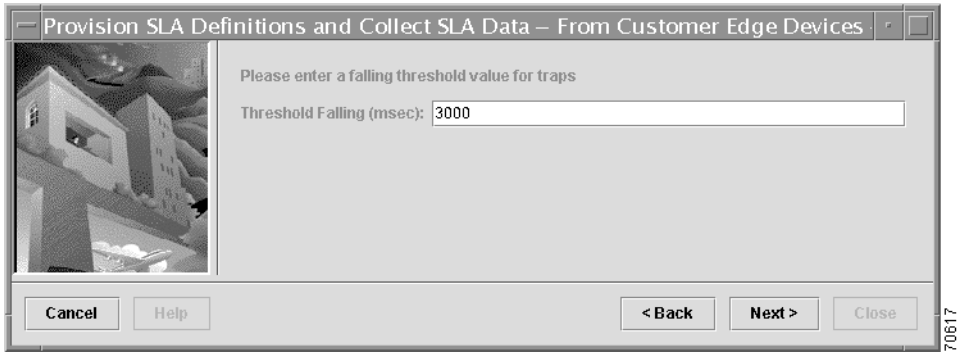

Step 5 Enter the value for the falling threshold in milliseconds, then click **Next**.

If you enable traps for an SLA, you must specify the *Falling Threshold* value, which triggers a threshold resolution trap. The default is 3000 milliseconds.

The following screen is displayed (see Figure 7-44).

Figure 7-44 Selecting the SLA to Enable the Trap

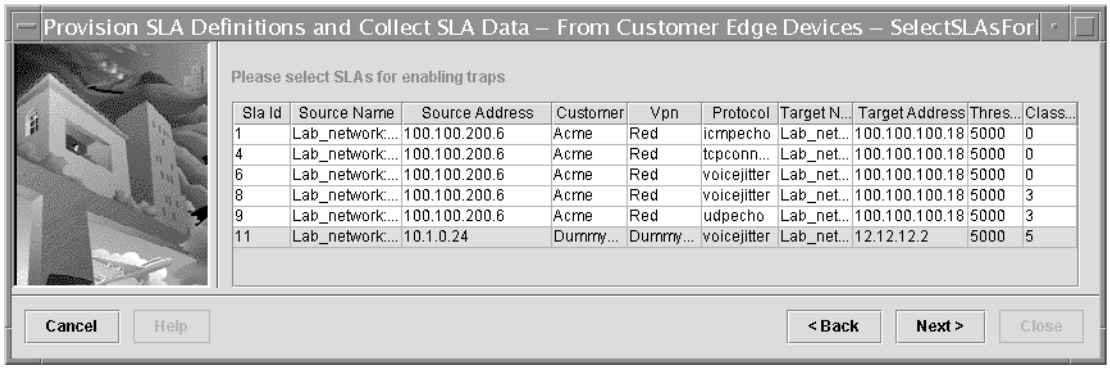

Step 6 Select the SLA that you want to enable traps for, then click **Next**.

The following screen is displayed (see Figure 7-44).

Figure 7-45 Specifying the Trap's Security Level

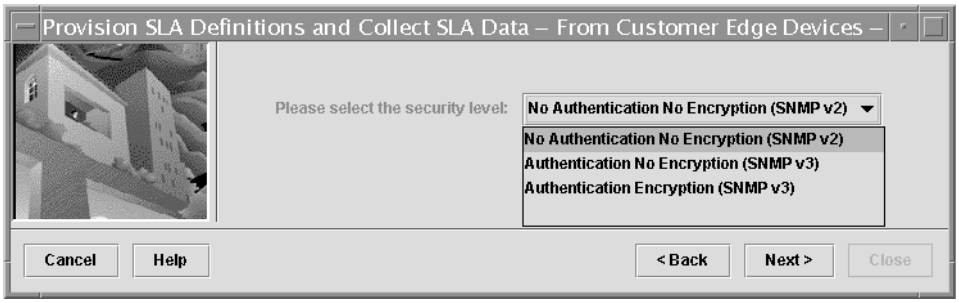

Step 7 From the drop-down list, select the appropriate SNMP security level.

- *No Authentication, No Encryption (SNMPv2)*
- *Authentication, No Encryption (SNMPv3)*
- *Authentication, Encryption (SNMPv3)*

When you have selected the SNMP security level for this trap, click **Next**.

Step 8 Enter a unique task name, then click **Next**.

The next dialog box asks if and when you want to schedule the task.

You have three options:

- *Now*. The task is scheduled to run immediately.
- *Future*. The Schedule dialog box appears.
- *No*. The SLA task is canceled.

Step 9 To run the task now, choose **Now**; to schedule the task, choose **Future**, then click **Next**.

If you choose to schedule the task for some time in the future, the Schedule dialog box appears.

Step 10 Set all the pertinent scheduling information in the Schedule dialog box, then click **Add**.

The task is added to the Schedule List (and displayed in the upper pane).

Step 11 Click **Next** twice, then click **Close**.

Disabling Traps

An indication as to whether traps are sent on each SLA is recorded in the Repository. When a router reboots, VPNSC recreates the SLA and configures the traps according the data in the Repository.

In VPN Solutions Center (MPLS mode), you can disable traps on three types of network devices:

- Customer Edge devices (CEs)
- Provider Edge devices (PEs)
- Any SA Agent-enabled device in the provider networks

Because the process for disabling traps is almost identical for all three device types, this section describes the procedure for CEs only. The other procedures vary only in the specific devices selected.

To disable traps for SLA data:

Step 1 From the VPN Console, choose **Monitoring > Provision SLA Definitions and Collect SLA Data**.

Step 2 Choose one of the device options:

- **From Customer Edge Devices**
- **From Provider Edge Devices**
- **From Any SA Agent Device**

The introductory screen is displayed. Click **Next**.

The next screen lets you select the SLA operation (see Figure 7-46).

Figure 7-46 Selecting the Disable Traps Option

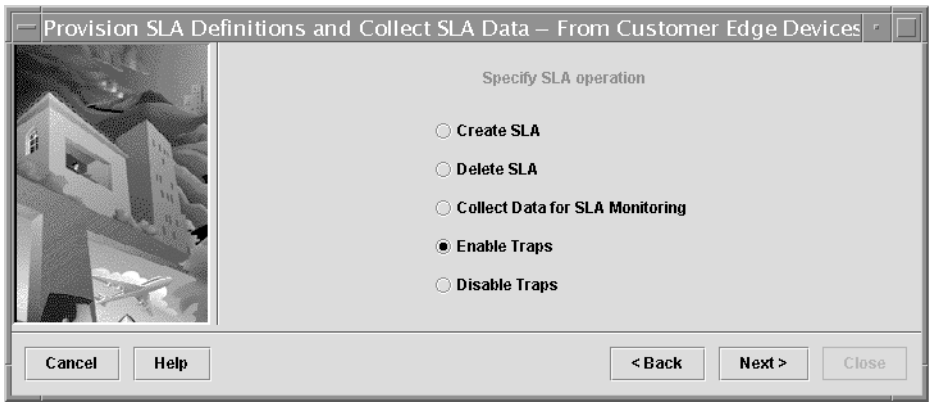

Step 3 Choose **Disable Traps**, then click **Next**.

The following screen is displayed (see Figure 7-47).

Figure 7-47 Selecting the Devices for Disabled Traps

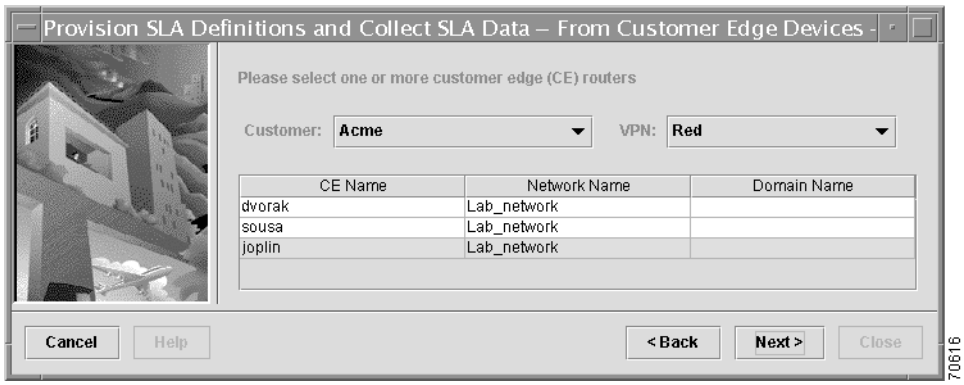

Step 4 Select one or more devices.

a. *Customer*: From the drop-down list, select the pertinent Customer.

b. *VPN*: From the drop-down list, select the pertinent VPN.

c. From the list of devices for the selected Customer and VPN, select the devices where the traps are currently running.

d. Click **Next**.

The next screen lets you select the pertinent SLA (see Figure 7-48).

Figure 7-48 Selecting the SLA Where the Trap is Running

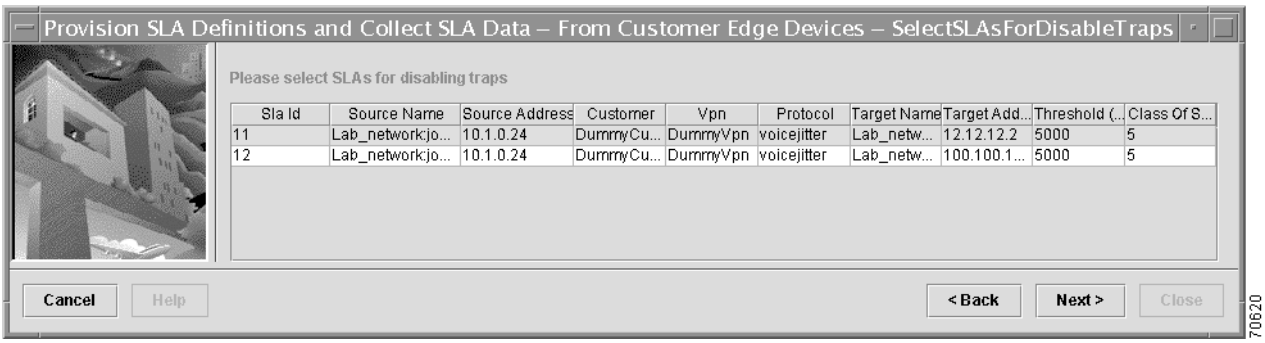

Step 5 Select the SLA that you want to disable traps for, then click **Next**.

The next screen asks you to specify the SLA trap's security level (see Figure 7-49).

Figure 7-49 Specifying the Trap's Security Level

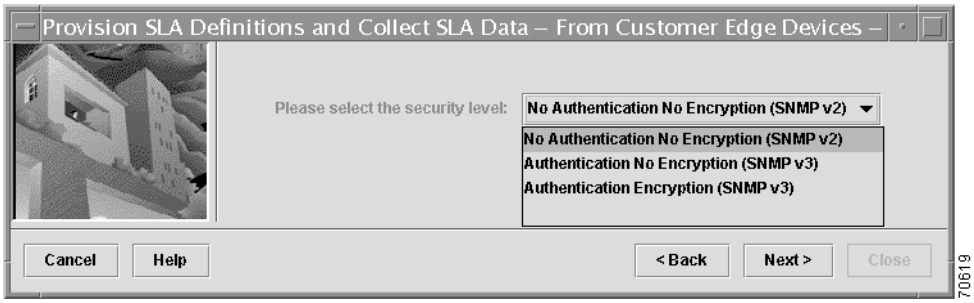

Step 6 From the drop-down list, select the appropriate SNMP security level.

- *No Authentication, No Encryption (SNMPv2)*
- *Authentication, No Encryption (SNMPv3)*
- *Authentication, Encryption (SNMPv3)*

When you have selected the SNMP security level for this trap, click **Next**.

Step 7 Enter a unique task name, then click **Next**.

The next dialog box asks if and when you want to schedule the task.

You have three options:

- *Now*. The task is scheduled to run immediately.
- *Future*. The Schedule dialog box appears.
- *No*. The SLA task is canceled.

Step 8 To run the task now, choose **Now**; to schedule the task, choose **Future**, then click **Next**.

If you choose to schedule the task for some time in the future, the Schedule dialog box appears.

Step 9 Set all the pertinent scheduling information in the Schedule dialog box, then click **Add**.

The task is added to the Schedule List (and displayed in the upper pane).

Step 10 Click **Next** twice, then click **Close**.

Viewing SLA Reports

After collecting SA Agent data for SLA, choose **Monitoring>View SLA Reports**, then select the specific type of report you require.

Note For details on each type of SLA report, refer to "View SLA Reports" in Chapter 9 of the *Cisco VPN Solutions Center: MPLS VPN User Reference*.

The specific report types are as follows:

- Summary Report

  These reports are time-based reports that show the following parameters: *Connectivity* as a percentage, *Maximum Delay* in milliseconds, and *Threshold Violation* as a percentage. These parameters are available in annual, monthly, weekly, daily, and hourly reports. For each parameter, you can generate detailed reports that show more related parameters. The reports can be organized by source router (the source CE of the SLA), SLA identifier, customer name, or VPN name.

- Jitter Report

  Displays statistics that are measured only by Voice Jitter SLAs originated in a selected router. The reports are time-based. They show hourly, daily, weekly, monthly, and annual data and can be organized by SLA ID, destination router, VPN, Customer, or Unspecified.

- HTTP Report

  Displays statistics that are measured only by HTTP SLAs. The reports are time-based, and they show data in the following time increments: hourly, daily, weekly, monthly, and annually. Data can be organized by SLA ID, source router, VPN, or Customer.

  The Summary HTTP Report displays the connectivity, maximum delay, and threshold violation (as in the Summary Report). The Stages HTTP Report displays the round trip time, timeouts, and the error distribution among different HTTP stages: DNS lookup, TCP connect, and Transaction.

- Customer Packet Drop (CE-CE) Report

  Shows the packet drop percentage among CEs of a specific customer. This information is measured only for the SLAs with the jitter protocol. The reports are organized by class of service. The reports are annually, monthly, weekly, daily, and hourly. You can navigate along the time scale.

- Customer Round Trip Delay (CE-CE) Report

  Shows the maximum, minimum, and average round-trip time (in milliseconds) among the CEs of a specific customer. The statistics are for all the probe types. The reports are organized by class of service. The reports are annually, monthly, weekly, daily, and hourly. You can navigate along the time scale.

- Network Packet Drop (PE-PE) Report

  Shows the packet drop percentage among all the shadow SA Agent CEs in the network. The network packet drop between PEs is measured by the shadow SA Agent CEs that are connected to the PEs. This information is measured only for the SLAs with the jitter protocol. The reports are organized by class of service. The reports are annually, monthly, weekly, daily, and hourly. You can navigate along the time scale.

- Network Round Trip Delay (PE-PE) Report

  Shows the maximum, minimum, and average round-trip time among shadow SA Agent CEs in the network. The statistics are for all the probe types. The reports are aggregated by class of service. The reports are annually, monthly, weekly, daily, and hourly. The user can navigate along the time scale.

- SLA Definition Report

  Shows all the SLAs on the SA Agent routers from which data was collected. The SLA Definition report shows the SLA ID given to each SLA. SLAs in the report may have been deleted but are kept in the SLA Definition to match the old collected data.

Querying for SA Agent and Interface Statistics Data

The VPN Solutions Center software periodically collects performance data such as Service Assurance Agent (SA Agent) data. VPNSC then places this data in the Repository. You can access this data through web-based data query tools, as well as through customized reports or through CORBA APIs. The performance data retrieved by the web-based data query tools is saved to a file in XML format.

VPNSC provides the following data query tools:

- SA Agent data
- Interface statistics

To access the VPNSC Data Query Tools, follow these steps:

Step 1 From the VPN Console menu, choose **Monitoring > Run XML Data Query Tool**.

The first time you access the web browser from the VPNSC software, you must log in.

Step 2 In the Netscape Password dialog box, enter the VPN Solutions Center administrative username and password, then click **OK**.

The default administrative username and password is **admin** (for both).

The VPN Solutions Center Data Query Tools page appears (see Figure 7-50).

Figure 7-50 MPLS Solution Data Query Tools Page

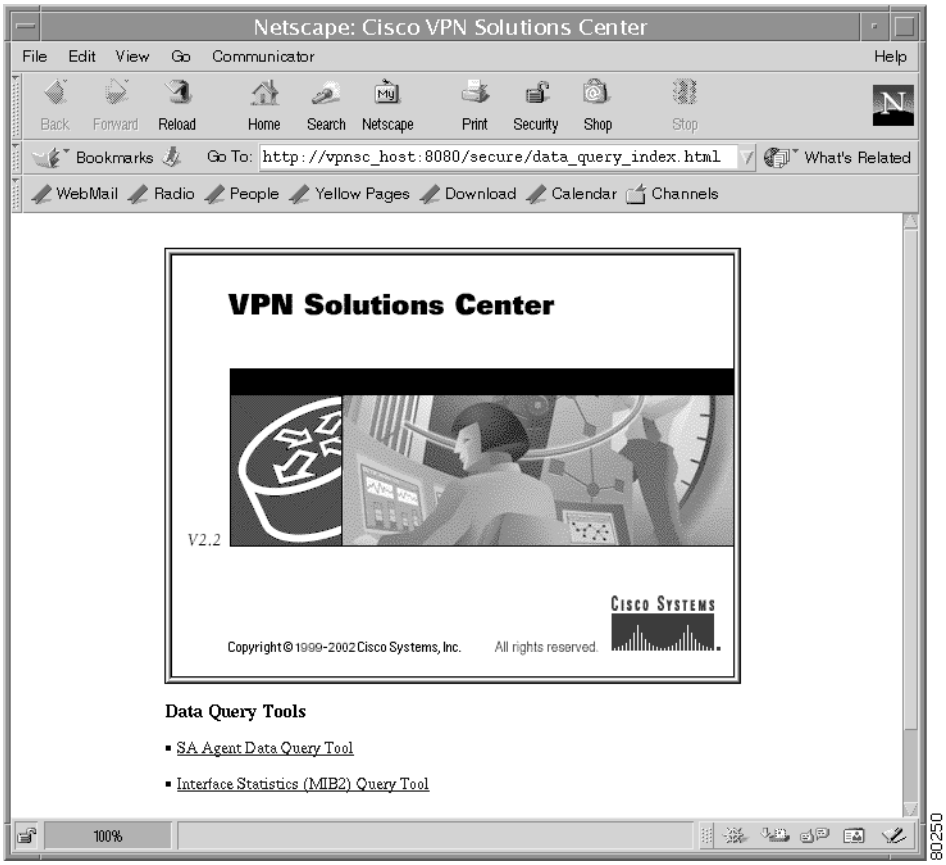

- For information on how to use the Data Query Tools to gather SA Agent data, see the "Retrieving SA Agent Data with the XML Data Query Tool" section on page 7-48.
- For information on how to use the Data Query Tools to gather interface statistics data, see the "Retrieving Interface Statistics with the XML Data Query Tool" section on page 7-58.

For additional details, refer to "XML Data Query Tool" in Chapter 8 of the *Cisco VPN Solutions Center: MPLS Solution User Reference*, Software Release 2.2.

Monitoring Performance Through Service Level Agreements

VPN Solutions Center software monitors performance through the service-level agreement (SLA) server. An SLA defines a service provided by a service provider to any customer. VPN Solutions Center monitors the service related performance criteria by provisioning and monitoring SLAs on routers that support the Service Assurance Agent (SA Agent) management information base (MIB). To provision the SLAs and to collect statistics for each SLA, the process of creating an SLA and collecting the data requires some user input, as described in this section.

The SLA server collects the relevant performance data, stores it persistently, and presents useful reports. The SLA server is based on the Service Assurance Agent (SA Agent) MIB. The MPLS VPN Solution software leverages the SA Agent MIB to monitor SLA performance. Service providers can monitor network traffic using any of the following protocols:

- Internet Control Message Protocol Echo (ICMP Echo)
- Transmission Control Protocol Connect (TCP Connect)
- User Datagram Protocol Echo (UDP Echo).
- Jitter (voice jitter)
- Dynamic Host Configuration Protocol (DHCP)
- Hyper text Transfer Protocol (HTTP)
- Domain Name System (DNS)

About the Service Assurance Agent Feature

The Service Assurance Agent (SA Agent) feature allows you to monitor network performance, network resources, and applications by measuring response times and availability. With this feature you can perform troubleshooting, problem notifications, and preventive analysis based on Service Assurance Agent statistics.

The SA Agent router uses the Cisco Round Trip Time Monitor (RTTMON) MIB. For more information on the RTTMON MIB, refer to the *Cisco MIB User Quick Reference*.

You can use the Service Assurance Agent feature to troubleshoot problems by checking the time delays between devices (such as between two CEs in a VPN) and the time delays on the path from the source device to the destination device at the protocol level.

You can use this feature to perform preventive analysis by scheduling the Service Assurance Agent and collecting the results as history and accumulated statistics. You can then use the statistics to model and predict future network topologies.

About SA Agent Traps

You can configure SA Agent traps per SLA probe. SA Agent can send three types of traps:

- *Connection Loss traps*. VPN Solutions Center sends a Connection Loss trap when an SLA probe detects a lost connection for a connection-oriented protocol. VPNSC sends a resolution trap the next time the operation is completed successfully.
- *Timeout traps*. When an operation delay exceeds the timeout value specified, VPNSC sends a Timeout trap.
- *Threshold traps*. When an operation delay meets a falling threshold value, VPNSC sends a Threshold trap.

The traps configuration encapsulates all three types of SA Agent traps.

In VPN Solutions Center, you can set traps per SLA either when you create an SLA or on an actively running SLA probe. When you configure traps during SLA creation, the traps are set before the SLA operation activates. In this case, VPNSC sends a trap in the event of a connection loss, a timeout, or threshold violation. When you configure a trap on an SLA probe that is already running, VPNSC does not send a trap after the first operation that triggers the trap until it sends the resolution trap.

An indication as to whether traps are sent on each SLA is recorded in the Repository. When a router reboots, VPNSC recreates the SLA and configures the traps according the data in the Repository.

Retrieving SA Agent Data with the XML Data Query Tool

VPN Solutions Center periodically collects Service Assurance Agent (SA Agent) performance data and places this data in the Repository. You can access the SA Agent data through web-based data query tools, as well as through customized reports or through CORBA APIs. The performance data retrieved by the web-based data query tools is saved to a file in XML format.

For related information regarding the retrieval of SLA definitions on SA Agent routers, see the "Retrieving SLA Data with the XML Data Query Tool" section on page 7-51.

To access the SA Agent Data Query Tools:

Step 1 From the VPN Console menu, choose **Monitoring > Run XML Data Query Tool**.

The Netscape browser comes up. The first time you access the web browser from the VPNSC software, you must log in.

Step 2 In the Netscape Password dialog box, enter the username and password for the VPN Solutions Center workstation, then click **OK**.

The VPN Solutions Center Data Query Tools page appears.

Step 3 Choose **SA Agent Data Query Tool**.

The SA Agent Data Query Tool page appears (see Figure 7-51).

Figure 7-51 SA Agent Data Query Tool Page

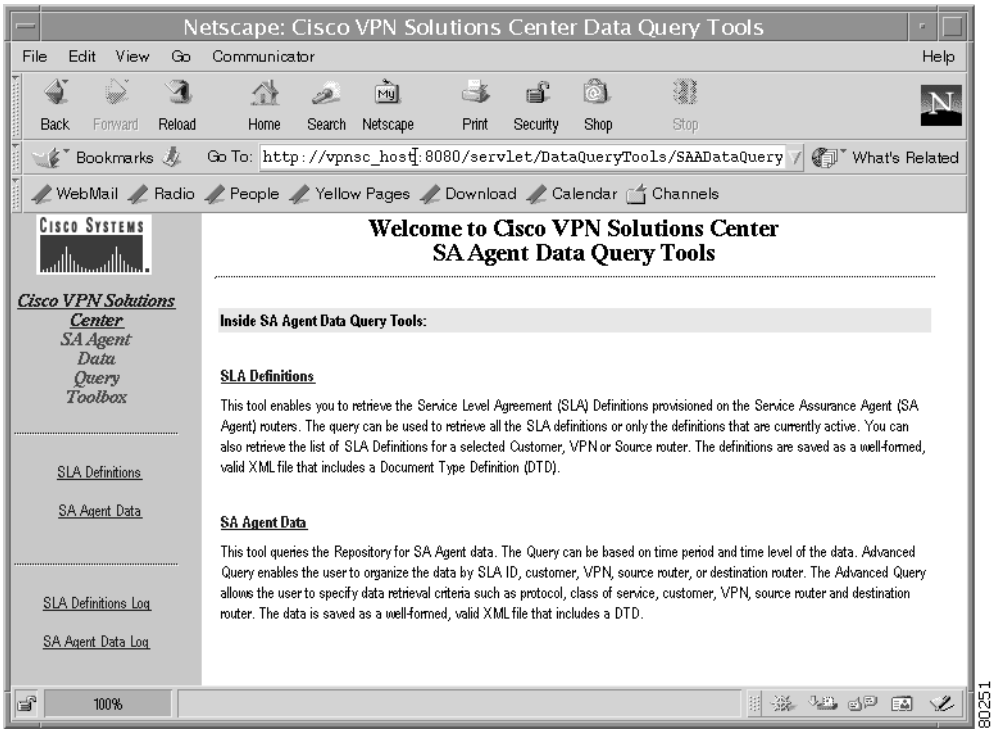

Step 4 Choose **SA Agent Data**.

The SA Agent Data Query page appears (see Figure 7-52).

Figure 7-52 SA Agent Data Query Page

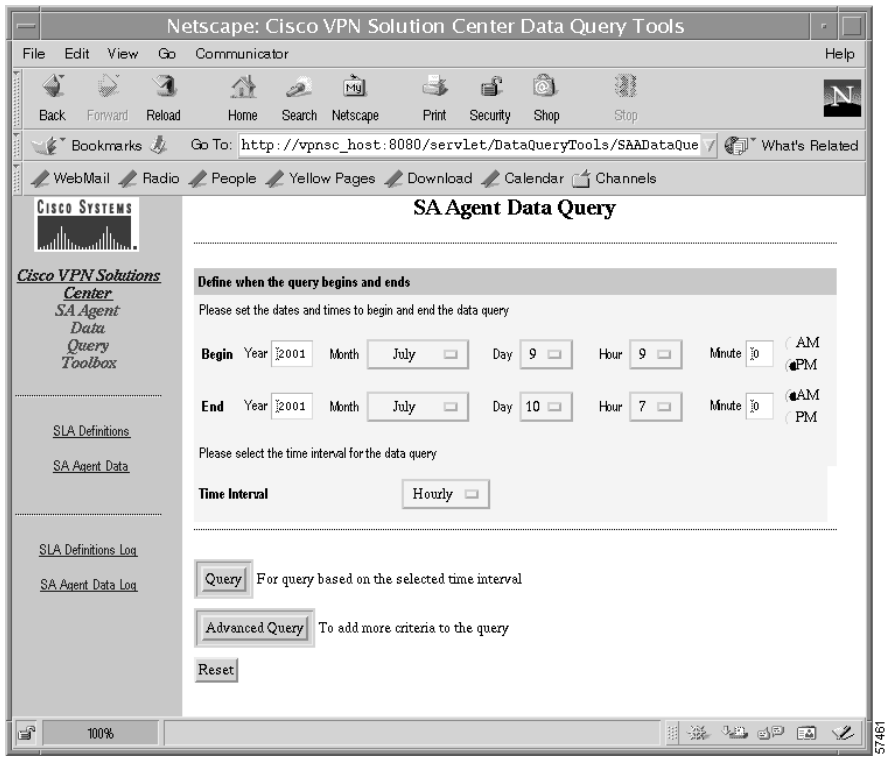

Step 5 In the *Begin* area, set the following parameters:

- **a.** Year to start the SA Agent data query
- **b.** Month to start
- **c.** Day to start
- **d.** Hour to start
- **e.** Minute to start
- **f.** A.M. or P.M.

Step 6 In the *End* area, set the same parameters outlined in Step 5 to indicate when you want the SA Agent data query to end.

Step 7 In the *Time Interval* area, select the appropriate interval for the query: *Hourly*, *Daily*, *Weekly*, *Monthly*, or *Annually*.

You have the option of proceeding with the data query by clicking the **Query** button or adding additional criteria to the data query by clicking the **Advanced Query** button.

Step 8 To initiate the SA Agent query with the current query parameters, click **Query**.

You receive the following message:

SA Agent Data Query is starting; it may take some time. Do you really want to continue?

Step 9 Click **OK** to start the data query.

To cancel the query, click **Cancel**.

If you click **OK**, the another page appears that provides the following options:

- To view the query status, choose the **Query SA Agent Log** link.
- To save the query result to a file, choose the **Query SA Agent Result** link.
- To stop the query process, choose the **Stop SA Agent Query** link.

Step 10 Choose the desired option to proceed.

Retrieving SLA Data with the XML Data Query Tool

VPN Solutions Center allows you to retrieve Service Level Agreement (SLA) definitions on the Service Assurance Agent (SA Agent) routers from which data has been collected. You can either retrieve all the SLA definitions on the SA Agent routers, or only the SLA definitions that are currently active.

You can access the SLA data through web-based data query tools, as well as through customized reports or through CORBA APIs. The data retrieved by the web-based data query tools is saved to a file in XML format that includes a Document Type Definition (DTD).

To access the interface statistics Data Query Tools, follow these steps:

Step 1 From the VPN Console menu, choose **Monitoring > Run XML Data Query Tool**.

The first time you access the web browser from the VPNSC software, you must log in.

Step 2 In the Netscape Password dialog box, enter your user name and password, then click **OK**.

The VPN Solutions Center Data Query Tools page appears.

Step 3 Choose **SA Agent Data Query Tool**.

The SA Agent Query Tools page appears. This page provides two options: **SLA Definitions** and **SA Agent Data**.

Step 4 From this page, choose **SLA Definitions**.

The SLA Definitions Query page appears (see Figure 7-53).

Figure 7-53 SLA Definitions Query Page

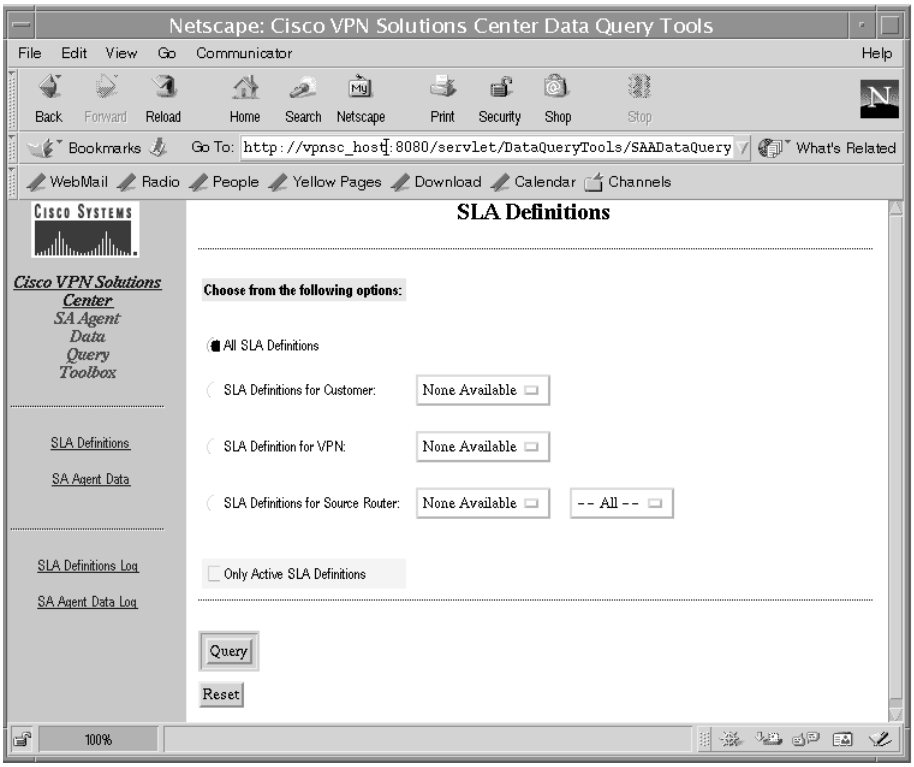

Step 5 Choose one of the following query options:

- **All SLA Definitions**
- **SLA Definitions for Customer**

  From the drop-down list, select the pertinent Customer name.
- **SLA Definitions for VPN**

  From the drop-down list, select the pertinent VPN name.
- **SLA Definitions for Source Router**

  From the drop-down list, select the pertinent router name or select **All**.
- **Only Active SLA Definitions**

Step 6 Click **Query**.

You receive the following message: *SLA Definitions Data Query is starting. Do you really want to continue?*

Step 7 Click **OK** to start the data query.

To cancel the query, click **Cancel**.

The next page that appears gives you the following options:

- To view the query status, choose **SA Agent Query Log**.
- To save the query result to a file, choose **Save Result**.

Step 8 Choose the desired option to proceed.

Viewing Data Reports

Data reports show all the collected data in the Repository. You can use the data reports as a debugging tool to determine whether your data collections were successful, and therefore available for other reports or applications.

You can generate three types of data reports: by *device*, by *network*, and by *dataset type*.

Data Report by Device

To view a data report by device:

Step 1 From the VPN Console menu, choose **Monitoring > View Data Reports > By Device**.

The Target Chooser dialog box appears (see Figure 7-54).

Figure 7-54 Selecting the Device for a Data Report

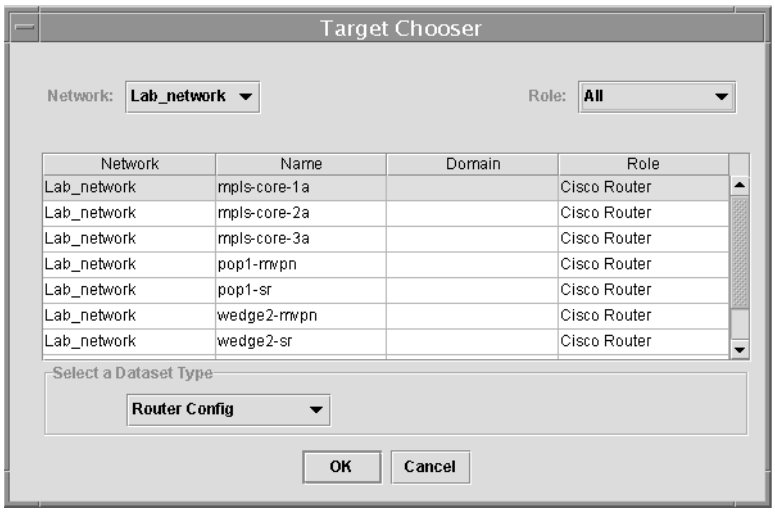

Step 2 In the Target Chooser dialog box, specify the following:

a. *Network*: Choose the network that the desired device is in.

b. *Role*: Specify the device's role.

You can choose **All** to display all the devices in the network, or choose one of the other device roles: **Cisco Router**, **Terminal Server**, **VPN 3000**, or a **PIX** device.

c. *Device*: From the list of displayed devices, select the device you want the data report for.

d. *Dataset Type*: From the drop-down menu, specify the dataset type for this device report.

You can select one of the following dataset types:

You can select one of the following dataset types:

- **Mediator Performance**. The performance of VPNSC itself. This dataset includes memory usage, CPU usage of each process, and so on. The Mediator Performance data is collected by the Watch Dog.
- **Router Configuration**. The Cisco IOS configuration file the selected router.

- **SA Agent Data**. The data collected by SLA probes.
- **VFIT Table**. The VPN Forwarding Information Table. The VFIT table consists of all the VPN routing-related tables on a Cisco router.

Step 3 Click **OK**.

The Data Report by Device is displayed (see (Figure 7-55).

Figure 7-55 Example of a Data Report by Device

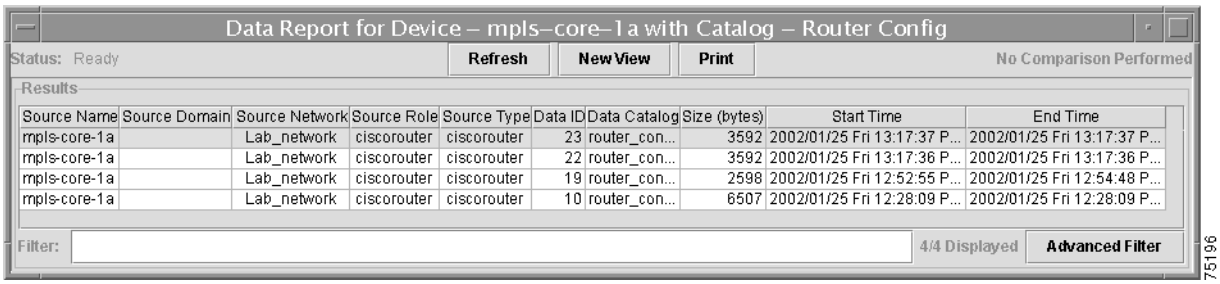

- *Refresh*: The **Refresh** button updates the data report with any new information available from the network for the selected device, network, or dataset type.
- *New View*: The **New View** button gives you a new window with the a copy of the current data report. This is useful when you want to keep the original report displayed while drilling down to other information or going to another area in VPNSC. You can then compare the new data report with the original one.
- *Print*: Click **Print** to print the current data report.
- *Advanced Filter*: The **Advanced Filter** option in the VPN Solutions Center Data Reports allows you to eliminate extraneous information initially displayed in the data reports and find specific items of interest. For details, see the next section, "Using the Advanced Filter for Report Data."

Using the Advanced Filter for Report Data

The Advanced Filter option in the VPN Solutions Center Data Reports allows you to eliminate extraneous information initially displayed in the data reports and find specific items of interest. Once you apply the Advanced Filter parameters, VPNSC redisplays the current data report with the selected information.

To use the filtering feature:

Step 1 From the Data Report window, click **Advanced Filter**.

The following dialog box is displayed (see Figure 7-56)

Figure 7-56 Advanced Filter Dialog Box

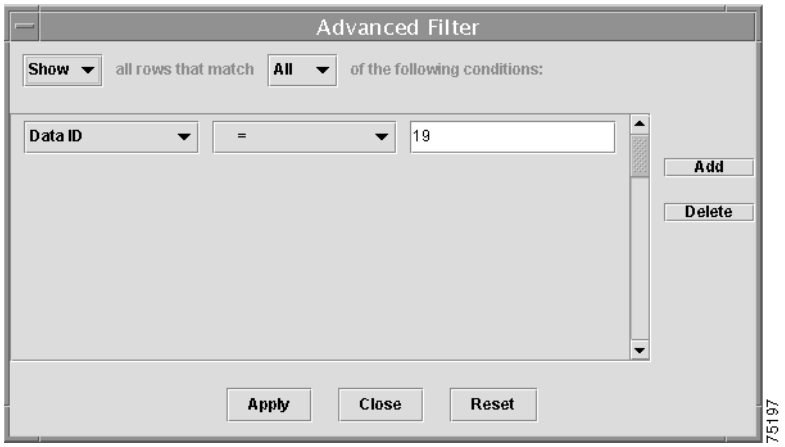

Step 2 Set the Advanced Filter parameters as follows:

The organizing statement for the Advanced Filter is displayed at the top of the dialog box:

Show (or Hide) all rows that match All (or Any) of the following conditions.

a. *Show*: From the Show drop-down menu, choose **Show** to display the rows of filtered information; or choose **Hide** to remove the rows filtered information from the updated data report.

b. *All*: From the All drop-down menu, choose **All** to display the rows that match all of the conditions you set; or choose **Any** to display the rows that match any of the conditions you set.

c. *Choose Column*: From the Choose Column drop-down menu, choose the name of the column that you want to sort the information on:

Source Name	Source Domain	Source Network	Source Role	Source Type
Data ID	Data Catalog	Size (bytes)	Start Time	End Time

d. *Choose Operator*: From the Choose Operator drop-down menu, choose one of these operators:
   - Equals
   - Contains

e. In the data field, enter the value you want to sort the information on.

f. When satisfied with your sort parameters, click **Apply**.

   VPNSC displays the row (or rows) or matching information in an updated data report.

   You can add additional filters by clicking **Add**. Then repeat Steps 2a. through 2f.

g. To close the Advanced Filter, click **Close**.

Data Report by Network

To view a data report by network:

Step 1 From the VPN Console menu, choose **Monitoring > View Data Reports > By Network**.

The Network Chooser dialog box appears (see Figure 7-57).

Figure 7-57 Selecting the Network for a Network Report

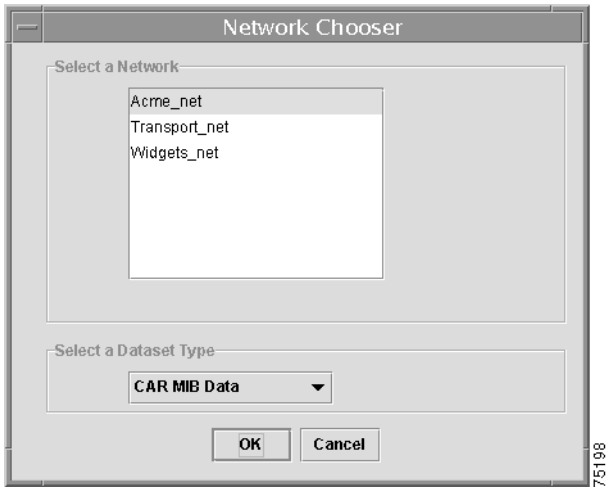

Step 2 In the Network Chooser dialog box, specify the following:

a. *Network*: Choose the network whose data you want to view

b. *Dataset Type*: From the drop-down menu, specify the dataset type for this device report.

You can select one of the following dataset types:

- **Mediator Performance**. The performance of VPNSC itself. This dataset includes memory usage, CPU usage of each process, and so on. The Mediator Performance data is collected by the Watch Dog.
- **Router Configuration**. The Cisco IOS configuration file the selected router.
- **SA Agent Data**. The data collected by SLA probes.
- **VFIT Table**. The VPN Forwarding Information Table. The VFIT table consists of all the VPN routing-related tables on a Cisco router.

Step 3 When satisfied with your selections, click **OK**.

The Data Report by Network is displayed (see (Figure 7-58).

Figure 7-58 Example of a Data Report by Network

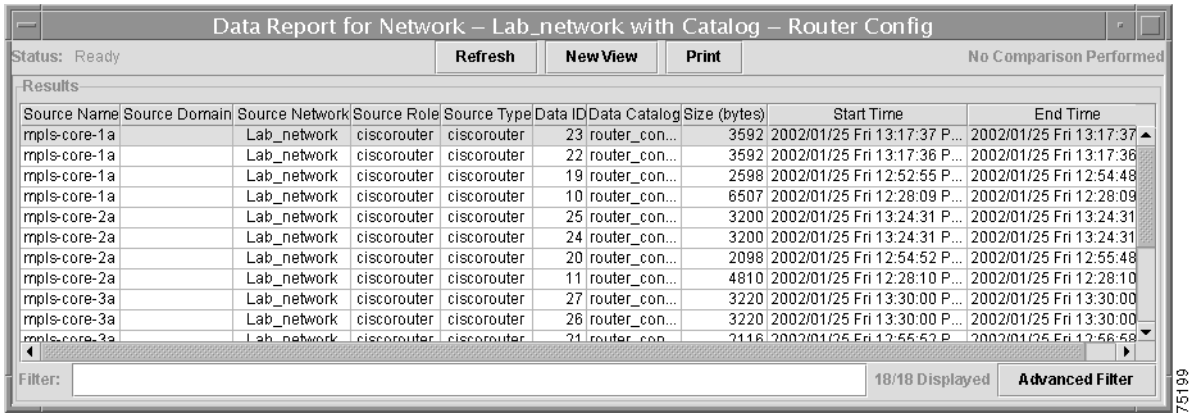

Data Report by Dataset Type

When you display data reports by dataset type, VPNSC presents a report of the collected data in the Repository organized by the specified dataset type.

To view a data report by dataset type:

Step 1 From the VPN Console menu, choose **Monitoring > View Data Reports > By Dataset Type**.

The menu of available dataset types appears.

Step 2 Choose one of the available dataset types:

- **Mediator Performance**. The performance of VPNSC itself. This dataset includes memory usage, CPU usage of each process, and so on. The Mediator Performance data is collected by the Watch Dog.
- **Router Configuration**. The Cisco IOS configuration file the selected router.
- **SA Agent Data**. The data collected by SLA probes.
- **VFIT Table**. The VPN Forwarding Information Table. The VFIT table consists of all the VPN routing-related tables on a Cisco router.

Step 3 When satisfied with your selections, click **OK**.

The Data Report by the specified dataset type is displayed.

Retrieving Interface Statistics with the XML Data Query Tool

VPN Solutions Center periodically collects interface statistics data and places this data in the Repository. You can access the interface statistics data through web-based data query tools. The data retrieved by the web-based data query tools is saved to a file in XML format that includes a Document Type Definition (DTD).

The data query tool collects and saves the interface statistics by router. The statistics include packet counters for router interfaces. You must identify the interfaces by *index number*, which is a unique and constant number, at least from one initialization of the router's network management system to another. The counters are wrapped around numbers with a maximum value of 2 to the power of 32 minus 1.

To access the interface statistics Data Query Tools, follow these steps:

Step 1 From the VPN Console menu, choose **Monitoring > Run XML Data Query Tool**.

The first time you access the web browser from the VPNSC software, you must log in.

Step 2 In the Netscape Password dialog box, enter your user name and password, then click **OK**.

The VPN Solutions Center Data Query Tools page appears.

Step 3 Choose **Interface Stats (MIB2) Query Tool**.

The Interface Stats (MIB2) Query Tool page appears.

Step 4 From this page, choose **Interface Statistics**.

The Interface Statistics Query page appears (see Figure 7-59).

Figure 7-59 Interface Statistics Query Page

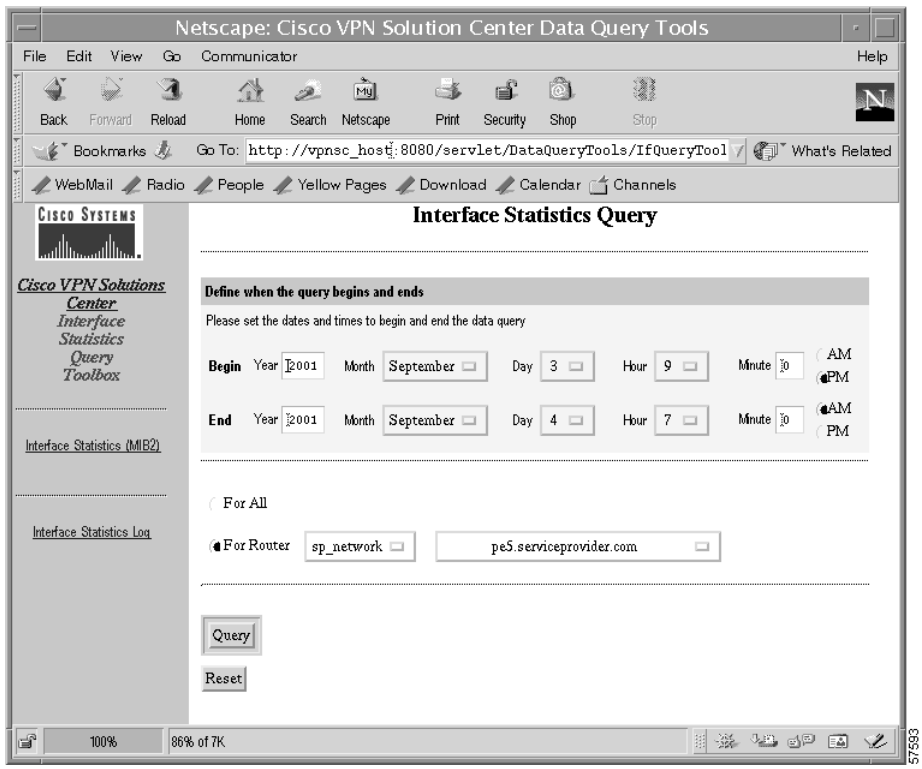

Step 5 In the *Begin* area, set the following parameters:

a. Year to start the Accounting data query

b. Month to start

c. Day to start

d. Hour to start

e. Minute to start

f. A.M. or P.M.

Step 6 In the *End* area, set the same parameters outlined in Step 5 to indicate when you want the Accounting data query to end.

Step 7 In the *Time Interval* area, select the appropriate interval for the query: *Hourly*, *Daily*, *Weekly*, *Monthly*, or *Annually.*

Step 8 You have the option of retrieving interface statistics for all the routers in the network or for a specific router.

- To retrieve interface statistics for all the routers, choose the **For All** radio button.
- To retrieve interface statistics for a specific router, choose the **For Router** radio button, then specify the network and router name.

Step 9 To initiate the interface statistics query, click **Query**.

You receive the following message:

```
Interface Statistics Query is starting; it may take some time. Do you really want to
continue?
```

Step 10 To start the data query, click **OK**.

To cancel the query, click **Cancel**.

If you click **OK**, the next page that appears gives you the following options:

- To view the query status, choose **Interface Statistics Query Log**.
- To save the query result to a file, choose **Interface Statistics Query Result**.
- To stop the query process, choose **Stop Interface Statistics Query**.

Step 11 Choose the desired option to proceed.

The VPNSC Management Network

This chapter provides the fundamental concepts and considerations, as well as our recommendations, for administering customer edge routers (CEs) in the context of the VPN Solutions Center management subnet. Before VPN Solutions Center software can be appropriately deployed to deliver services to customers, the question of whether the CEs are to be managed by the Service Provider or not must be answered. Finally, this chapter describes how to implement the Management VPN technique in the VPN Solutions Center software (see the "Implementing the Management VPN Technique" section on page 8-12).

Unmanaged CE Considerations

One of the options available to the Service Provider is to not manage the CEs connected to the Service Provider network. For the Service Provider, the primary advantage of unmanaged CEs is administrative simplicity.

If the CEs are unmanaged, the provider can use IPv4 connectivity for all management traffic. VPN Solutions Center software is not employed for provisioning or managing unmanaged CEs.

Figure 8-1 shows a basic topology with unmanaged CEs. Note that the network management subnet has a direct link to the Service Provider MPLS core network.

Figure 8-1 Service Provider Network and Unmanaged CEs

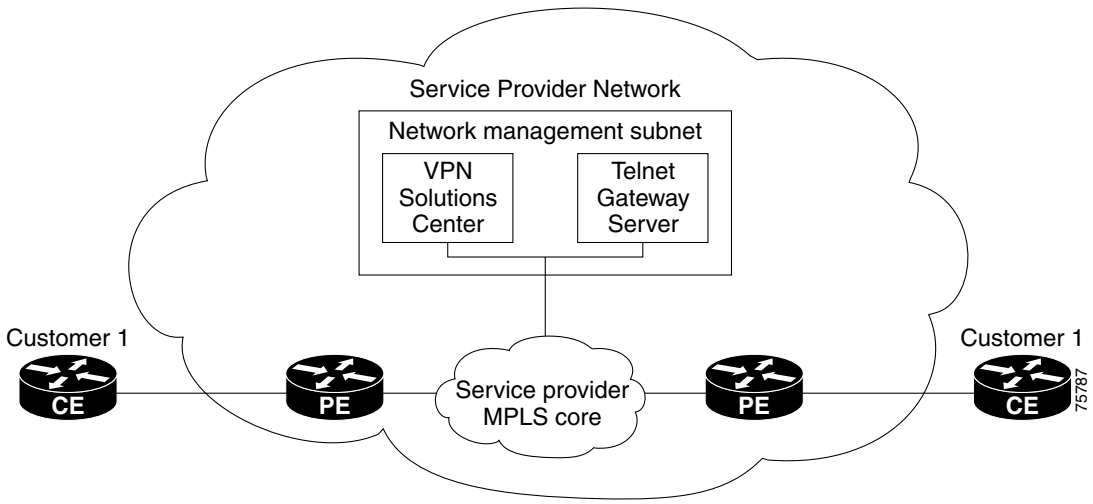

Regarding unmanaged CEs, Service Providers should note the following considerations:

- Because unmanaged CEs are outside the Service Provider's administrative domain, the Service Provider does not maintain or configure unmanaged CEs.
- The Service Provider does *not* administer the following elements on the unmanaged CE:
  - IP addresses
  - Host name
  - Domain Name server
  - Fault management (and timestamp coordination by means of the Network Time Protocol)
  - Collecting, archiving, and restoring CE configurations
  - Access data such as passwords and SNMP strings on the unmanaged CE
- Prototype CE configlets are generated, but they are not automatically downloaded to the router.
- There is no configuration management.
  - With no configuration management, no configuration history is maintained and there is no configuration change management.
  - Changes to a service request (on the PE-CE link) are not deployed to the CE.
- There is no configuration auditing because there is no means to retrieve the current CE configuration.
- You can perform routing auditing.
- You can use the Service Assurance Agent (SA Agent) to measure response times between shadow routers, but you *cannot* use SA Agent to measure response times between CEs.

Managed CE Considerations

The alternative to unmanaged CEs is managed CEs, that is, customer edge routers managed by the Service Provider. Managed CEs can be wholly within the Service Provider's administrative domain or co-managed between the provider and the customer, although CE co-management poses a number of ongoing administrative challenges and is not recommended.

For information on how you define a CE as a managed CE, refer to the "Specifying the Management Status for CE Routers" section on page 4-48.

Regarding managed CEs, Service Providers should note the following considerations:

- Managed CEs are within the Service Provider's administrative domain. Thus, some connectivity to the CEs from the Service Provider network is required.
- The Service Provider must administer the following elements on the managed CE:
  - IP addresses
  - Host name
  - Domain Name server
  - Access data such as passwords and SNMP strings
- The Service Provider should administer fault management (and timestamp coordination by means of the Network Time Protocol)
- The Service Provider can administer collecting, archiving, and restoring CE configurations.

- CE configlets are generated and downloaded to the managed CE.
- Changes to service requests are based on the current CE configuration and automatically downloaded.
- The CE configurations are audited.
- Customer routing and Service Provider routing must interact.
- Access from CEs to the management hosts on the network management subnet is required.
- Configuration auditing and routing auditing are both functional.
- You can use the Service Assurance Agent (SA Agent) to measure response times between CEs and between shadow routers.

The following sections discuss the concepts and issues required for administering a managed CE environment.

What Is the Network Management Subnet?

The *VPN Solutions Center Network Management Subnet* consists of the VPN Solutions Center workstation and one or more Telnet Gateway Server workstations connected on a LAN or subnet.

The Network Management Subnet is required when the provider's service offering entails the management of CEs. Once a CE is in a VPN, it is no longer accessible by means of conventional IPv4 routing unless one of the techniques described in this chapter is employed.

Figure 8-2 shows the VPN Solutions Center network management subnet and the devices that may be required to connect to it:

Figure 8-2 The VPN Solutions Center Network Management Subnet

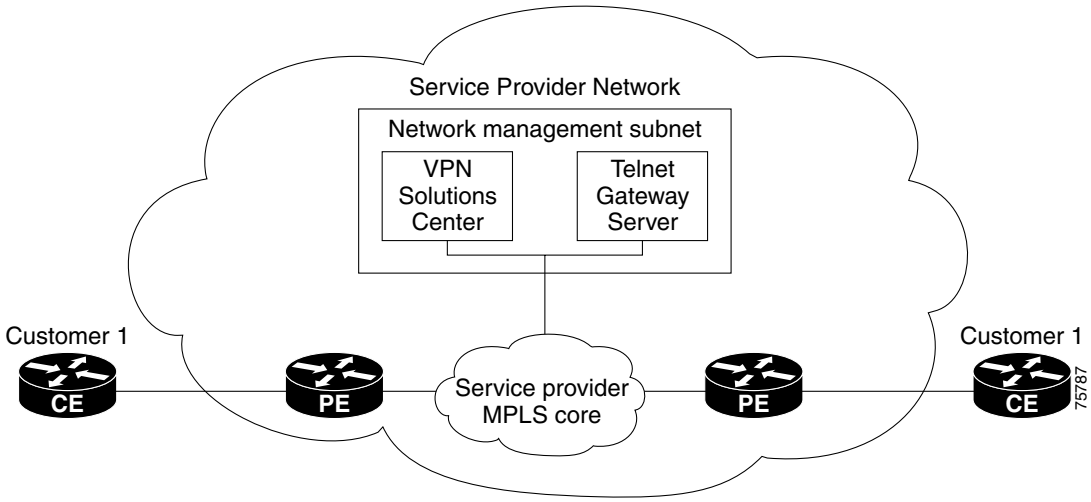

Issues Regarding Access to VPNs

The core issues with regard to gaining access to VPNs are as follows:

- How to keep provider space “clean” from unnecessary customer routes
- How to keep customer space “clean” from both the provider’s and other customer’s routes
- How to provide effective security
- How to prevent routing loops

  VPN Solutions Center does not handle any of these responsibilities—doing so must be designed and implemented by the Service Provider.
- Reachability changes as a direct consequence of employing VPN Solutions Center.

  Before you provision a CE in the VPN Solutions Center software, you might be able to reach the CE via IPv4 connectivity, but the moment the product deploys a service request, you cannot reach that CE any more—unless you have *first* implemented the network management subnet.

The Network Management Subnet Implementation Techniques

The network management subnet must have access to Management CEs (MCEs) and PEs.

The network management subnet is appropriate—and necessary—when there is an intent to have managed CEs connected via an in-band connection. *In-band* indicates a single link or permanent virtual circuit (PVC) that carries *both* the customer's VPN traffic, as well as the provider’s network management traffic.

Management CE (MCE)

The network management subnet is connected to the Management CE (MCE). The MCE *emulates* the role of a customer edge router (CE), but the MCE is in provider space and serves as a network operations center gateway router. The MCE is part of a management site as defined in the VPN Solutions Center software.

You configure the MCE by identifying the CE as part of the management LAN in the VPNSC software. For details on how to define a CE as an MCE within VPN Solutions Center software, see the “Implementing the Management VPN Technique” section on page 8-12.

Management PE (MPE)

The Management PE (MPE) *emulates* the role of a PE in the provider core network. The MPE connects the MCE to the provider core network. An MPE can have a dual role as both a PE and the MPE.

The MPE needs access to the following devices:

Device	Connectivity	Function
1. Customer Edge Routers (CEs)	Access from the network management subnet into the VPNs	Provision or change configuration and collect SA Agent performance data
2. Shadow CEs	Access from the network management subnet into the VPNs	A simulated CE used to measure data travel time between two devices. A shadow CE is connected directly to a PE via Ethernet.
3. Provider Edge Routers (PEs)	Standard IP connectivity	Provision or change configuration

At the current time, VPN Solutions Center recommends three main network management subnet implementation techniques:

- *Management VPN Technique*

  The MPE-MCE link uses a Management VPN (see the "Management VPN Technique" section on page 8-6) to connect to managed CEs. To connect to the PEs, the MPE-MCE link uses a parallel IPv4 link.

- *Extranet Multiple VPN Technique*

  The MPE-MCE link uses the Extranet Multiple VPN *technique* (see the "Extranet Multiple VPN Technique" section on page 8-8) to connect to managed CEs. To connect to the PEs, the MPE-MCE link uses a parallel IPv4 link. See also the "If the VPN Has CEs in Other VPNs (Extranets)" section on page 6-12.

- *Out-of-Band Technique*

  In the Out-of-Band technique, the MCE has IPv4 connectivity (that is, not MPLS VPN connectivity) to all the CEs and PEs in the network (see the "Out-of-Band Technique" section on page 8-9). In this context, *out-of-band* signifies a separate link between PEs that carries the provider's management traffic.

The network management subnet technique the provider chooses to implement depends on many factors, which are discussed later in this chapter.

Management VPN Technique

The Management VPN technique is the default method provisioned by VPN Solutions Center. A key concept for this implementation technique is that *all the CEs in the network are a member of the management VPN*. To connect to the PEs, the MPE-MCE link uses a parallel IPv4 link. Figure 8-3 shows a typical topology for the Management VPN technique.

Figure 8-3 Typical Topology for a Management VPN Network

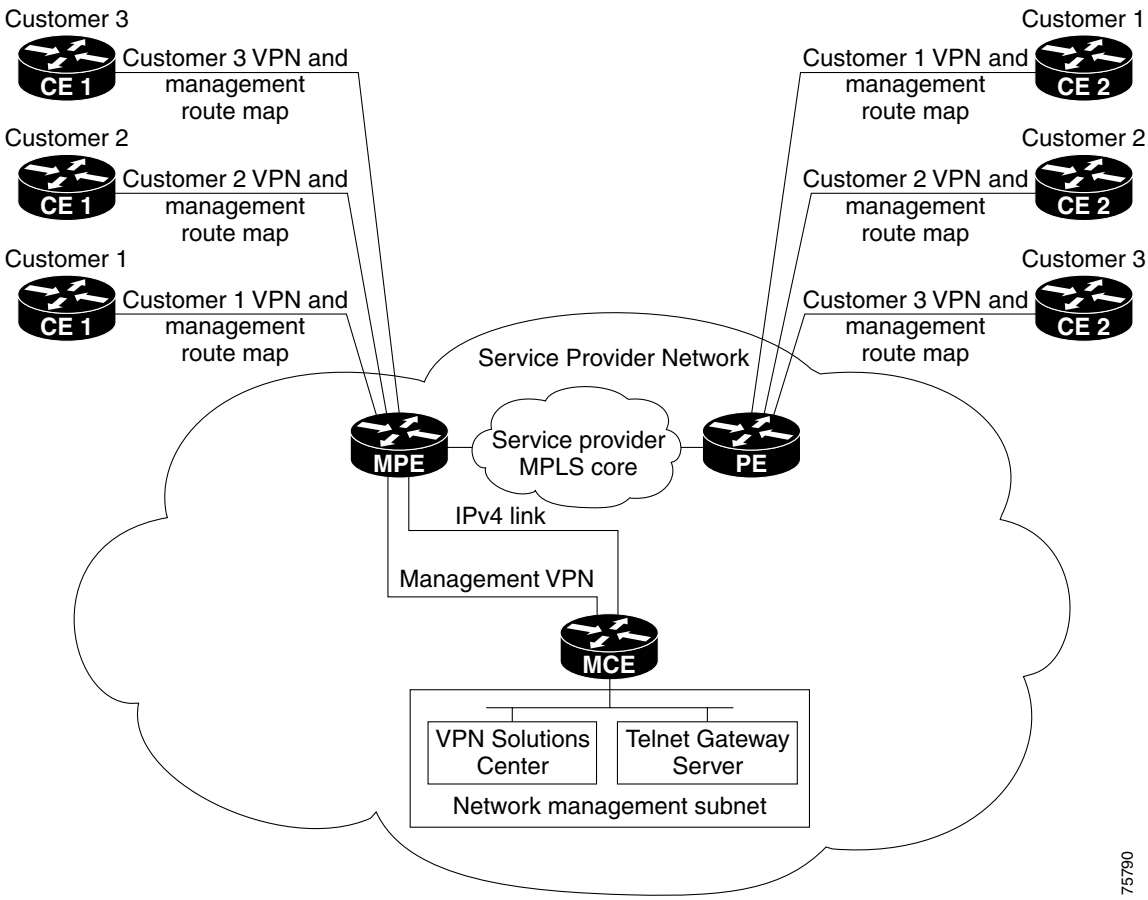

When employing the Management VPN technique, the MPE-MCE link uses a *management VPN* to connect to managed CEs. To connect to the PEs, the MPE-MCE link employs a parallel IPv4 link.

Each CE in a customer VPN is also added to the management VPN by selecting the **Join the management VPN** option in the service request user interface.

The function of the management route map is to allow only the routes to the specific CE into the management VPN. The Cisco IOS supports only one export route map and one import route map per VRF.

As shown in Figure 8-3, a second parallel non-MPLS VPN link is required between the MPE and MCE to reach the PEs.

For information on how to provision a Management VPN in VPN Solutions Center software, see the "Implementing the Management VPN Technique" section on page 8-12.

Note Implementation of the Management VPN technique requires Cisco IOS 12.07 or higher.

Advantages

The advantages involved in implementing the Management VPN technique are as follows:

- Provisioning with this method requires only one service request.
- The only routes given to the network management subnet are the routes to the CEs—that is, either the address of the CE link to the PE or the CE loopback address. General VPN routes are *not* given to the network management subnet.
- A CE in the Management VPN method is a spoke to the Management VPN regardless of which role the CE has within its own VPN. Therefore, CEs cannot be accidentally exposed to inappropriate routes. The only management routes the CEs can learn must come from a hub of the Management VPN.

Extranet Multiple VPN Technique

A key concept for this network management subnet technique is that *the MPE-MCE pair are part of all the customer's VPNs*. When you a add new VPN to the Extranet Multiple VPN, you must create a service request each time a VPN is defined to add that VPN to the MPE-MCE pair, and thus to the network management subnet. To connect to the PEs, the MPE-MCE link uses a parallel IPv4 link. Figure 8-4 shows a typical topology for the Extranet Multiple VPN.

Figure 8-4 Extranet Multiple VPN

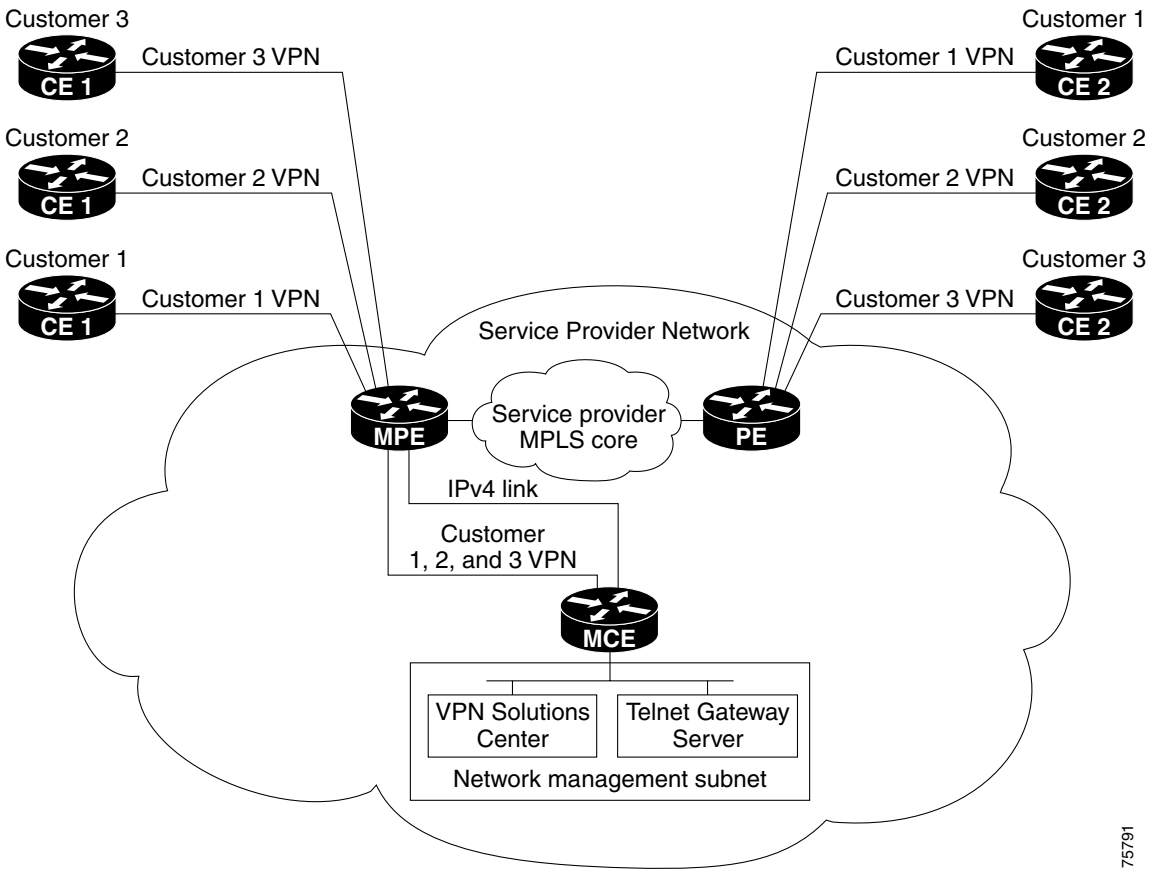

In the Extranet Multiple VPN (sometimes referred to as the *rainbow VPN*), several security and access list considerations exist, but these considerations are centralized at the MPE and MCE devices.

The MPE includes the BGP routes to *all* customer routes. This should be constrained such that only the CE subnet routes are imported to the interior gateway protocol.

Advantages

The advantages in implementing the Extranet Multiple VPN technique are as follows:

- Only the MPE has routes for all the VPNs—the PEs do not have the VPN routes; that is, all the customer routes are only in the MPE's VRF route tables.
- You must apply access lists on the MPE-MCE link only.
- It is easy to create another MPE-MCE pair if necessary.

Out-of-Band Technique

The Out-of-Band technique does not employ a management VPN to manage the CEs. Out-of-band connectivity is provided by IPv4 links. *Out-of-band* signifies a separate link between PEs that carries the provider's management traffic. As shown in Figure 8-5, the MCE provides separation between the provider's routes and the customer's routes.

Figure 8-5 Out-of-Band Technique

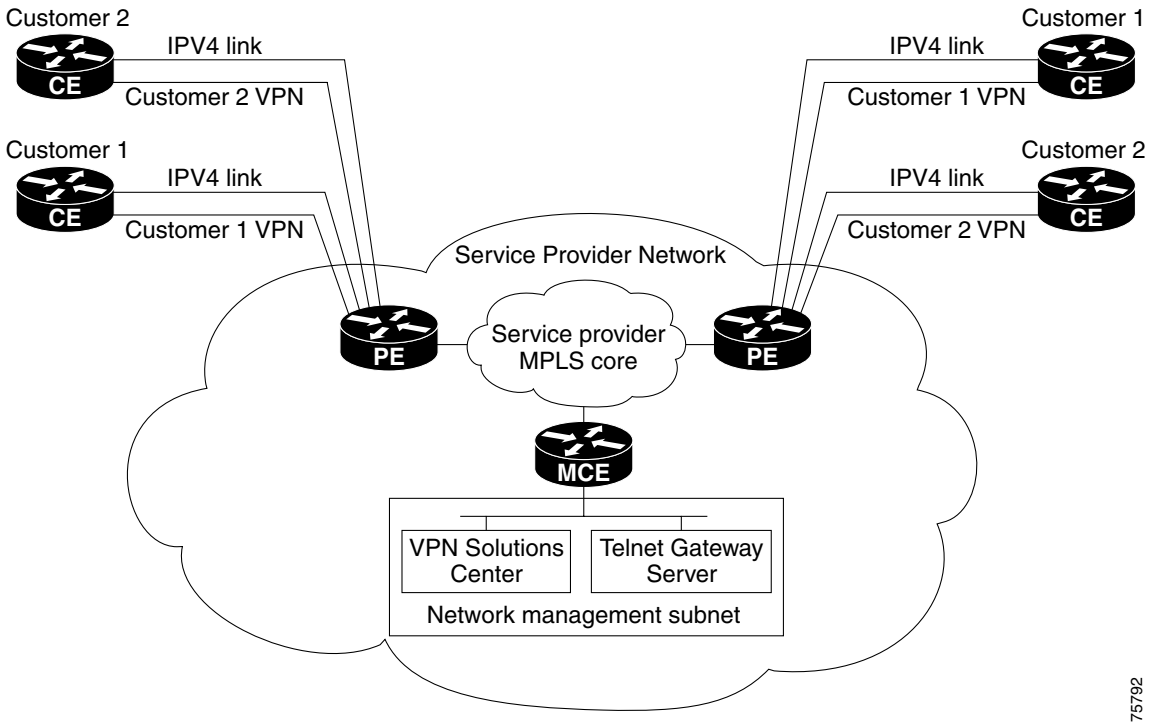

The Out-of-Band technique has the advantage of being relatively simple to set up, and no management VPN is required. However, its disadvantages are that it is expensive since it requires an IPv4 connection to each CE. Also, due to the delicate staging requirements for this technique, the Out-of-Band implementation does have a high degree of complexity.

Securing the Management Network

If you use VPN Solutions Center for IP allocation, you know the set of legal IP addresses for management access to CEs. Therefore, you can deny all packets that do not originate from a VPN Solutions Center IP address pool. (If the network employs non-auto-picked IP addressing, augment the notion of "pool" to cover all legal IP addresses for a CE interface to a PE.) You can also limit access precisely to those hosts on the network management subnet that need it.

The CE access lists between IP pool addresses and the network management subnet hosts should also specify the required ports for access (using Telnet). It is important to limit the port numbers. Those three ports are the only permissible ones. Access to the Orbix process running on the VPN Solutions Center host should particularly be denied.

Cisco recommends the following access rules of type:

- Permit from {VPNSC host} to anything in the "pool."
- Deny all others.

Apply these access rules outbound on the CE on its interface up to the PE so that only VPN Solutions Center can send packets, and then only to management addresses. Additional rules of type are as follows:

- Permit from "pool" to {VPNSC host} for TCP established.
- Deny all others.

These rules should apply on the CE as an input list on its link from the PE. Thus, only responses are allowed in—general CEs cannot start a session to the management machines—and then only from legal IP addresses.

Given these rules of type, only the CE can send packets into the network management subnet, and even those must be in response to a network management subnet query. Spoofing could be an issue, but for that Cisco recommends anti-spoofing access lists as part of the basic configuration of CEs at customer sites—deny all packets coming from within a site marked with a management address.) The CEs do not need the CE-PE link when returning management packets.

Another option is to suppress the network management subnet; that is, you can set up static addresses with /32 subnet masks on the PE—one for each host on the network management subnet needing to receive packets from CEs. At a minimum, that would be the VPN Solutions Center workstation. No other routes need to be allowed into the VRF supporting the network management subnet.

Build the local entries in the VRF like this:

```
IP route VRF Management VPNSC_host/32 CE_address
IP route VRF Management CIPM_host/32 CE_address
```

The term "VRF Management" is for illustration purposes only; VPN Solutions Center builds all this and picks a name for the VRF.

To prevent injections of inappropriate routes, it is helpful to add this command:

```
IP route VRF Management 0.0.0.0/0 Null0
```

That is all you want to put in the local VRF table. From there, it dynamically learns all the routes to the other CEs.

However, you cannot prevent it also knowing a directly connected route for the link between this PE and the Management CE. You must protect against customer attempts to gain access to the Management CE. The access lists described above control only *transit* traffic across that CE.

Therefore, Cisco recommends that the PE have an access list applied outbound on the link to the CE in the following form:

```
permit packets to {VPNSC Host}
deny everything else
```

This is simple and it prevents customers from gaining access to the Management CE.

Cisco recommends the following:

1. In the PE configuration file, enter the following commands:

   ip route vrf management *VPNSC_host* **ip/32 <mce**

   ip route vrf management *CIPM_host* **ip/32 <mce**

   To dump unknowns, add this command:

   ip route vrf management 0.0.0.0/0 Null0

2. Whenever possible, use statics on the Management CE too—use a static or set of statics covering legal management addresses, as discussed above.

   If dynamic routing is absolutely required (meaning it is not known which addresses might be used for CE-PE links), then you can use RIP. However, Cisco recommends doing so one way only: redistribute BGP into RIP on the PE, but do not redistribute back. VPN Solutions Center makes two-way redistributions in such cases, so add the RIP configuration manually when setting this up. Route maps could apply here, but as noted, running dynamic routing is generally undesirable.

3. The most important access lists are output and input lists on the Management CE.

   The output access list: On the Management CE, connected to Link B (with access to the VPNs), make an output access list as follows:

   ```
   permit {VPNSC host} to <pool>
   deny all
   ```

   The input access list is as follows:

   ```
   permit <pool to {VPNSC host} with tcp-established
   deny all
   ```

4. To protect the Management CE, create an output access list on the PE's link B interface:

   ```
   permit to {VPNSC host}
   deny all
   ```

5. If desired, also place an access list to protect the IPv4 link, depending on the Service Provider's own access needs to the network management subnet.

Implementing the Management VPN Technique

The Management VPN technique is the default method provisioned by VPN Solutions Center. A key concept for this implementation technique is that *all the CEs in the network are a member of the management VPN*. The Management VPN is a VPN that belongs to the service provider so that the service provider can manage the VPNs that belong to the provider's customers. Figure 8-6 shows a typical topology for the Management VPN technique.

Figure 8-6 Example of Management VPN Topology

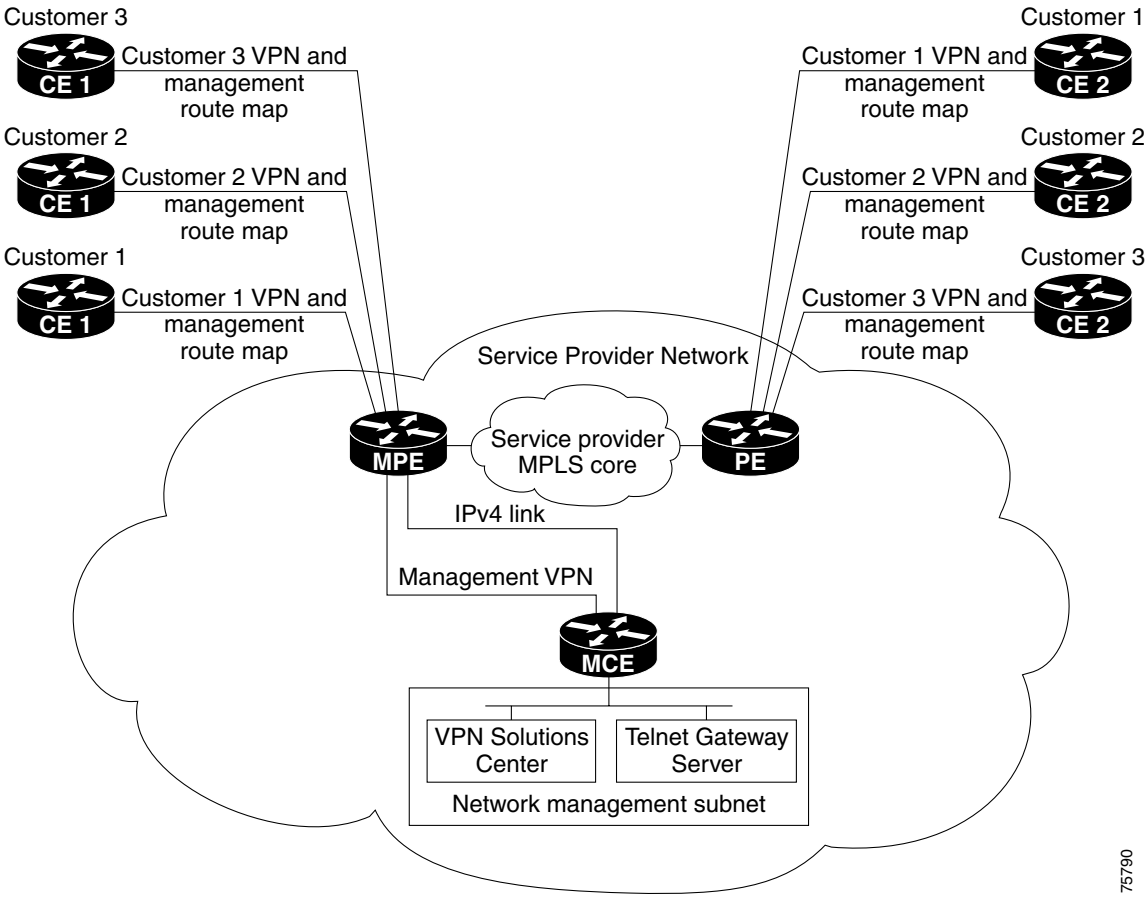

A Management VPN employs two devices called the *Management CE (MCE)* and the *Management PE (MPE).*

- The network management subnet is connected to the Management CE (MCE). The MCE *emulates* the role of a customer edge router (CE), but the MCE is a router in provider space that serves as a network operations center gateway router. The MCE is part of a management site as defined in the VPN Solutions Center software.
- The Management PE (MPE) is a router in service provider space that *emulates* the role of a PE in the provider core network. The MPE connects the MCE to the provider core network. An MPE can have a dual role as both a standard PE and the MPE.

The MPE needs access to the following devices:

Device	Connectivity	Function
1. Customer Edge Routers (CEs)	Access from the network management subnet into the VPNs	Provision or change configuration and collect SA Agent performance data
2. Shadow routers	Access from the network management subnet into the VPNs	A simulated CE used to measure data travel time between two devices
3. Provider Edge Routers (PEs)	Standard IP connectivity	Provision or change configuration

The MPE-MCE link uses a Management VPN (see the "Management VPN Technique" section on page 8-6) to connect to managed CEs. To connect to the PEs, the MPE-MCE link uses a parallel IPv4 link.

Provisioning a Management VPN

The procedure to provision a management VPN assumes that routers that are to function as the MPE and MCE already exist in the service provider network.

The first step is to create a VPN Customer specifically reserved as the Management VPN Customer. The Management VPN Customer should have a single site with a single CE—the router designated as the Management CE—assigned to the Management VPN Customer's site.

Note Prior versions of VPN Solutions Center used numbered access list entries. For versions 2.0 and after, the product employs named access list entries. When you redeploy existing service requests in MPLS VPN Solution 2.x, you will observe that each numbered access list entry automatically converts to a named access list entry. No action is required on the part of the service provider to effect the transition to named access list entries.

Defining a Management VPN in VPNSC Software

To define a management VPN in VPN Solutions Center software, follow these steps:

Step 1 From the VPN Console menu, choose **Setup > New VPN Customer**.

You can also **right-click** the VPN Customers folder and choose **New VPN Customer**.

The New VPN Customer dialog box appears (see Figure 8-7).

Figure 8-7 Creating the Management VPN Customer

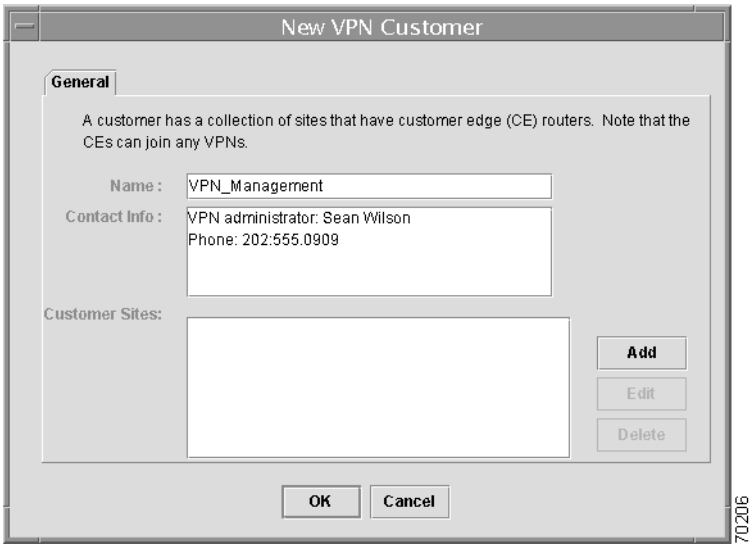

Step 2 Enter the name of the Management VPN Customer.

Remember that the Customer in this case is the Service Provider.

Step 3 Optionally, enter the contact information for the Service Provider network administrator.

Though it is not required, entering the contact information is recommended.

Step 4 To define the site for the Management VPN, click **Add**.

The Add Customer Site dialog box appears (see Figure 8-8).

Figure 8-8 Adding the Management VPN Customer Site

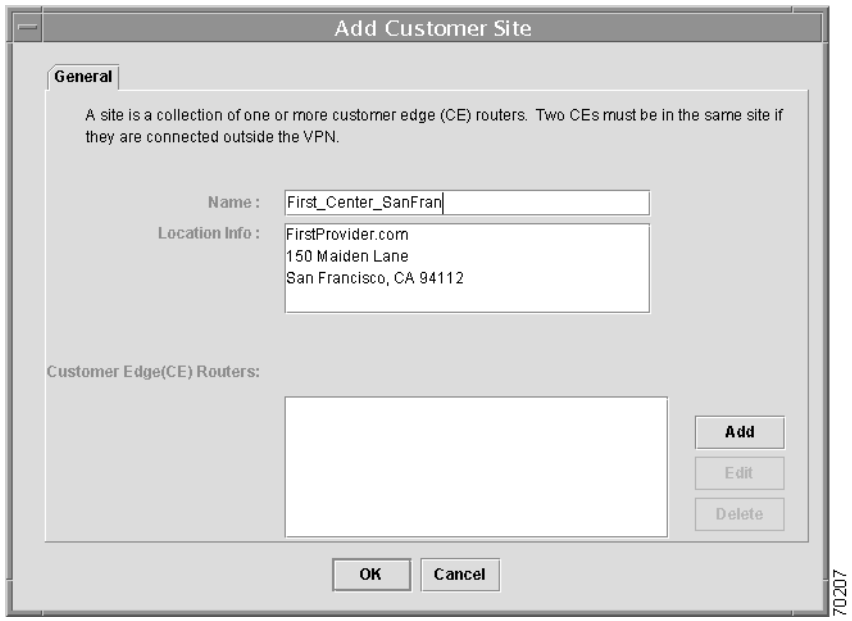

a. *Name*: Enter the management site's name.

b. *Location Info*: Enter the location information.

Step 5 To add the Management CE to the management site, click **Add**.

The Add Customer Edge Routers dialog box appears (see Figure 8-9).

Figure 8-9 Adding the MCE to the Management Site

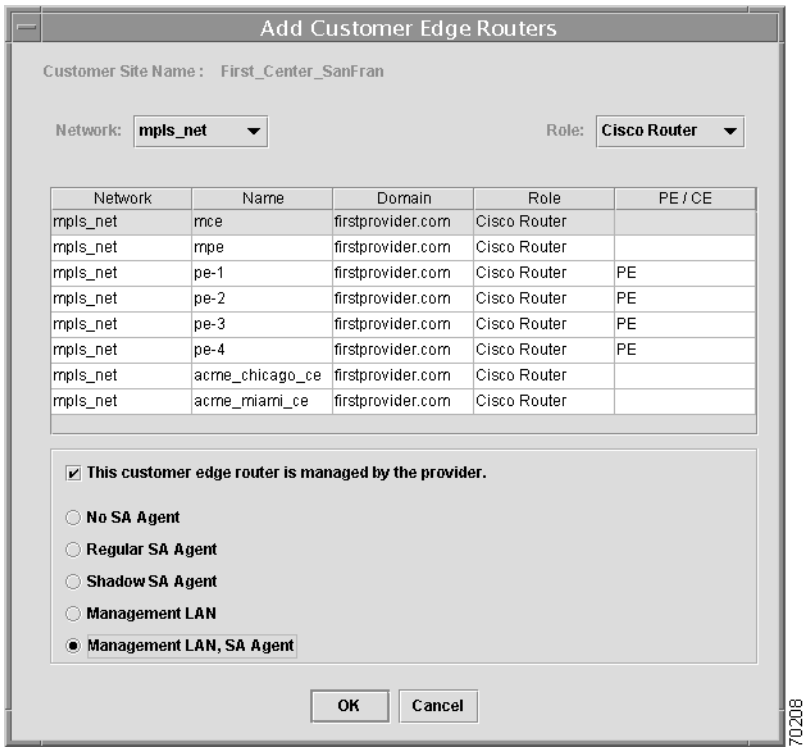

a. *Network*: Select the name of the appropriate network.

b. *Role*: Select **Cisco Router**.

c. From the list of routers, select the router that is to function as a Management CE (MCE).

d. Define the router as an MCE by choosing one of these two options:

- *Management LAN*
- *Management LAN, SA Agent*

  Selecting the *Management LAN, SA Agent* option defines the router as both an MCE and a CE with SA Agent enabled.

Step 6 Click **OK**. The selected router is designated as the MCE.

Step 7 To return to the VPN Console, click **OK** twice more.

Provisioning the Link Between the MCE and MPE

The next step is to provision a service request between the MCE and a PE designated as the Management PE (MPE). The recommended MPE to MCE connectivity is to use two subinterfaces or two physical interfaces. One interface should be used for the management VPN, and the other interface should be used for global table connectivity.

To provision the link between the MCE and MPE, follow these steps:

Step 1 From the VPN Console, choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor is displayed (see Figure 8-10).

Figure 8-10 The MPLS Service Request Editor

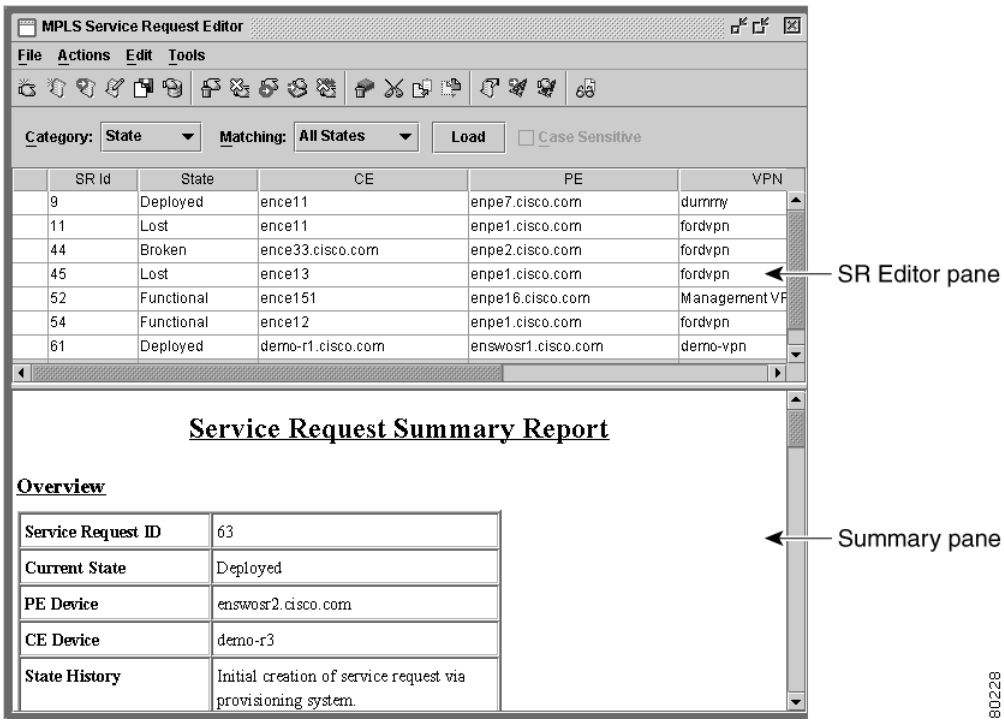

Step 2 From the MPLS Service Request Editor menu bar, choose **File > New Service Request(s)**.

The Open Profile dialog box appears (see Figure 8-11).

Figure 8-11 Selecting a Service Request Profile

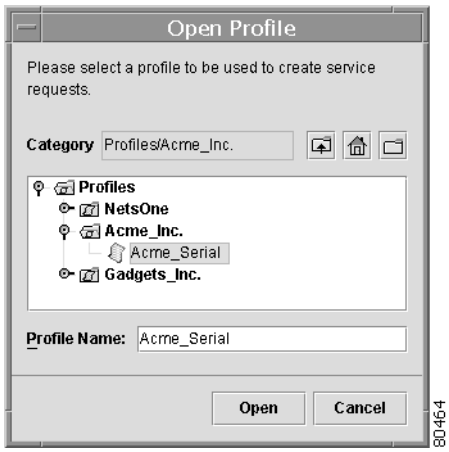

Step 3 From the list of service request profiles, select the appropriate profile for this service.

Step 4 Click **Open**.

The CE Device Chooser appears (see Figure 8-12).

When provisioning standard PE-CE links, this dialog box is used to select the CE in the PE-CE link. However, setting up a service request for the MCE is a special case, therefore, you can use this dialog box to select the router designated as the MCE.

Selecting the CEs for the Service Request

Figure 8-12 Selecting the CEs

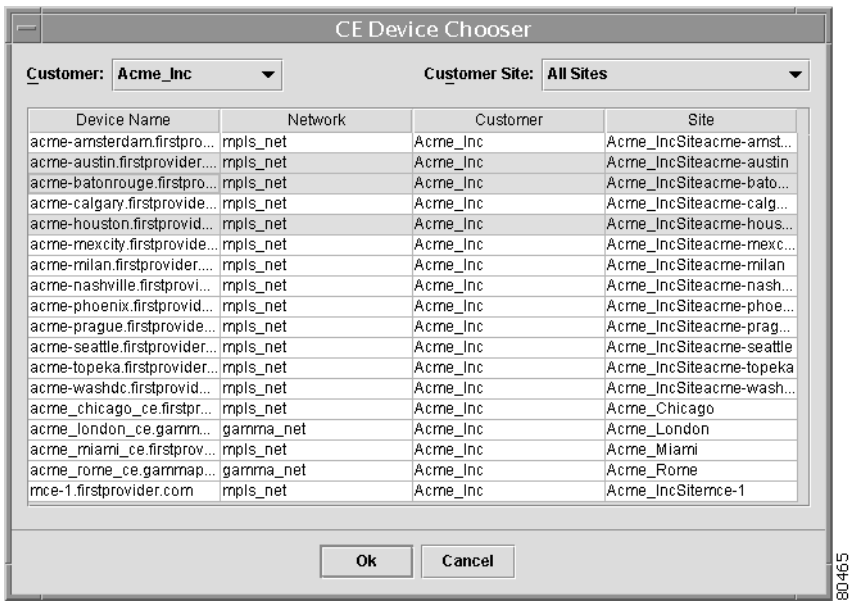

Step 5 Select the router that functions as the MCE.

a. *Customer*: From the Customer drop-down list, select the name of the VPN Customer.

b. *Customer Site*: From the Customer Site drop-down list, select the name of the site you want to see, or choose **All Sites** to see the list of all the sites for the selected Customer.

c. Select the MCE from CE Device Chooser.

d. Click **OK**.

You return to the Service Request Editor, where the CE information is now displayed.

When provisioning standard PE-CE links, the next dialog box is used to select the PE in the PE-CE link. However, for this operation, use this dialog box to select the router designated as the Management PE (MPE).

Selecting the MPE for the Service Request

Now that you have specified the MCE for the service request, the next step is to specify the MPE for the PE-CE link.

Step 6 Choose **Actions > Set PE**.

The PE Device Chooser appears (see Figure 8-13).

Figure 8-13 Selecting the MPE

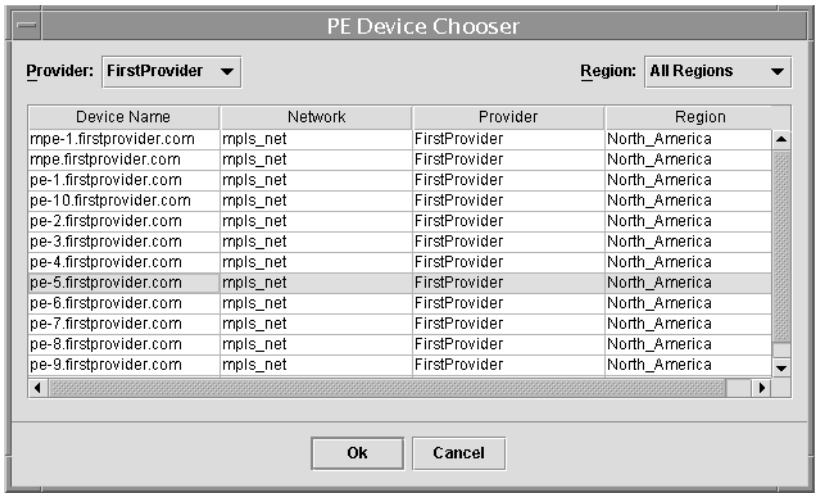

Step 7 Select the MPE for this service request.

a. *Provider*: From the Provider drop-down list, select the name of the service provider.

b. *Region*: From the Region drop-down list, select the name of the region the PE is in, or choose **All Regions** to see the list of all the PEs for the selected provider.

c. Select the MPE from the PE Device Chooser.

d. Click **OK**.

If you're creating multiple service requests, you will receive the following **Set PE** informational prompt:

```
Changed the PE device for x service requests.
```

e. Click **OK**.

You return to the Service Request Editor, where the name of the selected PE is now displayed.

Specifying the VPN Membership Information

Step 8 In the Editor pane of the MPLS Service Request Editor, select one or more of the PE-CE pairs you want to associate with a particular VPN.

Step 9 Choose **Actions > Set VPN Memberships**.

The VPN Memberships dialog box appears (see Figure 8-14).

Figure 8-14 Specifying the VPN Membership for the Devices

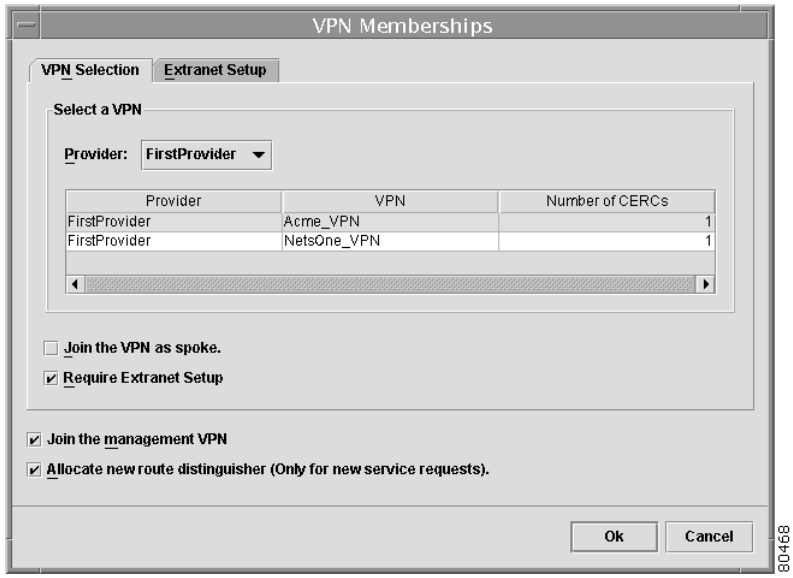

Specifying the VPN

Step 10 Select the VPN for this service request.

a. *Provider*: From the Provider drop-down list, select the name of the service provider.

b. From the list of VPNs displayed for the selected service provider, select the appropriate VPN.

Specifying a Hub-and-Spoke or Full Mesh VPN

Step 11 If you are building a VPN with a hub-and-spoke topology, enable the **Join as Spoke** option.

- If you want the MCE to *not* have access to all the other sites in the VPN, be sure to enable the **Join as spoke** option.
- If you want the MCE to have access to all the other sites in the VPN, do *not* enable the **Join as spoke** option.

Joining the Management VPN

Step 12 To add the MCE to the *management VPN*, enable the **Join the management VPN** option.

For more information on the management VPN, see Chapter 8, "The VPNSC Management Network."

When you use the VPN Solutions Center software to define a management VPN, the software automatically generates an *export route map* for the management VPN.

Set Up Load Balancing for a Multihomed CE

The **Allocate new route distinguisher** option allows the selected PEs to have unique route distinguishers (RDs) within the current VPN. Having unique RDs on each PE lets BGP load balance the traffic in the case where you have a dual-homed CE with links to two PEs.

Figure 8-15 shows a service provider BGP MPLS network that connects two networks (or subnets) to PE-1 and PE-2. Both of these PEs are configured for VPNv4 unicast iBGP peering. Network 2 is a multihomed subnet that is connected to both PE-1 and PE-2. Network 2 also has Extranet VPN services configured with Network 1. Both Network 1 and Network 2 are configured for eBGP peering with the PEs.

Figure 8-15 Unique RDs Assigned to PEs in the Same VPN

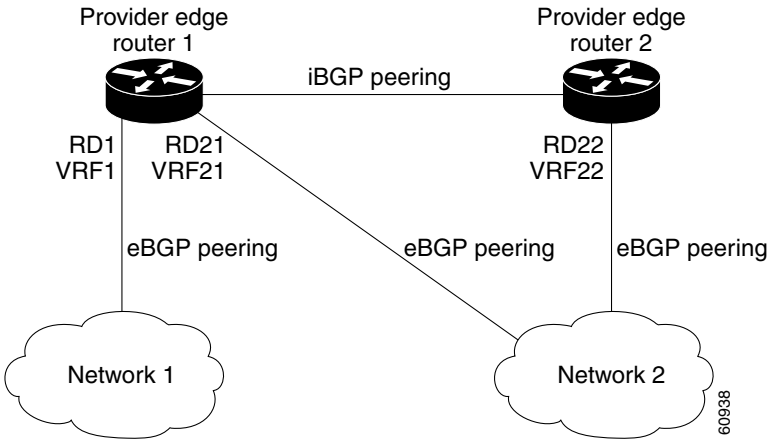

As shown in Figure 8-15, PE-1 and PE-2 have unique RDs (RD 21 and RD 22 respectively). The multipaths between Network 2 and PE-1 and PE-2 performs load balancing when the **Allocate new route distinguisher** option is enabled. Any prefix that is advertised from Network 2 will be received by RD 21 and RD 22. Thus, any traffic to Network 2 will be load balanced, with half the traffic going through PE-1 and half the traffic going through PE-2.

Step 13 To enable load balancing as described above, enable the **Allocate new route distinguisher** option.

If the VPN Has CEs in Other VPNs (Extranets)

Step 14 If you are building a VPN with CEs that are members of multiple VPNs (also referred to as *extranets*), enable the **Require Extranet Setup** option.

Extranet provisioning provides a way to create multiple VPN connectivity to a single VRF. You can add multiple CERCs to your VPN in any topology to form extranets. You can join an extranet in such a way that a CE can be a spoke in one VPN and a hub in another VPN.

a. Enable the **Require Extranet Setup** option.

When you do so, the Extranet Setup tab is enabled.

b. Click the **Extranet Setup** tab (see Figure 8-16).

Figure 8-16 Extranet Setup

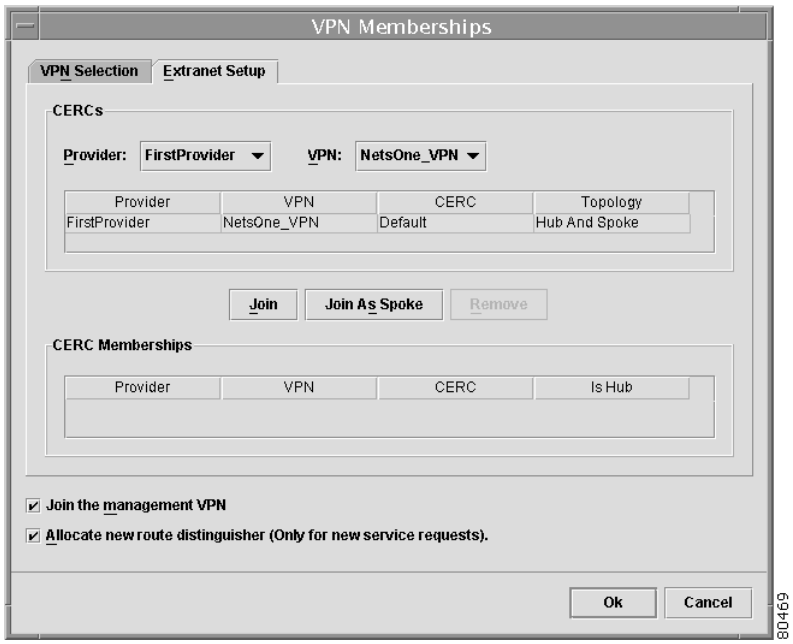

a. *Provider*: Specify the service provider name.

b. *VPN*: Specify the name of the other VPN that this CERC is a member of.

c. Specify whether the selected CERC is a hub or a spoke.

- If the selected CERC is a hub, click **Join**.

  The selected *hub CERC* is now displayed in the CERC Memberships panel. Note that the "Is Hub" checkbox is enabled (see Figure 8-17).

Figure 8-17 Hub CERC Added to Another VPN

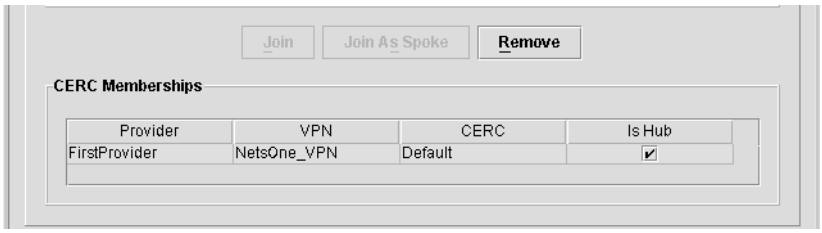

- If the selected CERC is a spoke, click **Join As Spoke**.

  The selected *spoke CERC* is now displayed in the CERC Memberships panel (see Figure 8-18)

Figure 8-18 Spoke CERC Added to Another VPN

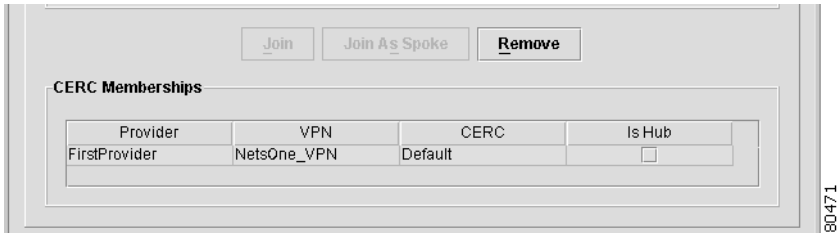

d. When satisfied with the Extranet settings, click **OK**.

   If you're creating multiple service requests, you will receive the following **Set VPN** informational prompt:

   ```
   Changed the VPN memberships for x service requests.
   ```

e. Click **OK**.

You return to the Service Request Editor, where all of the fields are filled in, indicating you are ready to deploy the service requests (see Figure 8-19).

Figure 8-19 Parameters Completed for Service Requests

If You Wish to Delay Service Request Deployment

You are not required to immediately deploy the service requests you have set up. If you wish to delay service request deployment for any reason, you can commit to the Repository the existing service requests in their current state. You can either commit some or all of the existing service requests to the Repository.

To commit new service requests to the Repository:

1. If you wish to commit selected service requests only, from the MPLS Service Request Editor, select the service requests you want to commit.
2. Depending on whether you are committing some or all of the new service requests, do one of the following:
   - For selected service requests, choose **File > Commit to Repository**
   - For all the new service requests, choose **File > Commit All to Repository**.

   The message bar at the bottom of the VPN Console displays the message:

   ```
   Committed x service requests to the Repository.
   ```

   VPN Solutions Center moves the selected service requests to the Requested state and assigns service request IDs to each (see Figure 8-20).

Figure 8-20 Service Requests Committed to the Repository

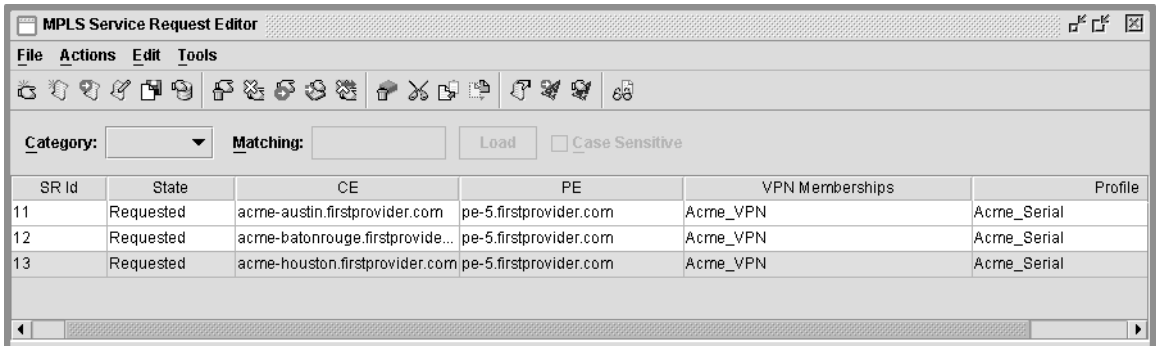

You can now deploy the committed service requests as your convenience.

Provisioning MPLS VPN Cable Services

This chapter provides a conceptual summary of the MPLS VPN Cable feature as implemented through the VPN Solutions software. It also describes how to use VPN Solutions software to provision cable services. The main topics presented in this chapter are as follows:

MPLS VPN Cable Feature Overview

Using MPLS VPN technology, service providers can create scalable and efficient private networks using a shared Hybrid Fiber Coaxial (HFC) network and Internet Protocol (IP) infrastructure. The cable MPLS VPN network consists of the following two major elements:

- The Multiple Service Operator (MSO) or cable company that owns the physical infrastructure and builds VPNs for the Internet Service Providers (ISPs) to move traffic over the cable and IP backbone.
- ISPs that use the HFC network and IP infrastructure to supply Internet service to cable customers.

The Cable MPLS VPN Network

As shown in Figure 9-1, each ISP moves traffic to and from a subscriber's PC, through the MSO's physical network infrastructure, to the ISP's network. MPLS VPNs, created in Layer 3, provide privacy and security by constraining the distribution of a VPN's routes only to the routers that belong to its network. Thus, each ISP's VPN is insulated from other ISPs that use the same MSO infrastructure.

In the MPLS-based cable scheme, a VPN is a private network built over a shared cable plant and MPLS-core backbone. The public network is the shared cable plant or backbone connection points. A cable plant can support Internet access services and carry traffic for an MSO and its subscribers, as well as for multiple Internet Service Providers (ISPs) and their subscribers.

An MPLS VPN assigns a unique VPN Routing/Forwarding (VRF) instance to each VPN. A VRF instance consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine the contents of the forwarding table.

Each PE router maintains one or more VRF tables. If a packet arrives directly through an interface associated with a particular VRF, the PE looks up a packet's IP destination address in the appropriate VRF table. MPLS VPNs use a combination of BGP and IP address resolution to ensure security.

Figure 9-1 Example of an MPLS VPN Cable Network

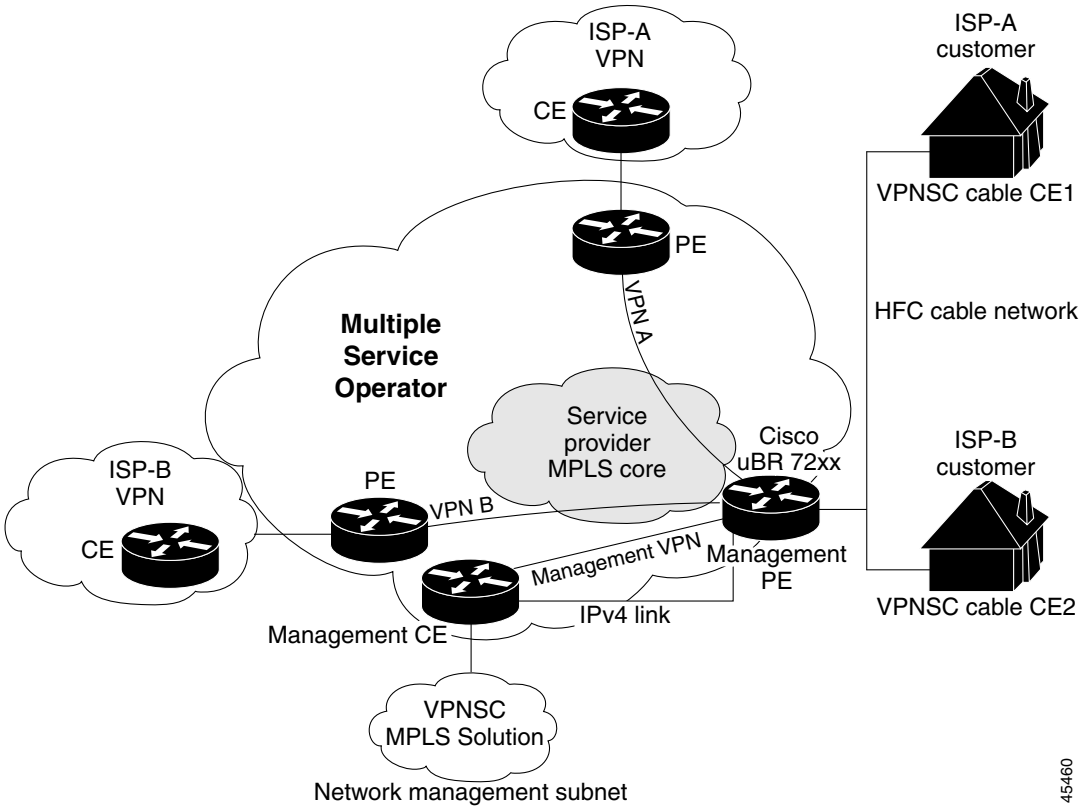

The routers in the cable network are as follows:

- *Provider (P) router*—Routers in the MPLS core of the service provider network. P routers run MPLS switching, and do not attach VPN labels (MPLS labels in each route assigned by the PE router) to routed packets. VPN labels direct data packets to the correct egress router.
- *Provider Edge (PE) router*—A router that attaches the VPN label to incoming packets based on the interface or subinterface on which they are received. A PE router is connected to a CE router. In the MPLS-VPN approach, each Cisco uBR7200 series router acts as a PE router.
- *Customer (C) router*—A router in the ISP or enterprise network.
- *Customer Edge (CE) router*—Edge router on the ISP's network that connects to the PE router on the MSO's network. A CE router must interface with a PE router.
- *Cable CE*—The cable CE is an object with the VPN Solutions Center only. In the VPN Solutions Center software, the cable CE represents a group of cable modems and its associated hosts for a particular site (see the "Creating a Cable-CE in VPNSC Software" section on page 9-5).
- *Management PE (MPE) router*—The MPE *emulates* the role of a PE in the provider core network. The MPE connects the MCE to the provider core network. An MPE can have a dual role as both a PE and the MPE.

- *Management CE (MCE) router*—The network management subnet is connected to the Management CE (MCE). The MCE *emulates* the role of a customer edge router (CE), but the MCE is in provider space and serves as a network operations center gateway router. The MCE is part of a management site as defined in the VPN Solutions Center software.

The shared cable plant supports Internet connectivity from ISP A to its subscribers and from ISP B to its subscribers.

The Management VPN in the Cable Network

The MPLS network has a unique VPN that exclusively manages the MSO's devices called the *management VPN*. It contains servers and devices that other VPNs can access. The management VPN connects the Management CE (MCE) router and the management subnet to the MSO PE router (a Cisco uBr72xx router or equivalent). VPN Solutions Center and the management servers, such as Dynamic Host Configuration Protocol (DHCP), Cisco Network Registrar (CNR), and Time of Day (ToD) servers are part of the management subnet and are within the management VPN for ISP connectivity.

As shown in Figure 9-1, the management VPN is comprised of the network management subnet (where the VPN Solutions Center workstation resides), which is directly connected to the Management CE (MCE). The management VPN is a special VPN for the MCE and the cable VPN gateway. The cable VPN gateway is usually a Cisco uBR 72xx router that functions as both a regular PE and a Management PE. Notice that there is also a parallel IPv4 link between the MCE and the MPE.

Cable VPN Configuration Overview

Cable VPN configuration involves the following:

- An MSO domain that requires a direct peering link to each enterprise network (ISP), provisioning servers for residential and commercial subscribers, and dynamic DNS for commercial users. The MSO manages cable interface IP addressing, Data Over Cable Service Interface Specifications (DOCSIS) provisioning, cable modem hostnames, routing modifications, privilege levels, and usernames and passwords.
- An ISP or enterprise domain that includes the DHCP server for subscriber or telecommuter host devices, enterprise gateway within the MSO address space, and static routes back to the telecommuter subnets.

Note Cisco recommends that the MSO assign all addresses to the end user devices and gateway interfaces. The MSO can also use split management to let the ISP configure tunnels and security.

To configure MPLS VPNs for cable services, the MSO must configure the following:

- *Cable Modem Termination System (CMTS)*

  The CMTS is usually a Cisco uBR72xx series router. The MSO must configure Cisco uBR72xx series routers that serve the ISP.

- *PE routers*

  The MSO must configure PE routers that connect to the ISP as PEs in the VPN.

Tip When configuring MPLS VPNs for cable services, you must configure the cable maintenance subinterface on the PE. The cable maintenance interface is the means by which the cable device retrieves its own IP address. For this reason, the maintenance subinterface must be configured before cable services provisioning can take place. See the "Provisioning the Cable Maintenance Subinterface" section on page 9-9.

- *CE routers*
- *P routers*
- *One VPN per ISP*
- *DOCSIS servers for all cable modem customers*

  The MSO must attach DOCSIS servers to the management VPN and make them visible to the network.

The MSO must determine the *primary IP address range*. The primary IP address range is the MSO's address range for all cable modems that belong to the ISP subscribers.

The ISP must determine the *secondary IP address range*. The secondary IP address is the ISP's address range for its subscriber PCs.

To reduce security breaches and differentiate DHCP requests from cable modems in VPNs or under specific ISP management, MSOs can use the **cable helper-address** command in Cisco IOS software. The MSO can specify the host IP address to be accessible only in the ISP's VPN. This lets the ISP use its DHCP server to allocate IP addresses. Cable modem IP address must be accessible from the management VPN.

In VPN Solutions Center 2.1 software, you specify the maintenance helper address (see the "Specifying the Cable Maintenance Helper Addresses" section on page 9-15), and the host helper address and the secondary addresses for the cable subinterface (see "Specifying the Cable Helper Secondary Addresses" section on page 9-26).

Cable VPN Interfaces and Subinterfaces

In the cable subscriber environment, several thousand subscribers share a single physical interface. Configurations with multiple logical subinterfaces are a vital part of the MPLS VPN network over cable. You can configure multiple subinterfaces and associate a specific VRF with each subinterface. You can split a single physical interface (the cable plant) into multiple subinterfaces, where each subinterface is associated with a specific VRF. Each ISP requires access on a physical interface and is given its own subinterface. The MSO administrator can define subinterfaces on a cable physical interface and assign Layer 3 configurations to each subinterface.

The MPLS VPN approach of creating VPNs for individual ISPs or customers requires subinterfaces to be configured on the cable interface. One subinterface is required for each ISP. The subinterfaces are tied to the VPN Routing/Forwarding (VRF) tables for their respective ISPs.

You must create the *maintenance subinterface* on the cable interface and tie it to the management VPN. The maintenance interface is for the ISP's use, and it is used for VPN connectivity, as well as the management VPN using an extranet between the ISP and the management VPN (for details, see "Provisioning the Cable Maintenance Subinterface" section on page 9-9).

VPN Solutions Center software automatically selects the subinterface number based on the VRF. If a subinterface that is associated with the current VRF does not yet exist, VPNSC software creates a subinterface and assigns it to the correct VRF. The subinterface number is incremented to 1 greater than the largest subinterface currently assigned for the selected cable interface.

The network management subnet (which includes the CNR, ToD, and VPN Solutions Center) can reply to the cable modem because the management VPN allows connectivity for one filtered route from the ISP's VPN to the Management CE (MCE). Similarly, in order to forward the management requests (such as DHCP renewal to CNR), the ISP VPN must import a route to the MCE in the management VPN.

Cisco uBR7200 series software supports the definition of logical network layer interfaces over a cable physical interface. The system supports subinterface creation on a physical cable interface.

Subinterfaces allow traffic to be differentiated on a single physical interface and associated with multiple VPNs. Each ISP requires access on a physical interface and is given its own subinterface. Using each subinterface associated with a specific VPN (and therefore, ISP) subscribers connect to a logical subinterface, which reflects the ISP that provides their subscribed services. Once properly configured, subscriber traffic enters the appropriate subinterface and VPN.

Creating a Cable-CE in VPNSC Software

The tasks you must complete to provision cable services in VPN Solutions Center software are as follows:

- Create a stand-in cable-CE (described in this section). This cable-CE is an object within VPNSC software only. It represents the cable modem and its associated hosts for a particular site.
- Add the PE that has cable interfaces to the appropriate Region (for procedural details, see the "Defining Provider Administrative Domains" section on page 4-30).
- Generate a service request to provision the cable maintenance interface on the PE (see the "Provisioning the Cable Maintenance Subinterface" section on page 9-9). You need only generate this service request once for each physical cable interface.
- Generate a second service request to provision the MPLS-based cable service ("Provisioning the Cable Link" section on page 9-20). You must generate this cable service request for each VPN.

When using the VPN Solutions Center to provision cable services, there are no CEs in the same sense there are when provisioning a standard MPLS VPN. Thus, you must create an unmanaged *cable-CE* that "stands in" for a CE in the provisioning process. You need define only one cable-CE per customer site.

To create a cable-CE in VPN Solutions Center software, follow these steps.

Step 1 In the VPN Console, open the Networks folder and select the pertinent network. The Network window appears (see Figure 9-2).

Figure 9-2 The Network Window

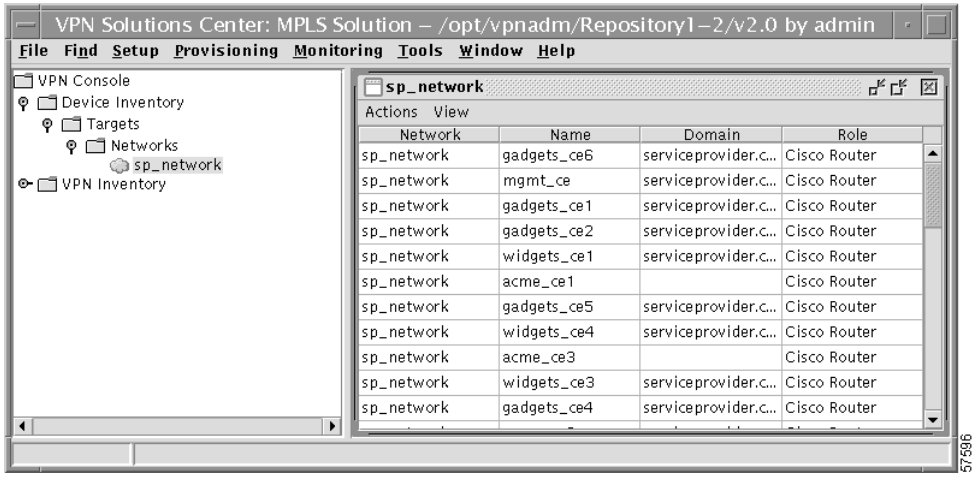

Step 2 From the Network window, choose **Actions > New Target**. The New Target dialog box appears (see Figure 9-3).

Figure 9-3 Creating a Cable-CE

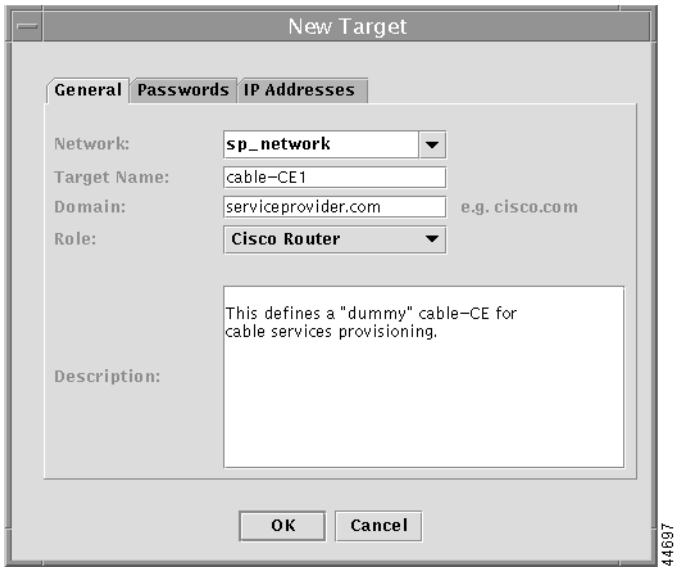

Step 3 Complete the fields displayed in the General tab. You do not need to complete the fields in the Passwords and IP Addresses tabs.

a. In the *Target Name* field, enter the name of the cable-CE.

This cable-CE is an object within VPNSC software only. It represents the cable modem and its associated hosts for a particular site.

b. In the *Domain* field, enter the name of a nonexistent domain.

c. Though optional, we recommend that you enter any pertinent information about the cable-CE.

d. Click **OK**.

Step 4 In the VPN Console, open the VPN Customers folder.

Step 5 Select the appropriate customer, then to display the list of sites for that customer, open the customer icon.

Step 6 **Double-click** the customer site where you want to place the cable-CE.

The Edit Customer Site dialog box appears (see Figure 9-4).

Figure 9-4 Editing the Customer Site

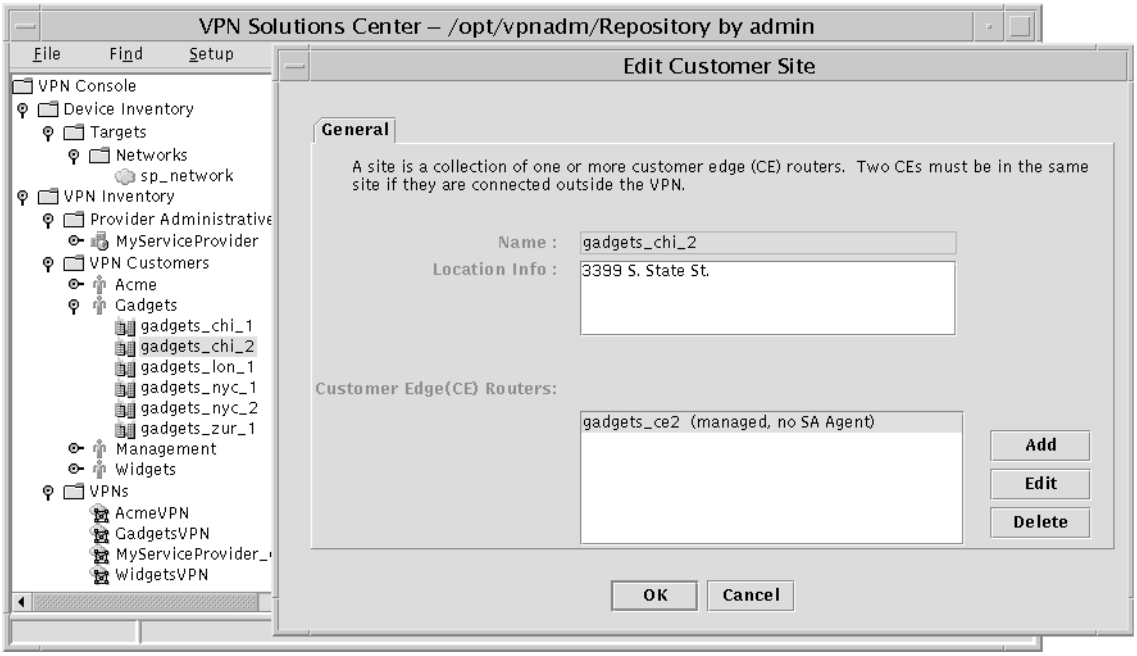

Step 7 In the Edit Customer Site dialog box, click **Add**.

The Add Customer Edge Routers dialog box appears. Figure 9-5 shows only the lower portion of this dialog box.

Figure 9-5 Adding the Cable-CE to the Network

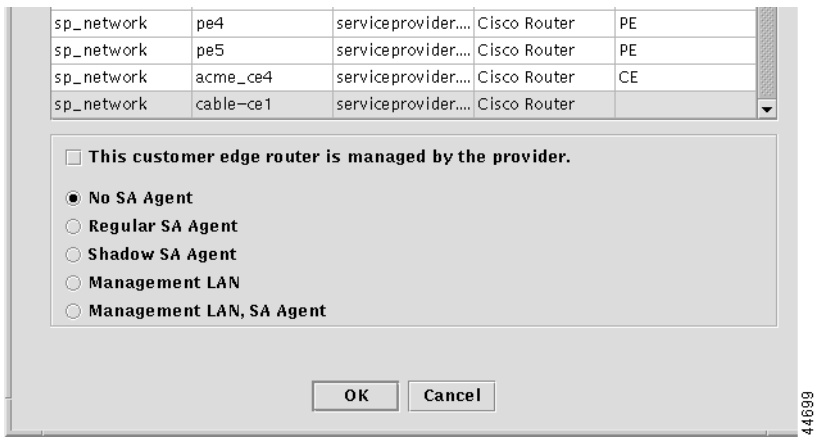

Step 8 In the Add Customer Edge Routers dialog box, do the following:

a. From the list of devices displayed, select the name of the cable-CE.

b. A cable-CE must be an unmanaged CE, so be sure that the **This customer edge router is managed by the provider** check box is *not* checked.

c. Choose the **No SA Agent** option.

SA Agent can gather performance information from managed CEs only.

Note Do *not* choose either of the **Management LAN** options.

d. When finished, click **OK**.

Provisioning the Cable Maintenance Subinterface

The cable maintenance subinterface on the PE is the means by which the cable device retrieves its own IP address. For this reason, the maintenance subinterface must be configured before cable services provisioning can take place.

This procedure assumes that the PE configuration files that define the cable maintenance interfaces have been imported into VPN Solutions Center. For a description of this procedure, see the "Importing Provider Edge Routers into VPN Solutions Center" section on page 4-4.

Setting Up the Cable Maintenance Interface on the PE

For this procedure, we will use the VPNSC v2.1 service request user interface.

To set up the cable maintenance subinterface on the PE:

Selecting the Cable-CE

When configuring a service for a cable link, the specified CE should be an *unmanaged CE*.

Step 1 From the VPN Console, choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor is displayed (see Figure 9-6).

Figure 9-6 MPLS Service Request Editor

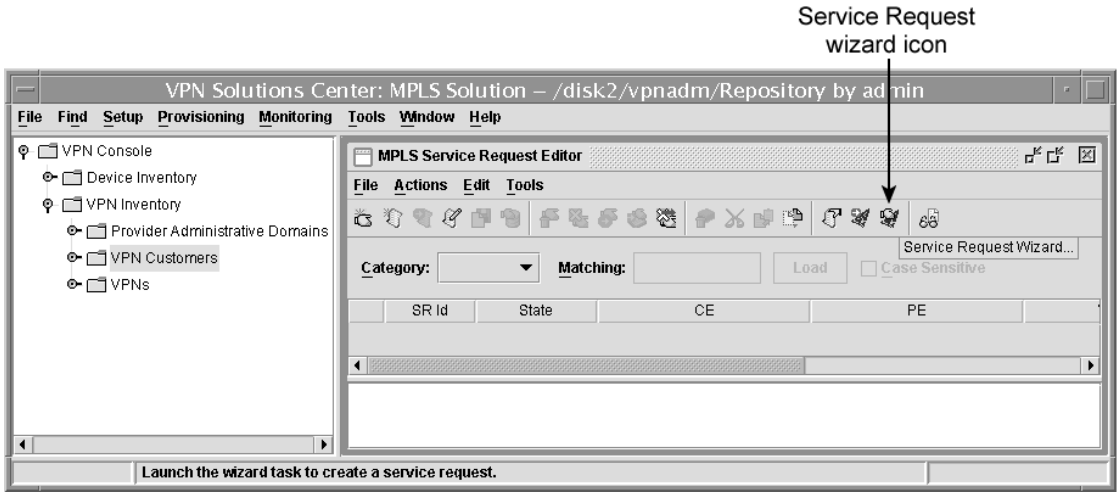

Step 2 To switch to the VPNSC 2.1 service request wizard, click the Service Request wizard icon.

The first—and informational only—screen appears.

Step 3 Click **Next**.

The Select CE dialog box appears (see Figure 9-7).

Figure 9-7 The Select CE Dialog Box

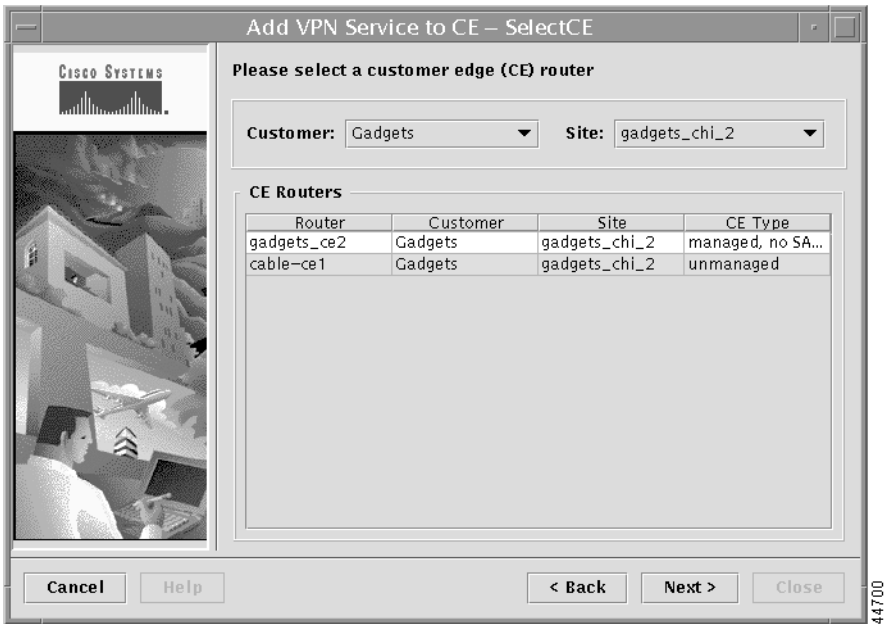

Step 4 From the Select CE dialog box, do the following:

a. From the Customer drop-down list, select the appropriate customer.

b. From the Site drop-down list, select the appropriate site.

c. From the CE Routers list, select the name of the cable-CE.

d. Click **Next**. The Select PE dialog box appears (see Figure 9-8).

Selecting the PE for Cable Service

Step 1 From the Select PE dialog box, select the provider edge router for this cable link.

Figure 9-8 Selecting the Cable-PE

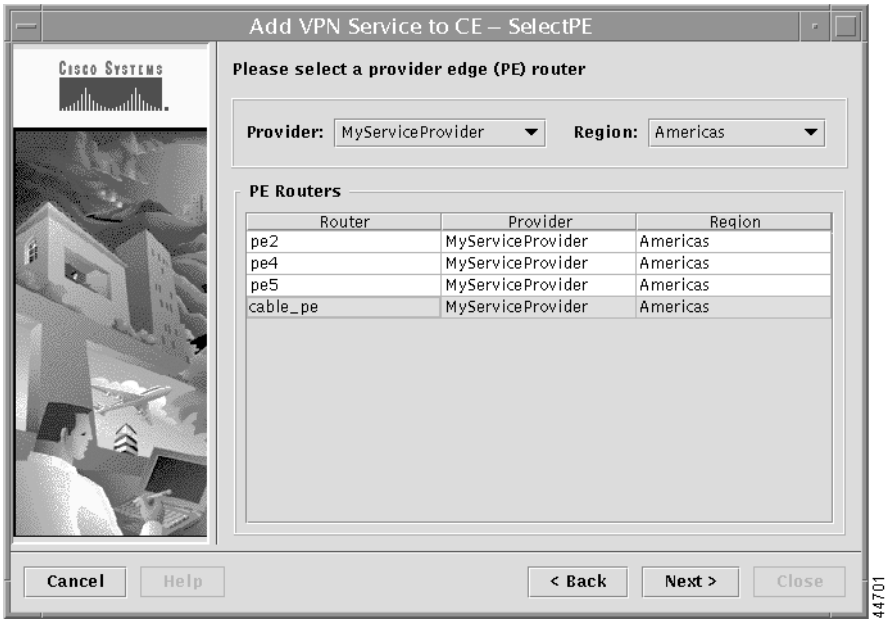

Step 2 From the Provider drop-down list, select the appropriate provider name.

Step 3 From the Region drop-down list, select the appropriate region.

Step 4 From the PE Routers list, select the PE.

Step 5 When finished entering the necessary information, click **Next**. The Select VPN dialog box appears (see Figure 9-9).

Selecting the VPN

Step 1 From the Select VPN dialog box shown in Figure 9-9, select the VPN that the cable maintenance interface is associated with.

Figure 9-9 Selecting the VPN

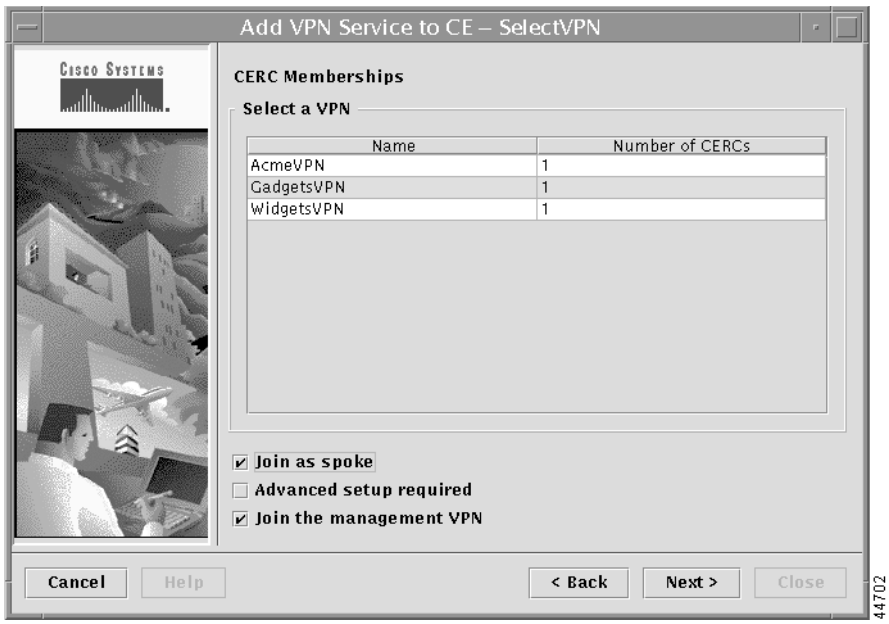

The most common types of VPNs are *hub-and-spoke* and *full mesh*. These two basic types of VPNs—full mesh and hub and spoke—can be represented with a single CERC.

For additional information on CE routing communities, see the "CE Routing Communities" section on page 1-18 and the "Defining CE Routing Communities" section on page 5-3.

Step 2 If you are building a VPN with a hub-and-spoke topology, check the **Join as Spoke** check box.

- A hub-and-spoke CERC is one in which one or a few CEs act as hubs, and all spoke CEs talk only to or through the hubs, never directly to each other.
- A full mesh CERC is one in which every CE connects to every other CE.

Step 3 If you are building a VPN with CEs that are members of multiple VPNs (extranets), check the **Advanced setup required** check box.

Extranet provisioning provides a way to create multiple VPN connectivity to a single VRF.

Step 4 When provisioning a cable service maintenance interface, joining the management VPN is required. Therefore, check the **Join the management VPN** check box. For more information on the management VPN and VPN Solutions Center software, see the "About Service Request Profiles" section on page 5-6.

When you use the VPN Solutions Center software to define a management VPN, the software automatically generates an *export route map* for the management VPN.

Step 5 When finished entering the necessary information, click **Next**. The Select Routing Policy dialog box appears (see Figure 9-10).

Specifying No Routing Protocol for the Cable CE

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

Step 1 From the list of routing protocol options (see Figure 9-10), choose **None**.

Figure 9-10 Specifying No Routing Protocol

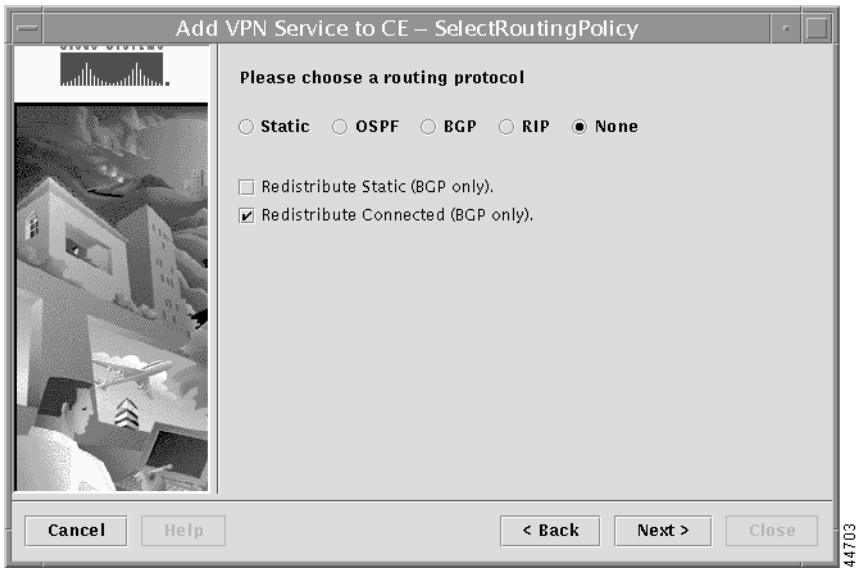

Step 2 If you want to distribute static routes into the provider core network (which runs BGP), check the **Redistribute Static (BGP only)** check box.

Step 3 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.

Step 4 When finished entering the necessary information, click **Next**. The Select PE-CE Interface dialog box appears (see Figure 9-11).

Specifying the Cable Maintenance Interface on the PE

You can now specify the interface on the PE that will host the cable maintenance subinterface.

Figure 9-11 Selecting the Cable Maintenance Interface

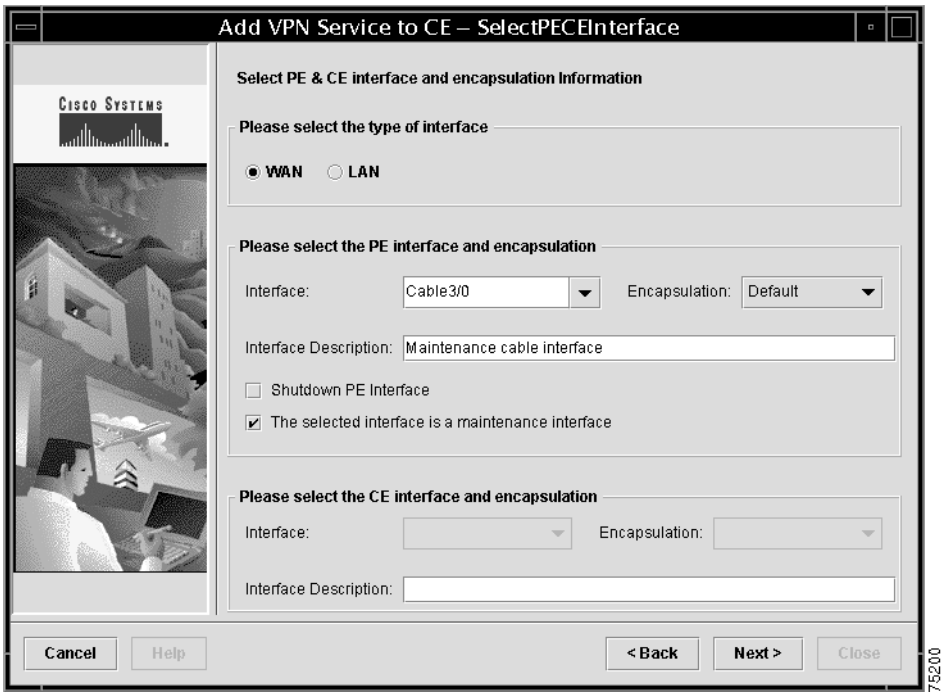

Step 1 *Interface Type*: Specify whether the interfaces for the PE-CE link are for a Wide Area Network (WAN) or Local Area Network (LAN).

Step 2 *PE Interface*: Select the interface on the PE that hosts the cable maintenance subinterface.

The encapsulation method is set to *Default* for the cable interface.

Step 3 *Interface Description*: Optionally, you can enter a description of the cable interface. The interface description entered here is also added to the configuration file.

Step 4 *Shutdown PE Interface*: Enable the **Shutdown PE Interface** option if desired:

When you check the **Shutdown PE Interface** checkbox, the specified PE interface will be configured in a shut down state.

Step 5 *Maintenance Interface*: Be sure to check the **Selected interface is a maintenance interface** check box.

Checking this option provisions the cable maintenance interface; this interface is always configured as subinterface 1 (for example, if the selected cable interface is **3/0**, the maintenance subinterface is **3/0.1**).

Step 6 When finished with these settings, click **Next**. The Select Cable Parameters dialog box appears (see Figure 9-12).

Specifying the Cable Maintenance Helper Addresses

The *maintenance helper address* is the IP address of the DHCP server in the Multiple Service Operator (MSO) network. You can add three types of maintenance helper addresses:

- Cable host helper address

  The IP address of the DHCP server of the Internet Service Provider (ISP) to which the customer belongs. A cable host helper address specifies that only cable host UDP broadcasts are forwarded.

- Cable modem helper address

  The IP address of the DHCP server in the MSO's network. The modem helper address assigns the IP address of the cable modem interface. A cable modem helper address specifies that only cable modem UDP broadcasts are forwarded.

- Cable helper address that is both the host and modem address

Figure 9-12 Setting the Maintenance Helper Address

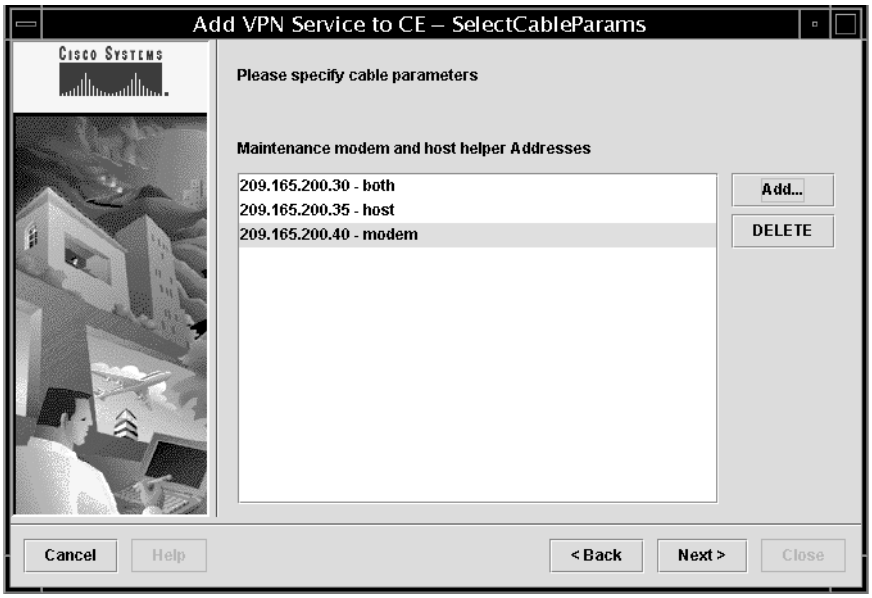

Step 1 To specify the maintenance helper addresses, click **Add**.

The following dialog box appears (see Figure 9-13).

Figure 9-13 Specifying Modem and Host Helper Addresses

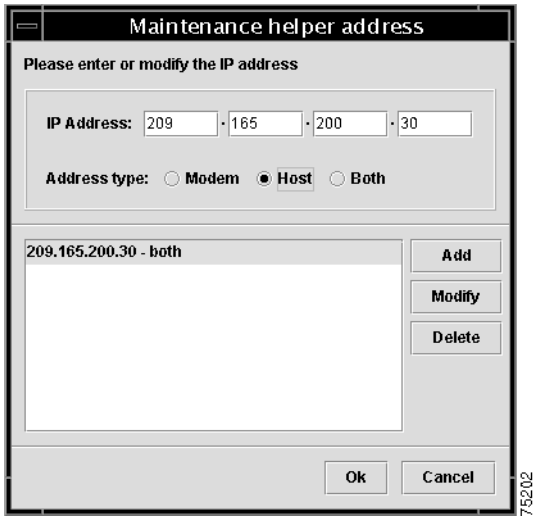

Step 2 Select the type of the maintenance helper address:

- Modem
- Host
- Both

Step 3 Enter the IP address in the fields provided.

Step 4 When finished with this helper address, click **Add**.

Step 5 To add additional maintenance helper addresses, repeat Steps 2, 3, and 4.

Step 6 When you are finished adding helper addresses, click **OK**.

You return to the dialog box shown in Figure 9-12 on page 9-15, where the maintenance helper addresses you entered here are displayed.

You can delete any of the helper addresses by selecting the address and clicking **Delete**.

Step 7 When satisfied with the cable maintenance helper address settings, click **OK**.

The Select IP Addresses dialog box appears (see Figure 9-14).

Specifying the IP Address for the Maintenance Subinterface

In the Select IP Addresses dialog box, you must specify the IP address for the cable maintenance subinterface on the PE.

Figure 9-14 Specifying the IP Address for the Maintenance Subinterface

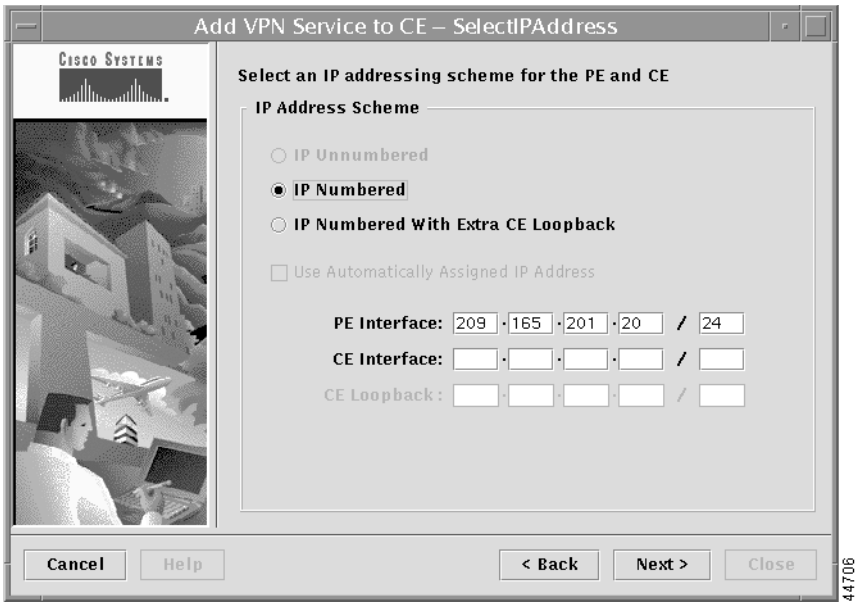

Step 1 For the IP addressing scheme, choose **IP Numbered**.

The *IP Numbered with Extra CE Loopback* option is not a viable option in a cable services configuration.

Step 2 In the **PE Interface** fields, enter the IP addresses on the PE for the cable maintenance subinterface.

Tip The IP address entered here must be different from the IP address entered for the cable subinterface. To be reachable, each subinterface must have its own IP address.

You do not need to enter an IP address for the CE interface.

Step 3 When finished entering the necessary information, click **Next**.

The Select VRF Parameters dialog box appears (see Figure 9-15).

Specifying the VRF Parameters

The Select VRF Parameters dialog box lets you set values for an import route map and the maximum number of routes in the VRF table. You can also enable NetFlow accounting.

Figure 9-15 Specifying the VRF Parameters

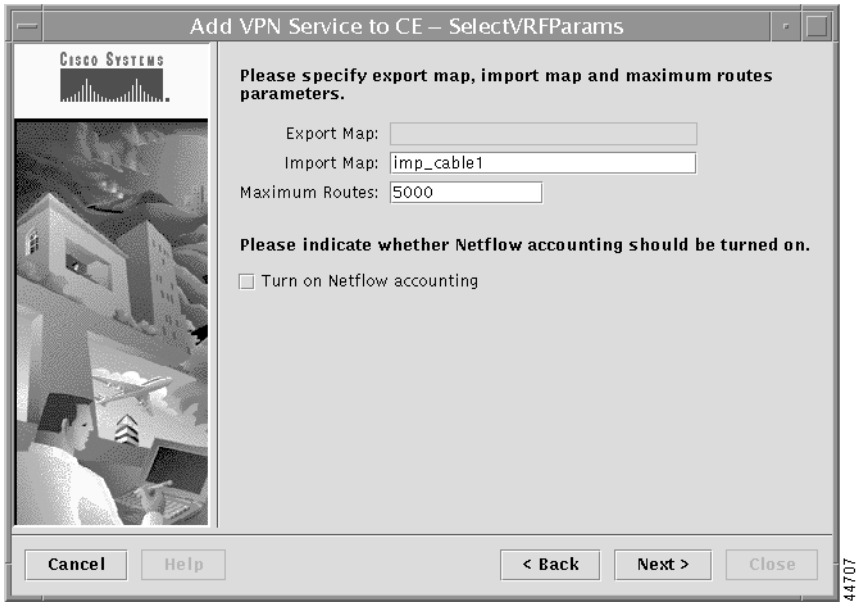

Step 1 In the *Import Map* field, enter the name of an existing import route map on the PE.

Note The Cisco IOS supports only one import route map per VRF (and therefore, 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 this PE to exclude the route.

For command reference details on the **import map** command, see the "import map" section on page B-4.

Step 2 In the *Maximum Routes* field, specify the maximum number of routes that can be imported into the VRF on this PE.

Step 3 To enable NetFlow accounting, check the **Turn on NetFlow accounting** checkbox.

Step 4 When you have completed the fields as necessary in the Specify VRF Parameters dialog box, click **Next**. The Class of Service (CoS) dialog box appears.

Selecting a Class of Service Profile for the PE-CE Link

Step 1 If desired, select a Class of Service (CoS) profile to assign to the PE-CE link.

You can create a Class of Service (CoS) profile when you define the Provider Administrative Domain. For information on creating a CoS Profile, see the "Defining a Class of Service Profile" section on page 4-35.

Class of Service profiles are applied to the Provider Edge Router (PE), but the CoS definition is enforced across the PE-CE link on both the PE and CE.

Step 2 Click **Next**. The Confirm dialog box appears (see Figure 9-16).

Confirming the Cable VPN Service Settings

VPN Solutions Center displays a summary of settings defined for this cable services VPN.

Figure 9-16 Viewing the Service Settings

Step 1 Verify that the service request information is correct, then click **Next**. The wizard displays the following message:

Your request to "Add VPN Service to CE" has been submitted with ID number *n*. This service request can be deployed by using the "Deploy Service Requests" wizard or by using the "Deploy VPN Service" item under the "Provisioning" option of a VPN service request report.

Step 2 Press **Close**. You have now queued a service request. It is entered into the product database and is in the initial state of "Requested."

Provisioning the Cable Link

When you have completed the tasks to create the cable-CE, set up the PE that has cable interfaces, and provisioned the cable maintenance subinterface on the PE as described in the previous sections, you can proceed to provision the cable link.

To provision the cable link:

Selecting the Cable-CE

Step 1 From the VPN Console, choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor is displayed (see Figure 9-17).

Figure 9-17 MPLS Service Request Editor

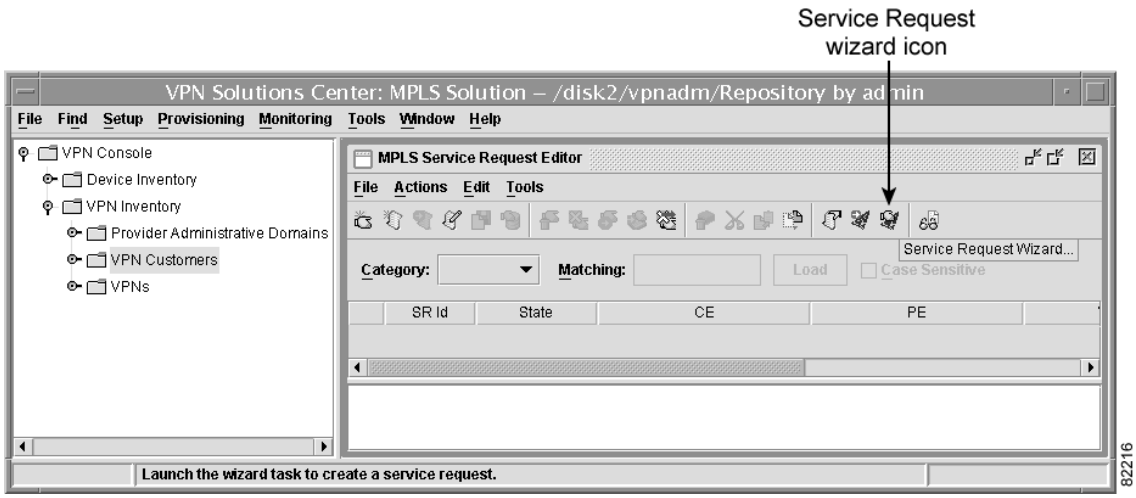

Step 2 To switch to the VPNSC 2.1 service request wizard, click the Service Request wizard icon.

The first—and informational only—screen appears.

Step 3 Click **Next**.

The Select CE dialog box appears (see Figure 9-18).

Figure 9-18 Selecting the Cable-CE

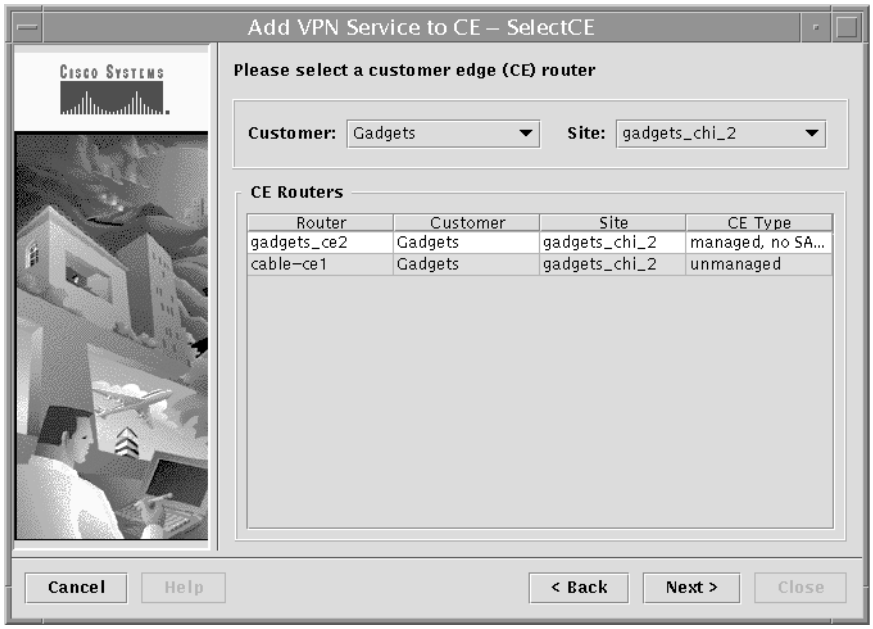

Step 4 From the Select CE dialog box, do the following:

a. From the Customer drop-down list, select the appropriate customer.

b. From the Site drop-down list, select the appropriate site.

c. From the CE Routers list, select the name of the cable-CE.

Note When configuring a service for a cable link, the specified CE should be an unmanaged CE.

d. Click **Next**. The Select PE dialog box appears (see Figure 9-19).

Selecting the PE for Cable Service

Step 1 From the Select PE dialog box, select the provider edge router for this cable link.

Figure 9-19 Selecting the Cable-PE

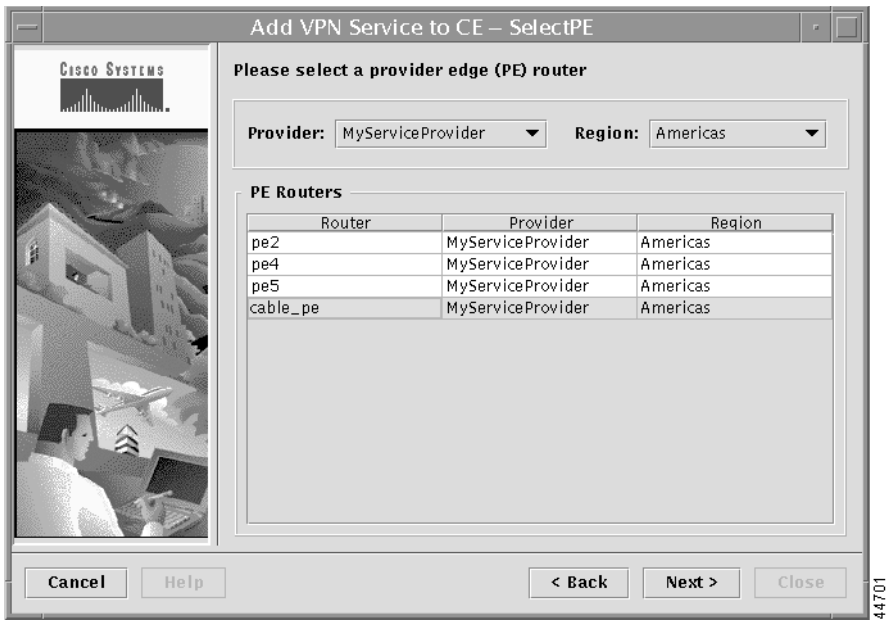

Step 2 From the Provider drop-down list, select the appropriate provider name.

Step 3 From the Region drop-down list, select the appropriate region.

Step 4 From the PE Routers list, select the PE.

Step 5 When finished entering the necessary information, click **Next**. The Select VPN dialog box appears (see Figure 9-20).

Selecting the VPN

Step 1 From the Select VPN dialog box, select the VPN that the cable interface is associated with.

Figure 9-20 Selecting the VPN

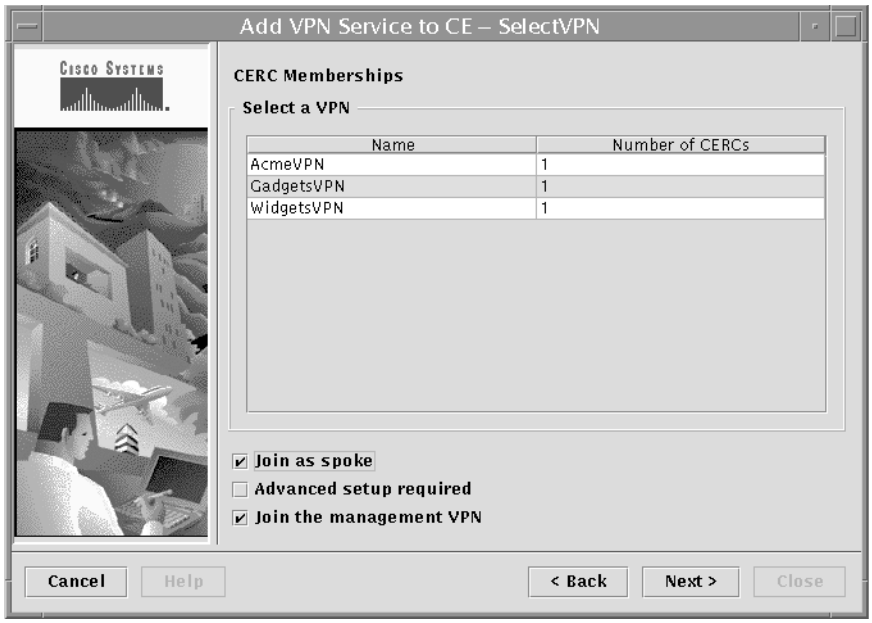

The most common types of VPNs are *hub-and-spoke* and *full mesh*. These two basic types of VPNs—full mesh and hub and spoke—can be represented with a single CERC.

For additional information on CE routing communities, see the "CE Routing Communities" section on page 1-18 and the "Defining CE Routing Communities" section on page 5-3.

Step 2 If you are building a VPN with a hub-and-spoke topology, check the **Join as Spoke** check box.

- A hub-and-spoke CERC is one in which one or a few CEs act as hubs, and all spoke CEs talk only to or through the hubs, never directly to each other.
- A full mesh CERC is one in which every CE connects to every other CE.

Step 3 If you are building a VPN with CEs that are members of multiple VPNs (extranets), check the **Advanced setup required** check box.

Extranet provisioning provides a way to create multiple VPN connectivity to a single VRF.

Step 4 When provisioning a cable service, we recommend that you check the **Join the management VPN** check box. For more information, see the "About Service Request Profiles" section on page 5-6.

When you use the VPN Solutions Center software to define a management VPN, the software automatically generates an *export route map* for the management VPN.

Step 5 When finished entering the necessary information, click **Next**. The Select Routing Policy dialog box appears (see Figure 9-21).

Specifying No Routing Protocol for the Cable Link

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

Step 1 From the list of routing protocol options, choose **None**.

Figure 9-21 Specifying No Routing Protocol

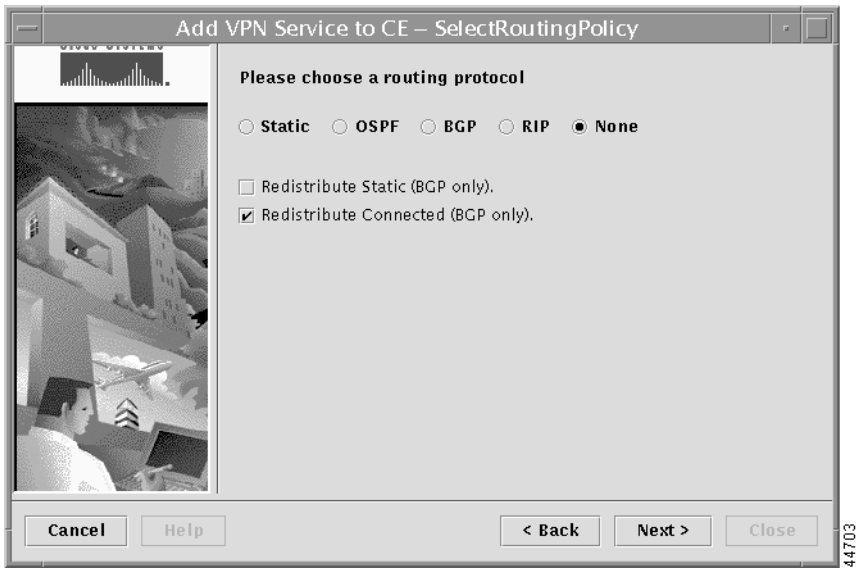

Step 2 If you want to distribute static routes into the provider core network (which runs BGP), check the **Redistribute Static (BGP only)** check box.

Step 3 Because there is no routing protocol on the cable link, we recommend that you redistribute the connected routes to all the other cable CEs in the VPN. To do so, check the **Redistribute Connected (BGP only)** check box.

Step 4 When finished entering the necessary information, click **Next**. The Select PE-CE Interface dialog box appears (see Figure 9-22).

Specifying the Cable Interface

You can now specify the interface on the PE that will host the cable subinterface.

Figure 9-22 Selecting the Cable Interface

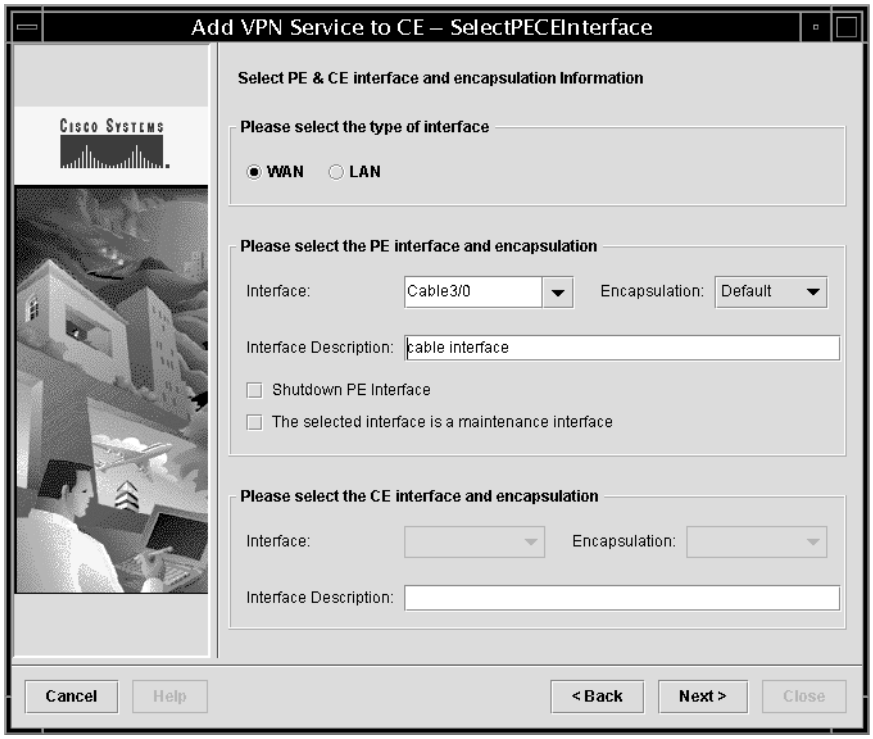

Step 1 Specify whether the interfaces for the PE-CE link are for a Wide Area Network (WAN) or Local Area Network (LAN).

Step 2 Select the interface on the PE that hosts the cable subinterface.

VPN Solutions Center software automatically selects the subinterface number based on the VRF. If a subinterface that is associated with the current VRF does not yet exist, VPNSC software creates a subinterface and assigns it to the correct VRF. The subinterface number is incremented to 1 greater than the largest subinterface currently assigned for the selected cable interface.

Tip Be sure to select the same interface that you chose for the maintenance subinterface. For example, if you chose *Cable 3/0* for the maintenance subinterface, choose *Cable 3/0* here as well.

The encapsulation method is set to *Default* for the cable interface.

Step 3 Optionally, you can enter a description of the interface for the cable link.

Step 4 Enable the **Shutdown PE Interface** option if desired:

When you check the **Shutdown PE Interface** checkbox, the specified PE interface will be configured in a shut down state.

Tip Be sure to *not* check the **Selected interface is a maintenance interface** check box. You should enable this option only when you want to provision the cable maintenance interface (see"Provisioning the Cable Maintenance Subinterface" section on page 9-9).

Step 5 When finished with these settings, click **Next**. The Select Cable Parameters dialog box appears (see Figure 9-23).

Specifying the Cable Helper Secondary Addresses

In this dialog box, you can specify the secondary addresses, if necessary.

Secondary addresses are IP addresses that are used for routing packets to the host devices connected to the cable modem. All the host devices on that cable subnet can use a single secondary address.

Figure 9-23 Setting the Cable Helper Addresses

Step 1 If desired, enter a secondary address by clicking **Add**.

The Secondary Addresses dialog box appears (see Figure 9-24).

Figure 9-24 Entering a Secondary Address

Step 2 In the *IP Address* fields, enter the secondary IP address, then click **Add**.

The address you entered appears in the display field.

Step 3 If the address is correct, click **OK**.

If you need to edit the secondary address, click **Modify**, then edit the address as necessary.

You return to the Select Cable Parameters dialog box.

Step 4 Click **Next**.

The Select IP Addresses dialog box appears (see Figure 9-25).

Specifying the IP Address for the Cable Subinterface

In the Select IP Addresses dialog box, you must specify the IP address for the cable subinterface on the PE.

If two or more cable modems belong to a particular ISP, the cable modems are connected to the same subnet on the PE, and that subnet is in the ISP's VPN.

Figure 9-25 Specifying the Address for the Cable Subinterface

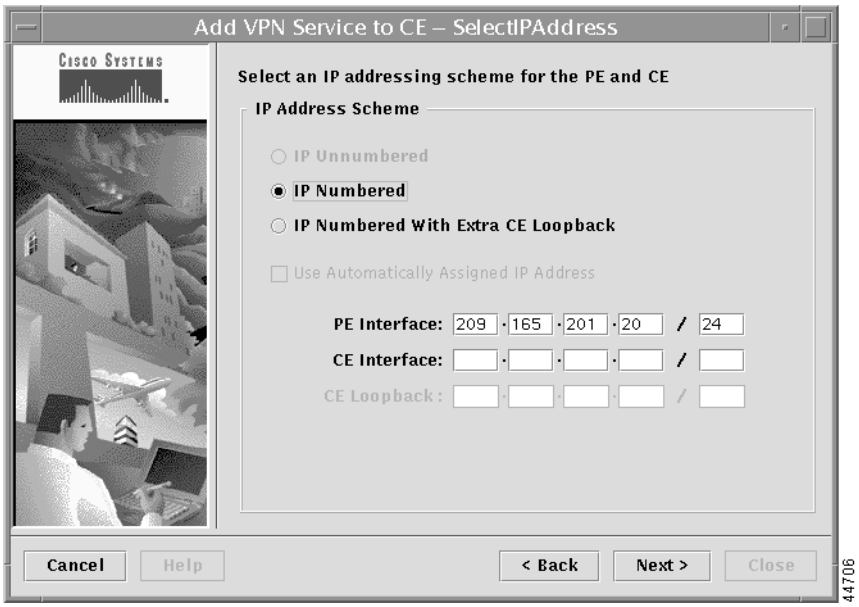

Step 1 For the IP addressing scheme, choose **IP Numbered**.

The *IP Numbered with Extra CE Loopback* option is not a viable option in a cable services configuration.

Step 2 In the **PE Interface** fields, enter the IP address on the PE for the subinterface that the cable is connected to.

Note The IP address entered here must be different from the IP address entered for the maintenance subinterface. To be reachable, each subinterface must have its own IP address.

You do not need to enter an IP address for the CE interface.

Step 3 When finished entering the necessary information, click **Next**. The Select VRF Parameters dialog box appears (see Figure 9-26).

Specifying the VRF Parameters

The Select VRF Parameters dialog box lets you set values for an import route map and the maximum number of routes in the VRF table. You can also enable NetFlow accounting.

Figure 9-26 Specifying the VRF Parameters

Step 1 In the *Import Map* field, enter the name of an existing import route map on the PE.

Note The Cisco IOS supports only one import route map per VRF (and therefore, per VPN).

An import route map applies 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 this PE to exclude the route.

For command reference details on the **import map** command, see the "import map" section on page B-4.

Step 2 In the *Maximum Routes* field, specify the maximum number of routes that can be imported into the VRF on this PE.

Step 3 To enable NetFlow accounting, check the **Turn on NetFlow accounting** checkbox.

Step 4 When you have completed the fields as necessary in the Specify VRF Parameters dialog box, click **Next**. The Class of Service (CoS) dialog box appears.

Selecting a Class of Service Profile

Step 1 If desired, select a Class of Service (CoS) profile to assign to the PE-CE link.

You can create a Class of Service (CoS) profile when you define the Provider Administrative Domain. For information on creating a CoS Profile, see the "Defining a Class of Service Profile" section on page 4-35.

Class of Service profiles are applied to the Provider Edge Router (PE), but the CoS definition is enforced across the PE-CE link on both the PE and CE.

Step 2 Click **Next**. The Confirm dialog box appears.

Confirming the Cable VPN Service Settings

The Confirm dialog box displays a summary of settings defined for this cable services VPN.

Step 1 Verify that the service request information is correct, then click **Next**. The wizard displays the following message:

Your request to "Add VPN Service to CE" has been submitted with ID number *n*. This service request can be deployed by using the "Deploy Service Requests" wizard or by using the "Deploy VPN Service" item under the "Provisioning" option of a VPN service request report.

Step 2 Press **Close**. You have now queued a service request. It is entered into the product database and is in the initial state of "Requested."

CHAPTER 10

Provisioning with the VPN Solutions Center Template Manager

The Template Manager

The Template Manager in the VPN Solutions Center software is a provisioning system that provides fast, flexible, and extensible Cisco IOS command generation capability. The Template Manager defines standard templates to generate Cisco IOS configuration files for common provisioning tasks, such as common IPv4, QoS, and VPN provisioning.

- A *template file* is a file created by the Template Manager that stores a VPN Solutions Center template definition.
- A *template data file* is a text file that stores variable values to generate the template file. A valid data file contains name-value pairs for all the variables defined in a template. Each template file can be associated with multiple data files; however, note that each data file can only be associated with a single template. You can select which data file to use to generate a template. The filename suffix for data files is *.dat*.
- A *template configuration file* is an IOS configuration file that stores the Cisco IOS commands created by the Template Manager. A template configuration file can be either a partial or complete configuration file. When you generate a template configuration file using a particular data file, the template configuration filename is the same as the data file's name.

The template data files are tightly linked with its corresponding template. You can use a data file and its associated template to create a template configuration file. The template configuration file is merged with (either appended to or prepended to) the VPNSC configlet. VPN Solutions Center downloads the combined configlet to the edge device router.

You can apply the same template to multiple edge devices, assigning the appropriate template data file for each device. Each template data file includes the specific data for a particular device (for example, the management IP address or host name of each device).

The template files and data files are in XML format. The template file, its data files, and all template configuration file files are mapped to a single directory.

- VPN Solutions Center creates the initial VPNSC configlet. Through the Template Manager, you can create a template configuration file. You can then associate a template configuration file with a service request, which effectively merges the VPNSC configlet and the template configuration file. For details on this process, see the "Templates" section on page 5-46.

  You can then download this merged VPNSC configlet to the target router (or routers).

Tip You can also create a template configuration file and download it directly to a router as described in the "Provisioning a Template Configuration File Directly to a Router" section on page 10-24.

Uses for the Templating Function

Service providers can use the Template Manager to enhance VPN Solutions Center functionality. Because the Telnet Gateway Server (TGS) supports console access to VPN Solutions Center targets, you can use the Template Manager to provide initial configuration for any service provider core device or edge device.

The Template Manager can be used as a stand-alone tool to generate complete configuration files that you can download to any VPN Solutions Center target.

Some of the additional uses for templating are as follows:

- IOS firewall provisioning
- Add a set of commands that VPN Solutions Center does not include to a service request; for example, provisioning ATM Class of Service.
- Use the templating feature to apply Class of Service using IP connectivity.
- Download a VPN Solutions Center service request and an Cisco IOS configuration file in one download operation through the console. This edge device staging method would create a template and apply the service request in one step.

Creating a Template for VPN Provisioning

To create a VPN Solutions Center template for VPN provisioning, follow these steps:

Step 1 Start VPN Solutions Center and bring up the VPN Console.

Step 2 From the VPN Console, choose **Tools > Template Console**.

The VPN Solutions Center Template Manager appears (see Figure 10-1).

Figure 10-1 The Template Manager Window

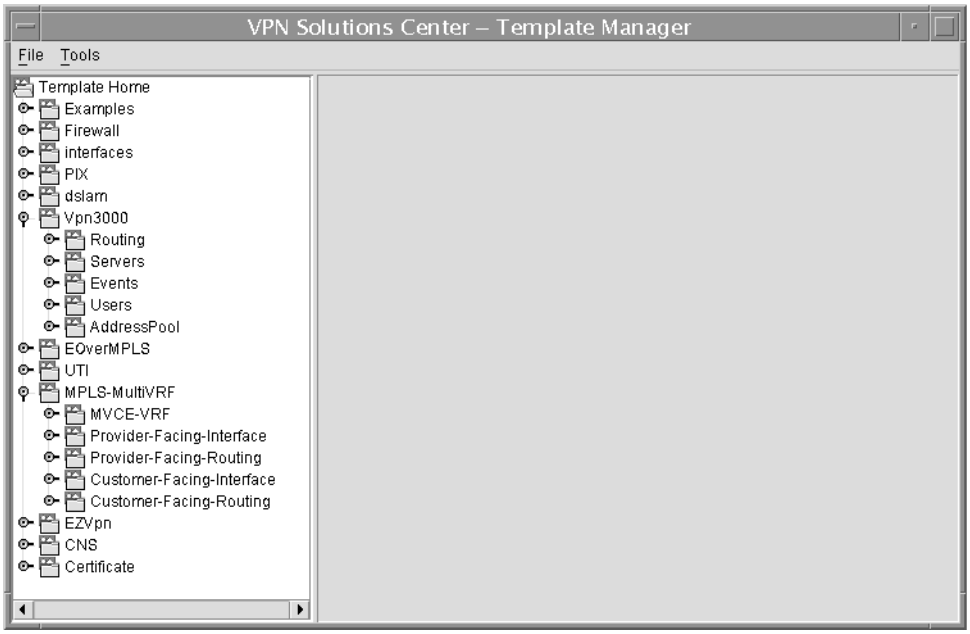

Template Home is an object that manages the life cycle and location of the Template objects.

Template folders logically organize templates into a hierarchy that facilitates navigation. As in a file system, the Template folders are like directories and the templates are equivalent to files. Many data files can be associated with a template, and any of the data files can be used to generate a template.

As you can see in Figure 10-1, the Template Manager provides a robust set of template examples. For a description of each of the template examples, see the "Example Templates Provided by VPN Solutions Center" section on page 10-50.

Step 3 To create a new folder at the top level, select the Template Home folder, then **right-click**.

The menu shown in Figure 10-2 appears.

Figure 10-2 Template Folders Menu

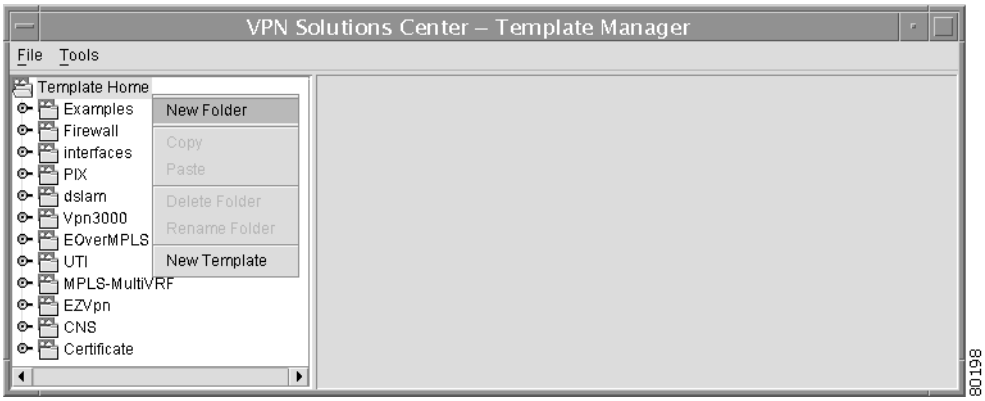

Step 4 From the Template Folders menu, choose **New Folder**.

The New Template Folder dialog box appears (see Figure 10-3).

Figure 10-3 New Template Folder Dialog Box

Step 5 Enter the name of the new template folder, then click **OK**.

The new template folder is added to the folders in the template tree.

Note You do not have to create a new folder to create a new template. You can select any template folder from which to create a new template.

Step 6 Select the folder where you want to place the new template, then **right-click**.

The Folders menu appears, but now additional commands are enabled, including **New Template** (see Figure 10-4).

Figure 10-4 Creating a New Template

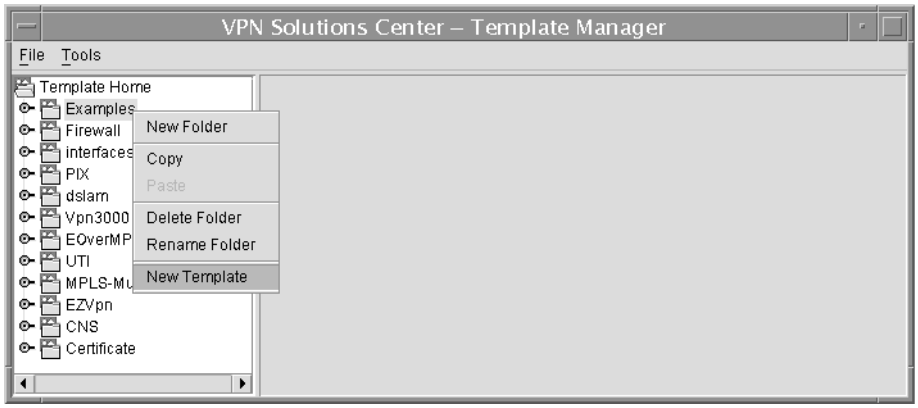

Step 7 From the menu, choose **New Template**.

The Template Editor appears in the data pane (see Figure 10-5).

Figure 10-5 The Template Editor

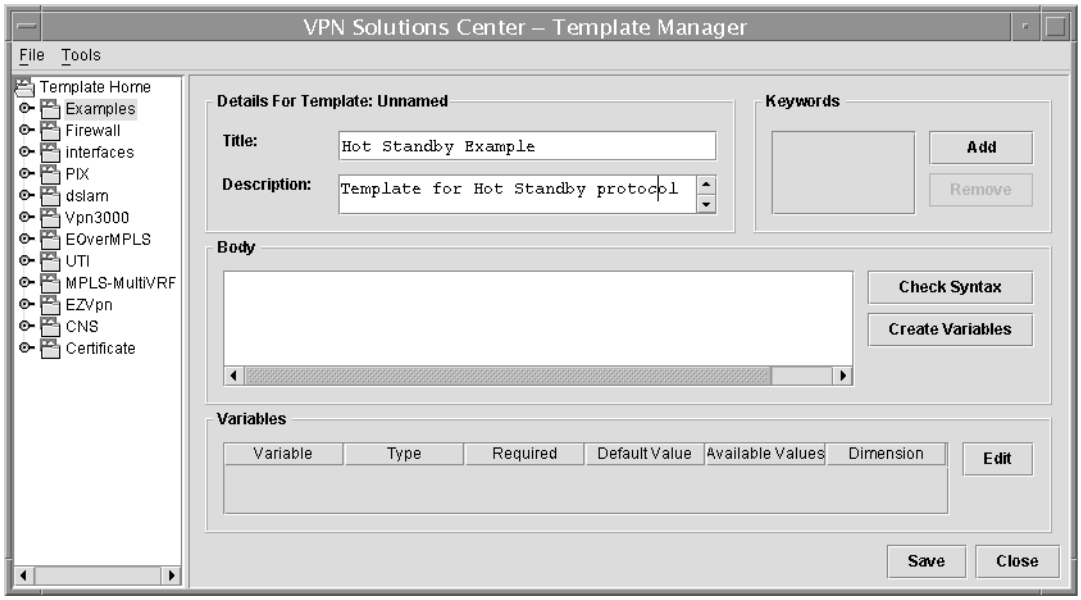

Step 8 Complete the *Title* , *Description*, and *Keywords* fields as described here.

a. *Title*: Enter a title for the new template.

The *Title* field is optional, but recommended. The title you enter here is not the name of the template; it is a high-level description of the template.

b. *Description*: enter a description of the template.

The *Description* field is optional, but recommended.

Step 9 *Keywords*: Enter one or more keywords in the *Keywords* area.

The keywords are to assist you in finding templates after they have been defined.

a. When you want to add a keyword, click **Add**.

   A new keyword field appears in the *Keywords* area.

b. Click into the *Keywords* area and enter the keyword.

Entering Configuration Commands in the Template Body

The Body area of the Template Manager is the place where you enter the configuration commands that you want to add to a configuration file. This area of the Template Manager allows you to enter the Cisco IOS commands, check the VPN Solutions Center template syntax of the command you have entered, and create the variables called by the IOS commands.

The result is a *template configuration file* that is added to the *VPNSC configlet* generated through standard VPN Solutions Center provisioning. VPN Solutions Center software downloads the combined template configuration file and VPNSC configlet to the target routers.

Note The use of interactive Cisco IOS commands is not supported in VPN Solutions Center templates.

For details on the VPN Solutions Center template language, see the "Template Language and Syntax Reference" section on page 10-39.

Step 1 In the Body area of the Template Manager, enter the Cisco IOS commands and the appropriate variables to create the template configuration file you need.

Tip For efficient data entry, you can copy existing commands from a terminal window and paste them (using **Ctrl-V**) into the Body area.

Step 2 When you have completed entering the IOS commands to your satisfaction, click **Create Variables**.

The Template Manager automatically checks the VPNSC template syntax—not the Cisco IOS syntax—and places the variables that were entered in the Body area into the Variables area (at the bottom of the dialog box)—see Figure 10-6.

If there is one or more errors in the template syntax, a syntax error message appears. The syntax error message states the type of error and the line on which the error occurs. Dismiss the error message and correct the syntax error.

Figure 10-6 Completed Template and Variables Displayed

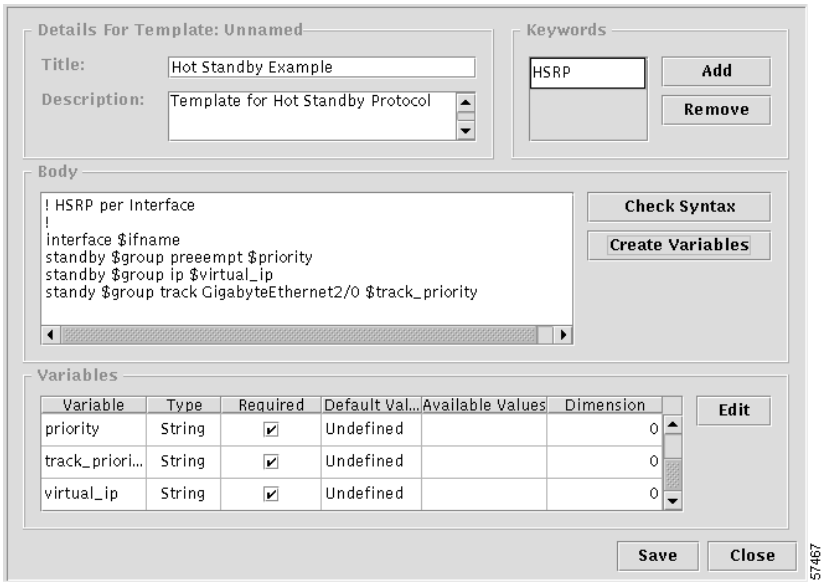

Step 3 Save the changes you have made by choosing **File > Save**.

The Save Template dialog box appears (see Figure 10-7).

Figure 10-7 Saving the Template Changes

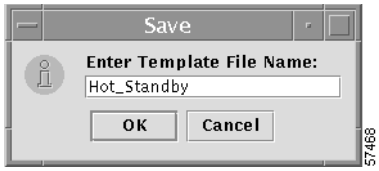

Step 4 Enter the template filename, then click **OK**.

When you save a template, the name of the template is added to the template tree under the appropriate template folder.

Assigning Attributes to Template Variables

Variables are strings that start with a $ sign. Variables are placeholders that are replaced with actual values from the template data files when the template is generated.

All variables entered into VPN Solutions Center templates are local variables. A variable declared in one template can only be used by that template.

To set attributes for the variables, do the following:

Step 1 From the Variables area in the Template Manager (see Figure 10-6 on page 10-7), select the variable you want to edit.

Step 2 Click **Edit**.

The Edit Variable dialog box appears (see Figure 10-8).

Figure 10-8 Edit Variable Dialog Box

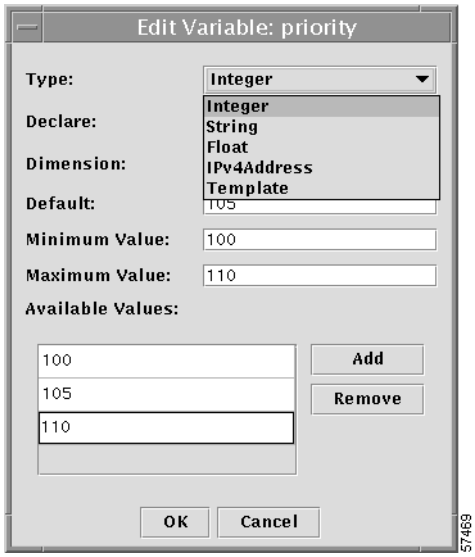

Setting the Variable Type

Step 3 From the Type drop-down menu, choose the variable type to be assigned to the selected variable.

As shown in Figure 10-8. there are five variable types you can choose from:

- Integer
- String
- Float
- IPv4 Address
- Template

When you select a variable type, the fields displayed in the Edit Variable dialog box change to accommodate the parameters you can set for the selected type. Some of the parameters are common to all the variable types, and those are described in the following steps.

For each variable type, there are a set of predefined attributes associated with it. For details on each of the variable types, see the "About the Variable Types and Their Attributes" section on page 10-9.

Declaring the Variable as Required or Optional

Step 4 From the Declare drop-down menu, specify whether the variable is *Required* or *Optional*.

- When you select **Required**, you must specify a value for this variable in the associated data file(s). This is the default setting.
- When you select **Optional**, you can omit the value for this variable in the associated data file(s).

Specifying the Dimension Attribute

Step 5 From the Dimension drop-down menu, specify the Dimension attribute.

The Dimension attribute is an optional attribute that creates an array (or list) of variables. The default value is 0, which indicates a single or enum variable.

If you set the Dimension attribute to **1** or **2**, the variable becomes a 1- or 2-dimension array. In this case, the constraint attributes are applied to all the elements in the array.

Specifying the Default Attribute

Step 6 In the *Default* field, specify the default value for the selected variable (if there is a default value). This is an optional attribute. If you also set the available values for this variable, the default value specified here must be one of the available values.

Specifying the Available Values

When you specify a set of available values for a variable, the values entered become the only values that are allowed to be assigned to the variable. If you also set a default value for this variable, the default value must be one of the available values.

Step 7 To specify the available values for the selected variable, follow these steps:

a. From the list of variables displayed in the Template Manager, select the variable you want to edit, then click **Edit**.

   The Edit Variable: *variable_name* dialog box appears.

b. Click **Add**.

   A new blank field is displayed in the Available Values area.

c. Enter the first available value.

d. To specify additional available values for the selected variable, click **Add**, then enter the additional available values until you have specified all the necessary values.

e. When finished entering the available values, click **OK**.

The values you entered in the Edit Variable dialog box are displayed for the selected variable in the Template Manager Variables area.

Step 8 Save your changes by clicking **Save**, then click **Yes** to accept the save operation.

About the Variable Types and Their Attributes

A variable is a symbol or name that stands for a value. VPN Solutions Center converts most variables to strings when the template is created, but you can set attributes for each variable type. For example, for an integer variable, you can set the minimum and maximum values allowed in the data.

As shown in Figure 10-8. there are five variables types you can choose from:

- Integer
- String
- Float
- IPv4 Address
- Template

Integer Variable Type

An integer is a whole number. The attributes included in the Integer variable type are as follows:

Table 10-1 Integer Variable Attributes

Attribute	Comments	Attribute	Comments
Name	A required attribute. The name must be unique.	**Minimum value**	Optional
Maximum value	Optional	**Available values**	Optional
Default	The default value if the operator does not provide data.	**Dimension**	The Dimension attribute of **1** or **2** creates an array (or list) of variables. Specify whether the integer variable is an array. The options are: **0**, **1**, or **2**.
Declaration	Specify whether it is a required or optional variable.		

String Variable Type

A string is a combination of characters considered as a group. The attributes included in the String variable type are as follows:

Table 10-2 String Variable Attributes

Attribute	Comments	Attribute	Comments
Name	A required attribute. The name must be unique.	**Minimum length**	Optional. If you specify a minimum length, the string cannot be less in length than the value specified here.
Maximum length	Optional. If you specify a maximum length, the string cannot be longer in length than the value specified here	**Available values**	Optional.
Default	The default value if the operator does not provide data.	**Dimension**	The Dimension attribute of **1** or **2** creates an array (or list) of variables. Specify whether the string variable is an array. The options are: **0**, **1**, or **2**.
Declaration	Specify whether it is a required or optional variable.	**Pattern**	The regular expression pattern of the string. For example, a pattern of **vpnsc[0-9]+** defines a string that starts with “vpnsc” followed by a number from 0 to 9.

Float Variable Type

A floating point number is a number that has no fixed number of digits before or after the decimal point. The attributes included in the Float variable type are as follows:

Table 10-3 Float Variable Attributes

Attribute	Comments	Attribute	Comments
Name	A required attribute. The name must be unique.	**Minimum value**	If you specify a minimum value, the values cannot be less than the value specified here.
Maximum value	If you specify a maximum value, the values cannot exceed the value specified here.	**Available values**	Optional
Default	The default value if the operator does not provide data.	**Dimension**	The Dimension attribute of **1** or **2** creates an array (or list) of variables. Specify whether the float variable is an array. The options are: **0**, **1**, or **2**.
Declaration	Specify whether it is a required or optional variable.		

IPv4 Address Variable Type

This variable type represents an IPv4 address. The default IPv4 Address variable is in string format, unless the operator provides the IP address data. Figure 10-9 shows an example of an IPv4 variable definition using the Edit Variable dialog box.

Figure 10-9 Example of an IPv4 Address Variable Definition

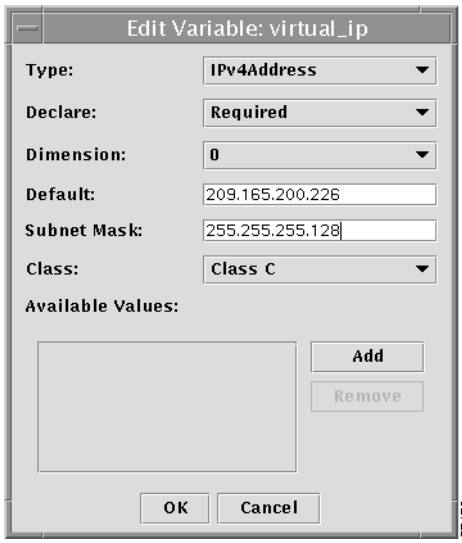

Note The *Subnet Mask* attribute is not supported in the VPN Solutions Center 2.x releases.

The attributes included in the IPv4 Address Variable type are as follows:

Table 10-4 IPv4 Address Variable Attributes

Attribute	Comments	Attribute	Comments
Name	A required attribute. The name must be unique.	**Class**	The class of the IP address. The options are: **A**, **B**, or **C**.
Subnet mask	The subnet mask of the IP address. This attribute is not supported in this release.	**Available values**	Optional. If desired, specify the addresses to be accepted as the available values.
Default	**String**: the default value if the operator does not provide data.	**Dimension**	The Dimension attribute of **1** or **2** creates an array (or list) of variables. Specify whether the address variable is an array. The options are: **0**, **1**, or **2**.
Declaration	Specify whether it is a required or optional variable.		

Template Variable Type

A template that is embedded within a template is called a *subtemplate*. This variable type specifies a subtemplate variable that can retrieve a subtemplate. VPN Solutions Center supports one level of subtemplates only. Thus, a subtemplate cannot include (or embed) another subtemplate.

Before you create a Template variable, you must create the subtemplate that the variable refers to. When you set the variable type to **Template**, you must browse for and specify the name and location of the indicated subtemplate.

The attributes included in the Template variable type are as follows:

Table 10-5 Template Variable Attributes

Attribute	Comments	Attribute	Comments
Name	The subtemplate name must be unique.	**Template location**	Specifies the location of the template in the VPNSC template hierarchy.

Creating a Template Data File Folder

You can create a folder to organize your template data files. Template data file folders logically organize data files into a hierarchy that facilitates navigation. The template data files are exclusively linked to the associated template.

To create a template data file folder:

Step 1 From the VPN Console menu bar, choose **Tools > Template Console**.

The Template Manager appears.

Step 2 In the Template Home hierarchy pane, expand the hierarchy until you can see the name of the template you want to create the new data file folder for.

Step 3 Select the template name, then **right-click**.

Step 4 From the menu, choose **New Folder** (see Figure 10-10).

Figure 10-10 New Data File Folder Option

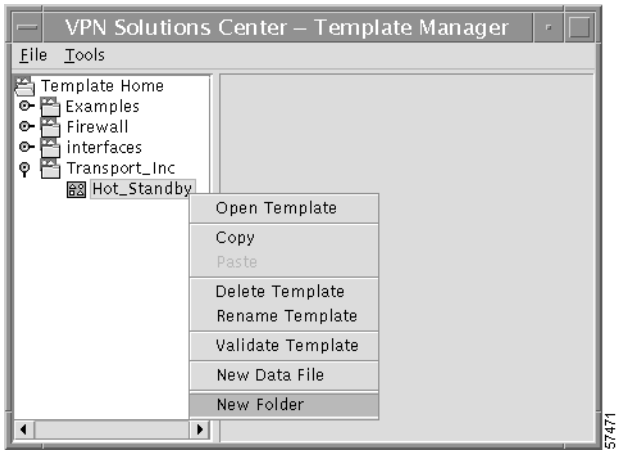

The New Data File Folder dialog box appears (see Figure 10-11).

Figure 10-11 Entering the Name of the New Data File Folder

Step 5 Enter the name of data file folder, then click **OK**.

The new data file folder is added to the template hierarchy (see Figure 10-12).

Figure 10-12 New Data File Folder Added

Creating a New Template Data File

A *template data file* is a text file that stores the variable values necessary to generate a template file. A valid template data file contains name-value pairs for all the variables defined in a template.

Each template can be associated with multiple data files; however, note that each data file can only be associated with a single template. You can select which template data file to use to generate a template.

To create a new template data file, follow these steps:

Step 1 From the VPN Console menu bar, choose **Tools > Template Console**.

The Template Manager appears.

Step 2 In the Template Home hierarchy pane, expand the hierarchy until you can see the name of the template you want to create the new data file for.

Step 3 Select the template name, then **right-click**.

Step 4 From the menu, choose **New Data File** (see Figure 10-13).

Figure 10-13 New Data File Menu Option

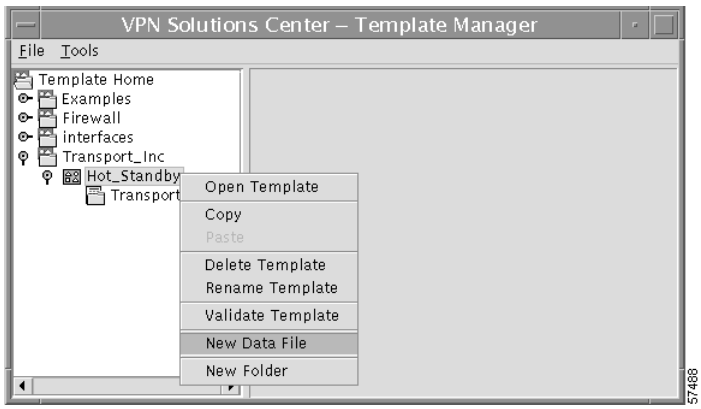

The new Template Data Files dialog box appears in the data pane on the right (see Figure 10-14).

In the Data Files area of the dialog box, the name of each of the variables appears with a cell beneath it for the designated value of that variable.

Figure 10-14 New Template Data Files Dialog Box

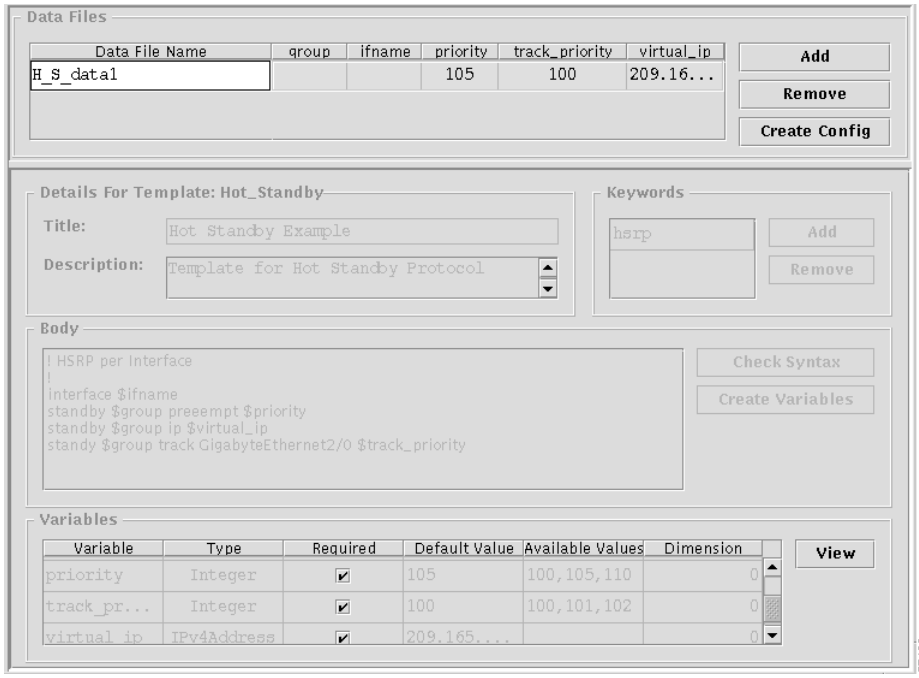

Step 5 If desired, change the name of the data file from the default name (*Data0*) to a more identifiable name.

a. Click the *Data File Name* cell.

b. Change the data file name as desired.

Step 6 Enter the appropriate data for each of the variables called in the Template body by clicking the cell below each variable name and entering the data.

The Template Manager displays the variables in standard spreadsheet format. The initial default values are displayed for each variable.

Depending on the type of variable, and its settings and attributes, the data values you enter will vary. The Template Manager provides type checking when you save the data; for example, if you enter non-numeric characters in a field defined as an integer, the Template Manager does not accept the data and alerts you to the discrepancy.

As shown in the track_priority variable in Figure 10-15, you can see the available values set for a particular variable by clicking the name of the variable.

Figure 10-15 A Variable's Available Values Listed

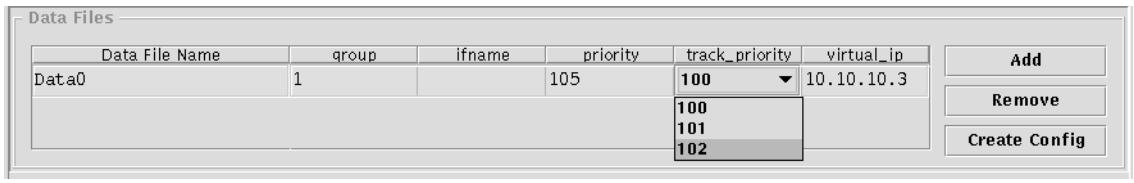

Entering Values for a One-Dimensional Array

A one-dimensional array creates a list of values that you can apply in your template for that variable. The number of values that can be substituted for any given one-dimensional variable is determined by the number of values you define for it (as shown below in Figure 10-17 on page 10-17).

The following is an example of program statements entered into a template that define a one-dimensional array:

```
!
{
access-list $ACL[$i] permit ip $Source_IP_Address[$i] $Source_Mask[$i]
$Dest_IP_Address[$i] $Dest_Mask[$i] precedence 7
}
```

The **$i** specified in the variables declared in the substatements above such as:

```
$Source_IP_Address[i]
```

indicates that it is a one-dimensional variable. For example, for each value specified for the **Source_IP_Address** variable, the template substitutes one of the values for **Source_IP_Address**.

To enter the values for a one-dimensional array:

Step 1 If the Template Data Files dialog box is not already open, expand the Template Manager hierarchy until you can see the name of the data file of interest.

Step 2 Select the template data file, then **right-click**.

Step 3 From the menu, choose **Open**.

The Template Manager Data Files dialog box is displayed (see Figure 10-16).

Figure 10-16 Selecting One-Dimensional Variable

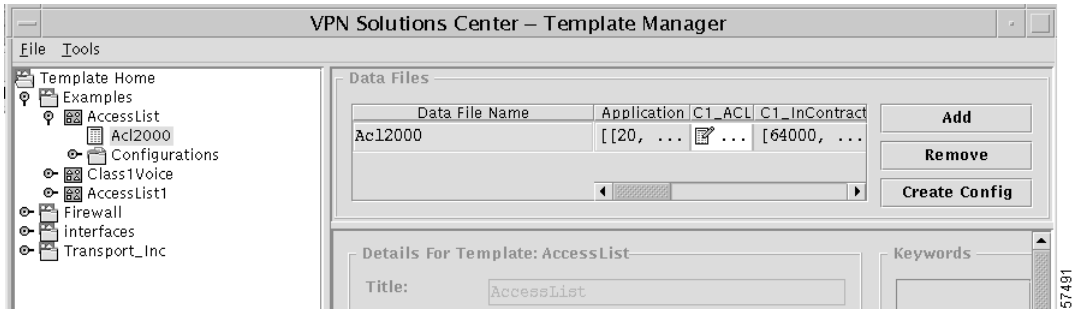

The name of the selected template data file and each of its variables are displayed. You can use the mouse to adjust the width of each variable cell.

Step 4 Scroll to the cell that displays the name of the variable you want to edit.

Step 5 Select the variable cell, then **double-click**.

The Edit Variable: *variable_name* dialog box appears (Figure 10-17).

Figure 10-17 Entering Data for a One-Dimensional Array Variable

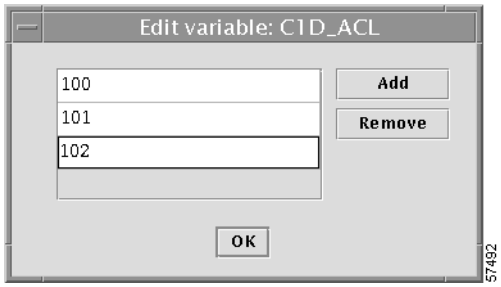

Step 6 In the Edit Variable dialog box, click **Add**.

A new field appears in the editing area.

Step 7 Enter the first value to be assigned to the variable.

Step 8 Click **Add**, then enter the appropriate values for each additional value to be assigned to the variable.

To delete a value from the list, select the value, then click **Remove**.

Step 9 When the data is entered to your satisfaction, click **OK**.

The values you entered are displayed in the corresponding Data Files variable cell, as shown in Figure 10-18.

Figure 10-18 One-Dimensional Array Added to Data File Variable Display

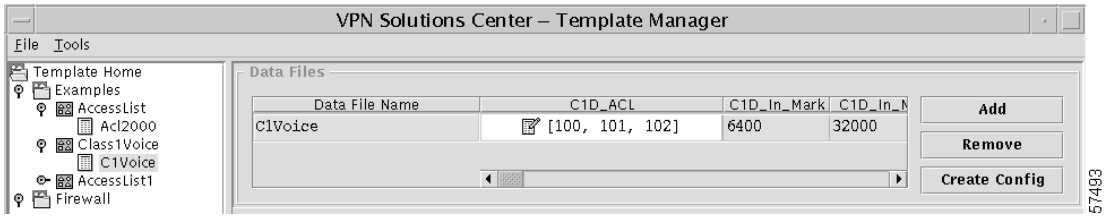

Step 10 To save your work, choose **File > Save**, then when prompted to save the file, click **Yes**.

Entering Values for a Two-Dimensional Array

A two-dimensional array creates a table of values that you can apply in your template for that variable. The number of values that can be substituted for any given two-dimensional variable is determined by the number of rows and columns that you define in the Edit Variable dialog box.

The indexing convention uses pairs of numbers that refer to positions in the indexing table, as shown in Table 10-6:

Table 10-6 Indexing Convention for a Two-Dimensional Array

0, 0	0, 1	0, 2
1, 0	1, 1	1, 2
2, 0	2, 1	2, 2

The rows and columns that you create with the Template Manager Variable Editor creates a table of cells that correspond to the indexing convention for two-dimensional variables. Substituting values occurs by linking the positions indicated by the indexing convention with two-dimensional variable statements.

For example, let us say that you have a Protocol variable that has four possible values: IP, TCP, RIP, and UDP. The example of the Edit Variable dialog box shown in Figure 10-19 shows these four protocol values in their relative positions.

Figure 10-19 Entering a Two-Dimensional Array

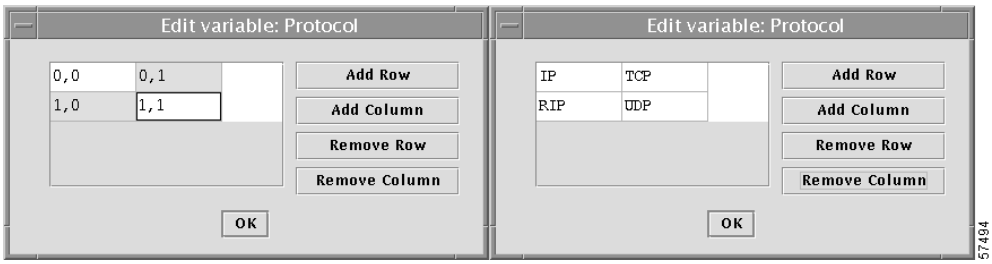

- "IP" is in the 0, 0 index position
- "TCP" is in the 0, 1 index position
- "RIP" is in the 1, 0 index position
- "UDP" is in the 1, 1 index position

To substitute the appropriate variable values, the variables defined in the template body would be defined as follows:

- To substitute the **IP** value: **Protocol [0] [0]**
- To substitute the **TCP** value: **Protocol [0] [1]**
- To substitute the **RIP** value: **Protocol [1] [0]**
- To substitute the **RIP** value: **Protocol [1] [1]**

To enter the values for a two-dimensional array, follow these steps:

Step 1 If the Template Data Files dialog box is not already open, expand the Template Manager hierarchy until you can see the name of the data file of interest.

Step 2 Select the template data file, then **right-click**.

Step 3 From the menu, choose **Open**.

The Template Manager Data Files dialog box is displayed. The name of the selected template data file and each of its variables are displayed. You can use the mouse to adjust the width of each variable cell.

Step 4 Scroll to the cell that displays the name of the variable you want to edit.

Step 5 Select the variable cell, then **double-click**.

The Edit Variable: *variable_name* dialog box appears (Figure 10-20).

Figure 10-20 Editing Two-Dimensional Array

When the Edit Variable dialog box for a two-dimensional array first appears, there are no fields displayed in the edit area.

Step 6 To begin entering data, click **Add Row**.

A row appears, as shown in Figure 10-20.

Step 7 To add a column to the variable table, click **Add Column**.

Step 8 Add rows and columns as necessary to build a variable index table that provides the number of cells you need to hold the values required by the variable.

You can also remove unneeded rows and columns by selecting a cell in the row or column you want to remove and clicking **Remove Row** or **Remove Column**.

Step 9 Click each cell to enter the variable values as necessary.

Step 10 When the data is entered to your satisfaction, click **OK**.

The values you entered are displayed in the corresponding Data Files variable cell, as shown in Figure 10-21.

Figure 10-21 Two-Dimensional Array Added to Data File Variable Display

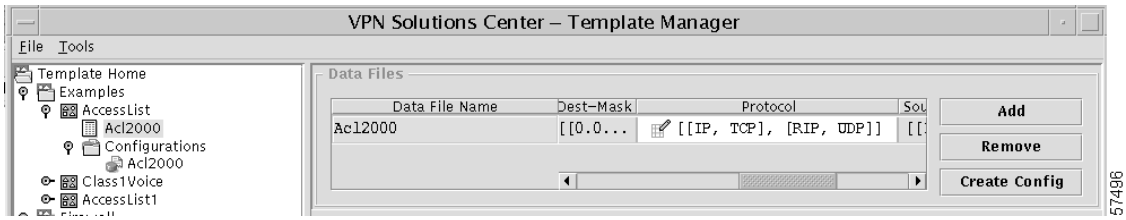

Step 11 To save your work, choose **File > Save**, then when prompted to save the file, click **Yes**.

Copying a Template Data File

You can copy a template data file from one location to another within the Template Manager. Keep in mind that template data files are tightly coupled with their associated template. Therefore, you would normally copy a template data file to the same template, then modify the variables accordingly.

To copy a template data file, follow these steps:

Step 1 Expand the template hierarchy until you can see the name of the data file you want to copy.

Step 2 Select the data file name, then **right click**.

Figure 10-22 Copying a Template Data File

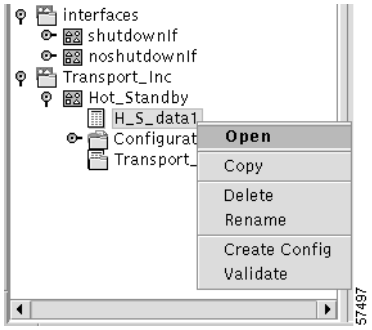

Step 3 From the menu, choose **Copy**.

Step 4 Select the template folder (or the data file folder) that you want to copy the data file to.

Step 5 From the menu, choose **Paste**.

The template data file is copied to the selected folder.

Deleting a Template Data File

You can delete a template data file from the Template Manager. Keep in mind that template data files are tightly coupled with their associated template. Therefore, deleting a data file can effectively disable a template.

To delete a template data file:

Step 1 Expand the template hierarchy until you can see the name of the data file you want to delete.

Step 2 Select the data file name, then **right click**.

Step 3 From the menu, choose **Delete**.

You receive the following warning:

Warning: Are you sure that you want to delete <filename> file?

Step 4 To delete the selected data file, click **Yes**.

To cancel the data file deletion operation, click **No**.

Creating a Template Configuration File

You can create a template configuration file based upon a particular template and a specific data file.

To create a template configuration file:

Step 1 From the VPN Console menu bar, choose **Tools > Template Console**.

The Template Manager appears.

Step 2 In the Template Home hierarchy pane, expand the hierarchy until you can see the name of the template and then the name of the associated template data file that you want to create the template configuration file from.

Step 3 Select the template data file of interest, then **right-click**.

Step 4 From the menu, choose **Open**.

The Template Data Files dialog box appears.

Step 5 Select the row where the data file name and variables are displayed (see Figure 10-23).

Figure 10-23 Selecting the Template Data File

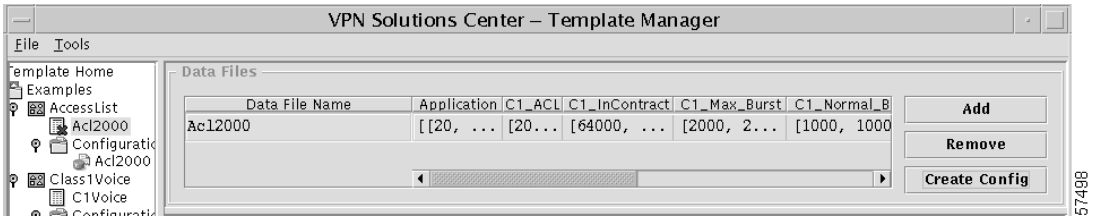

Step 6 Click **Create Config**.

You receive the following message:

The Data File will be saved before creating a configuration.

Step 7 Click **OK** to proceed.

The Configuration Created window appears, which displays the template configuration file that was created from the selected template and configuration file (see Figure 10-24).

Figure 10-24 Template Configuration File Displayed

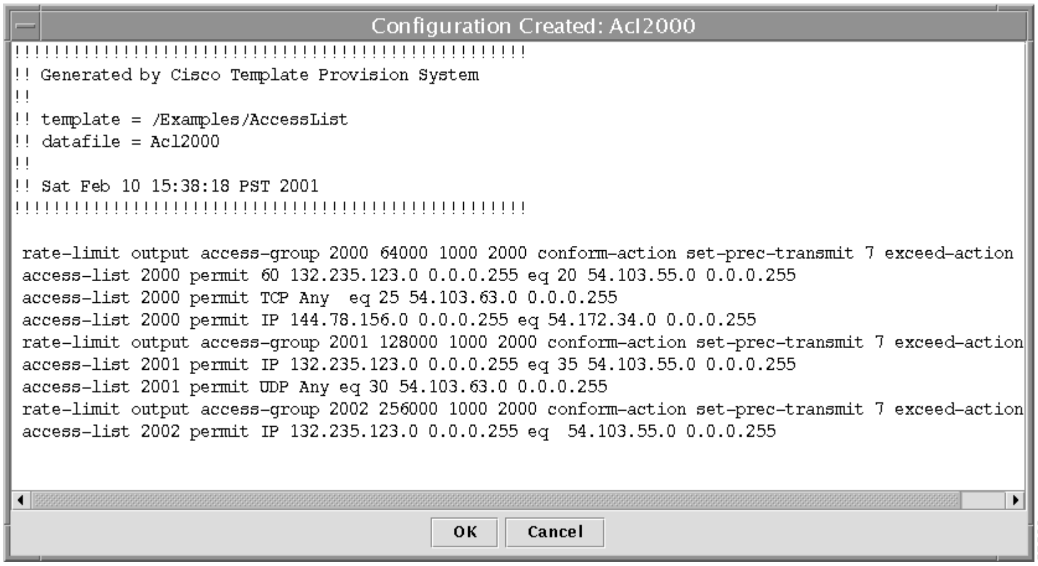

```
Configuration Created: Acl2000
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!! Generated by Cisco Template Provision System
!!
!! template = /Examples/AccessList
!! datafile = Acl2000
!!
!! Sat Feb 10 15:38:18 PST 2001
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

 rate-limit output access-group 2000 64000 1000 2000 conform-action set-prec-transmit 7 exceed-action
 access-list 2000 permit 60 132.235.123.0 0.0.0.255 eq 20 54.103.55.0 0.0.0.255
 access-list 2000 permit TCP Any  eq 25 54.103.63.0 0.0.0.255
 access-list 2000 permit IP 144.78.156.0 0.0.0.255 eq 54.172.34.0 0.0.0.255
 rate-limit output access-group 2001 128000 1000 2000 conform-action set-prec-transmit 7 exceed-action
 access-list 2001 permit IP 132.235.123.0 0.0.0.255 eq 35 54.103.55.0 0.0.0.255
 access-list 2001 permit UDP Any eq 30 54.103.63.0 0.0.0.255
 rate-limit output access-group 2002 256000 1000 2000 conform-action set-prec-transmit 7 exceed-action
 access-list 2002 permit IP 132.235.123.0 0.0.0.255 eq  54.103.55.0 0.0.0.255

OK  Cancel
```

The Template Manager creates a folder called "*Configurations*" in the appropriate template folder, and places the configuration file in the Configurations folder.

Step 8 To close the configuration file window, click **OK**.

Copying a Template

You can copy a VPN Solution Center template from one folder to another within the template hierarchy. When you copy a template, all the files and folders associated with that template—its data files and configuration files—are copied to the new location.

To copy a VPN Solution Center template:

Step 1 Expand the template hierarchy until you can see the name of the template you want to copy.

Step 2 Select the template name, then **right click**.

Step 3 From the menu, choose **Copy**.

Step 4 Select the template folder that you want to copy the template to.

Step 5 From the menu, choose **Paste**.

The template, along with the files and folders associated with it, is copied to the selected folder.

Deleting a Template

You can delete any template from the VPN Solutions Center Template Manager. When you delete a template, all the files and folders associated with that template—its data files and configuration files—are also deleted.

To delete a VPN Solution Center template, follow these steps:

Step 1 Expand the template hierarchy until you can see the name of the template you want to delete.

Step 2 Select the template name, then **right click**.

Step 3 From the menu, choose **Delete Template**.

You receive the following warning:

WARNING: All files in the selected template directory will be deleted. Do you want to continue?

Step 4 To delete the selected template, click **Yes**.

To cancel the template deletion operation, click **No**.

The template, along with the files and folders associated with it, is deleted from the Template Manager.

Provisioning a Template Configuration File Directly to a Router

You can download a template configuration file directly to a router. To do so, follow these steps:

Step 1 From the VPN Console menu bar, choose **Tools > Template Console**.

The Template Manager appears.

Step 2 In the Template Home hierarchy pane, expand the hierarchy until you can see the name of the template and then the name of the associated template configuration file that you want to download to a router.

Step 3 Select the template configuration file of interest, then **right-click** (see Figure 10-25).

Figure 10-25 Template Configuration File Menu

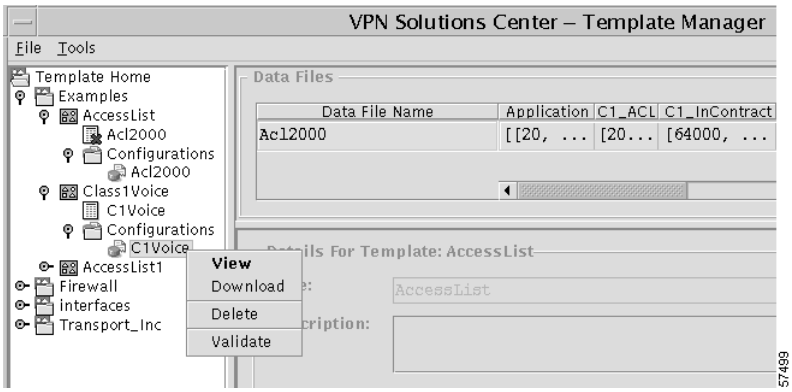

Step 4 From the menu, choose **Download**.

The VPN Solutions Center Network window appears (see Figure 10-26).

Figure 10-26 Selecting Router from the VPNSC Network Window

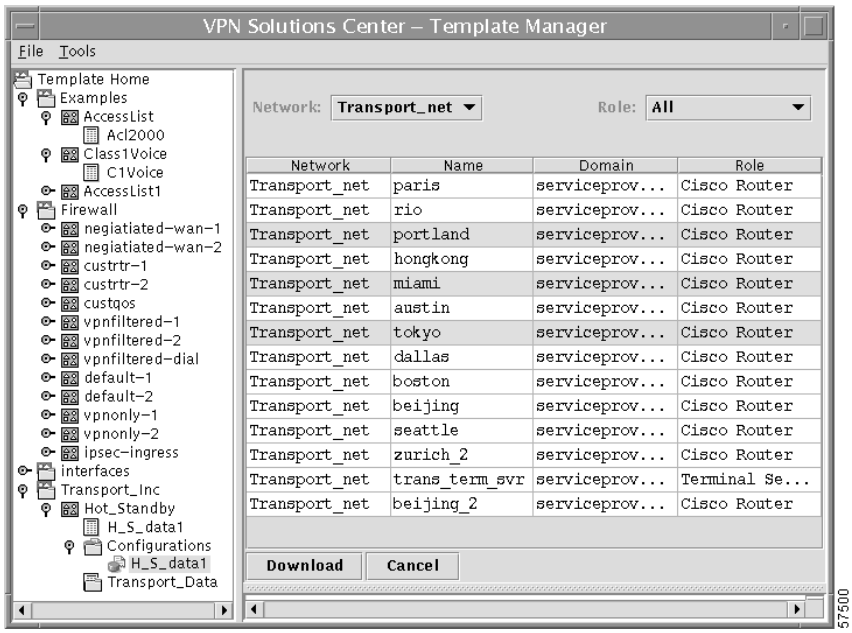

Step 5 From the Network window, specify the following:

a. From the Network drop-down menu, choose the appropriate network.

b. From the Role drop-down menu, choose **Cisco Router**.

c. From the list of routers, select one or more routers that you want to download the template configuration file to.

Step 6 When satisfied with your selections, click **Download**.

An informational window appears, which shows the status of the template configuration file download operation (see Figure 10-27).

Figure 10-27 Template Configuration File Download Status Window

Using VPN Solutions Center Repository Variables as Template Data

You can choose from a list of IPsec or MPLS variables that are part of the VPN Solutions Center service definition. Their data values are substituted when you associate a template with a service request. For details on this procedure, see the "Templates" section on page 5-46.

Caution When you use the Repository variables in a template, you cannot use the Template Manager to download the template configuration file to the router. This is because the data values for that variable are substituted when the VPN Solutions Center provisions a service request. Since this process bypasses the service request, the appropriate data is not substituted for the assigned Repository variable.

To use VPN Solutions Center Repository values as template data:

Step 1 From the VPN Console menu bar, choose **Tools > Template Console**.

Step 2 From the Template Manager hierarchy pane, expand the template hierarchy until you can see the name of the data file of interest.

Step 3 Select the name of the pertinent data file, then **right-click**.

Step 4 From the menu, choose **Open**.

The Data Files dialog box appears.

Step 5 Click the cell for the variable of interest, then **right-click**.

Step 6 From the menu, choose **MPLS**. The list of MPLS service values is displayed (see Figure 10-28).

Figure 10-28 Menu of MPLS Repository Variables

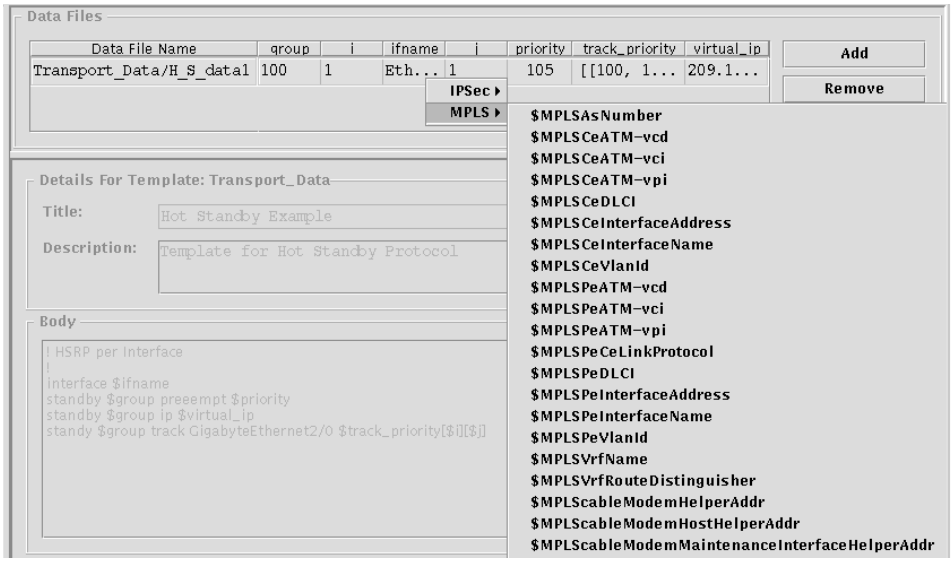

Step 7 From the list of variable values, choose the appropriate value.

Summary of MPLS Repository Variables

Table 10-7 provides a summary of the MPLS Repository variables available from the VPN Solutions Center Template Manager.

Table 10-7 MPLS Repository Variables

Repository Variable	Dimension	Description	Example
MPLSAsNumber	0	Domain autonomous system (AS) number	100
MPLScableModemHelperAddr	0	Cable modem helper address	209.165.201.1
MPLScableModemHostHelperAddr	0	Cable modem host helper address	209.165.201.2
MPLScableModemMaintenanceInterface HelperAddr	0	Cable modem maintenance interface helper address	209.165.201.3
MPLScableModemSecondaryIpAddrList	1	Cable modem secondary IP address list	209.165.201.4
MPLSCeATM	0	CE router ATM virtual circuit descriptor	200
MPLSCeATM-vci	0	CE router virtual circuit identifier	210
MPLSCeATM-vpi	0	CE router virtual path identifier	220
MPLSCeBGPASNumber	0	CE router BGP AS number	100
MPLSCeDLCI	0	CE router data link circuit identifier	100
MPLSCeInterfaceAddress	0	CE router interface address	209.165.201.5
MPLSCeInterfaceName	0	CE router interface name	Serial0
MPLSCeVlanId	0	CE router virtual LAN identifier	755
MPLSPeCeLinkProtocol	0	PE-CE link routing protocol	BGPRoutingType
MPLSPeATM-vcd	0	PE router ATM virtual circuit descriptor	200
MPLSPeATM-vci	0	PE router virtual circuit identifier	210
MPLSPeATM-vpi	0	PE router virtual path identifier	220
MPLSPeBGPASNumber	0	PE router BGP AS number	
MPLSPeDLCI	0	PE router ATM link circuit identifier	100
MPLSPeInterfaceName	0	PE router interface name	Serial0
MPLSPeInterfaceAddress	0	PE router interface address	209.165.201.6
MPLSPeOSPFArea	0	PE router OSPF area number	500
MPLSPeOSPFProcessId	0	PE router OSPF process ID number	75
MPLSPeVlanId	0	PE router virtual LAN identifier	755
MPLSVrfRouteDistinguisher	0	MP-BGP attribute: VRF route distinguisher	100:100
MPLSVrfName	0	VPN routing and forwarding table name	V2:Widgets-S

Summary of IPsec LAN-to-LAN Repository Variables

Table 10-8 provides a summary of the IPsec LAN-to-LAN Repository variables available from the VPN Solutions Center Template Manager.

Table 10-8 IPsec LAN-to-LAN Repository Variables

Repository Variable	Dimension	Description	Example
IPSecLocalProtectedIPAddress	1	List of IP address ranges protected by this edge device	209.165.202.129 209.165.202.130...
IPSecLocalProtectedIPAddressMask	1	List of IP address ranges and subnet masks protected by this edge device	209.165.20.129/30 209.165.20.130/30...
IPSecLocalProtectedIPAddress ReverseMask	1	List of reverse subnet masks of IP address ranges protected by this edge device	209.165.202.129 209.165.202.130...
IPSecLocalProtectedIPMask	1	List of subnet masks of IP address ranges protected by this edge device	209.165.20.129/30 209.165.20.130/30...
IPSecRoutingProtocol	0	List of IPsec VPN routing protocols	NONE, OSPF
IPSecSecureInterfaceAddress	1	List of IP addresses for CPE's secured interfaces	209.165.202.129 209.165.203.129
IPSecSecureInterfaceAddressMask	1	List of IP addresses and subnet masks for CPE's secured interfaces	209.165.20.130/30 209.165.21.130/30
IPSecSecureInterfaceAddress ReverseMask	1	Edge device's secured interface reverse subnet mask	0.0.0.3
IPSecSecureInterfaceMask	1	Edge device's secured interface subnet mask	255.255.255.252
IPSecSecureInterfaceName	1	List of CPE's secured interface names	Serial0/0, Serial0/1
IPSecSecureLoopbackInterfaceAddress	0	Edge device's secured tunnel endpoint interface IP address	171.23.45.33
IPSecSecureLoopbackInterfaceAddress Mask	0	Edge device's secured tunnel endpoint interface IP address and mask	171.23.45.33/32
IPSecSecureLoopbackInterfaceAddress ReverseMask	0	Edge device's secured tunnel endpoint interface reverse subnet mask	0.0.0.0
IPSecSecureLoopbackInterfaceMask	0	Edge device's secured tunnel endpoint interface subnet mask	255.255.255.255
IPSecSecureLoopbackInterfaceName	0	Edge device's secured tunnel endpoint interface name	Loopback0
IPSecUnsecureInterfaceAddress	1	List of edge device's nonsecured interface IP addresses	209.165.202.131 209.165.203.131
IPSecUnsecureInterfaceAddressMask	1	List of edge device's nonsecured interface IP addresses and subnet masks	192.23.45.33/24 192.23.46.33/24
IPSecUnsecureInterfaceAddressReverse Mask	1	Edge device's nonsecured interface reverse subnet mask	0.0.0.255

Table 10-8 IPsec LAN-to-LAN Repository Variables (continued)

Repository Variable	Dimension	Description	Example
IPSecUnsecureInterfaceMask	1	Edge device's nonsecured interface subnet mask	255.255.255.0
IPSecUnsecureInterfaceName	1	List of CPE's nonsecured interface names	Ethernet1, Ethernet0
IPSecUnsecureLoopbackInterface Address	0	List of CPE's nonsecured loopback interface IP addresses	10.1.93.1 10.1.94.1
IPSecUnsecureLoopbackInterface AddressMask	0	Edge device's nonsecured loopback interface IP address and mask	10.1.93.1/32 10.1.94.1/32
IPSecUnsecureLoopbackInterface AddressReverseMask	0	Edge device's nonsecured loopback interface reverse subnet mask	0.0.0.0
IPSecUnsecureLoopbackInterfaceMask	0	Edge device's nonsecured loopback interface subnet mask	255.255.255.255
IPSecUnsecureLoopbackInterfaceName	0	List of nonsecured loopback interface names	Loopback1 Loopback0
IPSecRemoteProtectedIPAddress	2	List of IP address ranges protected by remote peers	209.165.202.129 209.165.202.130...
IPSecRemoteProtectedIPAddressMask	2	List of IP address ranges and subnet masks protected by remote peers	209.165.20.129/30 209.165.20.130/30...
IPSecRemoteProtectedIPAddressReverse Mask	2	List of reverse subnet masks of IP address ranges protected by remote peers	0.0.0.255...
IPSecRemoteProtectedIPMask	2	List of subnet masks of IP address ranges protected by remote peers	209.165.20.129/30 209.165.20.130/30...
IPSecRemoteHostName	1	List of remote peer host names	cpe1-frankfurt

Summary of IPsec Remote-Access Repository Variables

Table 10-9 provides a summary of the IPsec Remote-Access Repository variables available from the VPN Solutions Center Template Manager.

Table 10-9 IPsec Remote-Access Repository Variables

Repository Variable	Dimension	Description	Example
$RA-AAServerNameList	1	List of authentication server names	North_Am_AA
$RA-GroupNameList	1	List of Group names	North_AM_Sales
$RA-LocalProtectedIPAddressMaskList	1	List of IP address ranges and subnet masks protected by this edge device	209.165.20.129/30 209.165.20.130/30...
$RA-SecureInterfaceIPAddressMaskList	1	List of secured interface IP addresses and their masks	192.209.10.10/30 192.209.11.10/30
$RA-SecureInterfaceNameList	1	List of CPE's secured interface names	Serial0/0, Serial0/1
$RA-SplitTunnelingNetworkLists	2	List of split-tunneling networks	List 0 (Sales Group) 10.1.1.0/24 10.1.2.0/24
$RA-SplitTunnelingTypeList	1	List of split-tunneling types	Entry 0(Acct Group) in-list Entry1(Mkt Group) in-list
$RA-UnsecureInterfaceIPAddressMaskList	1	List of unsecured interface IP addresses and their subnet masks	171.23.44.33/24 171.23.45.33/24
$RA-UnsecureInterfaceNameList	1	List of unsecured interface names	Ethernet0, Ethernet1

Using Ethernet Over MPLS in VPN Solutions Center

Ethernet Over MPLS (EoMPLS) is a form of a Layer 2 VPN technology that sends Ethernet frames over an MPLS core network. The EoMPLS technology extends the Ethernet LAN using point-to-point links between VPN sites. Unlike MPLS VPN, Ethernet over MPLS does not require routing from the CE to the PE—it requires only VLAN connectivity. VPN Solutions Center supporting EoMPLS in the this release using the Template Manager.

The EoMPLS service is a point-to-point service, much like Frame Relay, but based on Ethernet technology. VPNSC provisions EoMPLS one PE at a time using the VPNSC Template Manager. VPNSC requires that the MPLS LDP and core routing protocols are active for label distribution. The template can create multiple VLAN interfaces per PE to provision multiple EoMPLS end points on the PE. VPNSC templates are not designed to provision the CE unless you want to download a separate template for the CEs in the network.

Example of Configlet Generated

This section shows the configlet that would be generated by the Template Manager for OSR 1.

```
OSR1
mpls label protocol ldp
vlan 25
!
vlan 60
!
interface Loopback0
  ip address 10.1.1.1 255.255.255.255
  no ip route-cache
  no ip mroute-cache
!
interface GigabitEthernet4/1
  ip address 10.4.1.2 255.255.0.0
  no negotiation auto
  tag switching ip
!
interface GibabitEthernet4/2
switch port
switchport trunk vlan 25
!
interface GibabitEthernet4/3
switch port
switchport trunk vlan 60
!
interface Vlan25
  no ip address
  mpls 12transport route 10.1.1.2 1021
!
router ospf 100
network 10.1.0.0 0.0.255.255 area 0
```

The VPNSC Ethernet Over MPLS Template

The variables in the VPNSC template reference the tunnel end points PE1 and PE2. In the VPNSC Template Manager, the operator inputs information into the VPNSC template data file. The combination of the template data file and the template body generates the configuration.

The key variables in the template are one-dimensional arrays for the VLAN ID and loopback IP addresses. If the operator has an array of dimension N for these variables, you can provision N+1 PE end points by consecutively downloading the template to each PE. The data is the same for all PE's, and only a minor modification is needed for each template data file. Table 10-10 lists all the variables that are used in the Ethernet Over MPLS template body.

Table 10-10 Ethernet Over MPLS Template Variables

Variable Name	Description
$cust_intf_name	Name of the interface facing the customer with the VLAN.
$cust_vlan_num	VLAN ID facing the customer. This variable is a one-dimensional array.
$endpt	End point of the point-to-point link a value of 1 or 2. If this variable is 1, VPNSC provisions the first endpoint (or PE1). If the value is 2, VPNSC provisions the second endpoint (or PE2).
$intf_loop_num	Number of the loopback interface.
$MPLS_Vlan_id	VLAN ID to the MPLS core. Used for connectivity between the PEs. This variable is a one-dimensional array.
$PE1_loopback_addr	PE1 IP address for loopback interface. This interface is a tunnel end point. This variable is a one-dimensional array.
$PE2_loopback_addr	PE2 IP address for loopback interface. This interface is a tunnel end point. This variable is a one-dimensional array.

EoMPLS Template Contents

This section provides the contents for the Ethernet Over MPLS template.

```
mpls label protocol ldp
#if($endpt #eq 1)
{
!
      interface loopback $intf_loop_num
      ip address    $PE1_loopback[$i] 255.255.255.255
!
      #repeat ( $MPLS_Vlan_id, $i)
      {
      vlan $cust_Vlan_num[$i]
!
      interface VLAN $cust_Vlan_num[$i]
      mpls l2transport route $PE2_loopback[$i]  $MPLS_Vlan_id[$i]
      no shutdown
!
      interface $cust_intf_name[$i]
         switchport
         switchport trunk vlan $cust_Vlan_num[$i]
         no shutdown
!
      }
!
      router ospf 1
      network $PE1_loopback[$i] 0.0.0.0 area  10
      }

#else
      {
#if($endpt #eq 2)
      {
!
      interface loopback $intf_loop_num
         ip address   $PE2_loopback[$i] 255.255.255.255
!
      #repeat ($MPLS_Vlan_id, $i)
      {
      vlan $cust_Vlan_num[$i]
!
      interface VLAN $cust_Vlan_num[$i]
      mpls l2transport route $PE1_loopback[$i] $MPLS_Vlan_id[$i]
      no shutdown
!
      interface $cust_intf_name[$i]
         switchport
         switchport trunk vlan $cust_Vlan_num[$i]
         no shutdown
      }
!
```

```
      }
!
      router ospf 1
      network $PE2_loopback[$i] 0.0.0.0 area  10
!
      }
```

Configuration Examples

This section provides an example of provisioning EoMPLS using the Template Manager. The folder in the Template Manager hierarchy is named **EOverMPLS**; the template body name is **EthernetOverMPLS**. In the examples below, the Template Manager uses the template data files EndPoint1 and EndPoint2 to provision OSR1 and OSR2.

OSR1 and Data File EndPoint1

This configuration example has two end points with multiple VLANs between OSR CPEs.

The VPNSC template data file EndPoint1 provisions the configlet for the first PE—OSR1.

Note Be sure to set the dimension to 1 for the array variables.

Table 10-11 Variable Values for the EndPoint1 Template Data File

Variable Name	Value
$cust_intf_name	GigabitEthernet4/2, GigabitEthernet4/3
$cust_vlan_num	25, 60
$endpt	1
$intf_loop_num	0
$MPLS_Vlan_id	50, 1021
$PE1_loopback_addr	143.10.0.1, 143.10.0.1
$PE2_loopback_addr	143.20.0.1, 143.20.0.1

Configlet for OSR1 and the EndPoint1 Data File

Here is the generated configlet for the OSR1 edge device using the EndPoint1 template data file.

```
!Configlet - OSR 1
interface loopback 0
ip address 143.10.0.1 255.255.255.255
!
vlan 25
!
interface VLAN 25
```

```
mpls l2transport route 143.20.0.1 50
no shutdown
!
interface GigabitEthernet4/2
switchport
switchport trunk vlan 25
no shutdown
!
vlan 60
!
interface VLAN 60
mpls l2transport route 143.20.0.1 1021
no shutdown
!
interface GigabitEthernet4/3
switchport
switchport trunk vlan 60
no shutdown
!
router ospf 1
network 143.10.0.1 0.0.0.0 area 10
```

OSR2 and Data File EndPoint2

VPN Solutions Center template data file EndPoint2 provisions the configlet for the second PE—OSR2.

Table 10-12 Variable Values for the EndPoint2 Template Data File

Variable Name	Value
$cust_intf_name	GigabitEthernet4/2, GigabitEthernet4/3
$cust_vlan_num	25, 60
$endpt	2
$intf_loop_num	0
$MPLS_Vlan_id	50, 1021
$PE1_loopback_addr	143.10.0.1, 143.10.0.1
$PE2_loopback_addr	143.20.0.1, 143.20.0.1

Configlet for OSR2 and the EndPoint2 Data File

Here is the generated configlet for the OSR2 edge device using the EndPoint2 template data file.

```
!Configlet - OSR 2
     mpls label protocol ldp
!
     interface loopback 0
        ip address 143.20.0.1 255.255.255.255
!
     vlan 25
!
     interface VLAN 25
        mpls l2transport route 143.10.0.1 50
        no shutdown
!
     interface GigabitEthernet4/2
        switchport
        switchport access 25
        no shutdown
     vlan 60
!
     interface VLAN 60
        mpls l2transport route 143.20.0.1 1021
        no shutdown
!
     interface GigabitEthernet4/3
        switchport
        switchport access 60
        no shutdown
!
     router ospf 1
     network 143.20.0.1 0.0.0.0 area 10
!
```

Ethernet Over MPLS Removal Template

This section describes the template used to remove commands from the router(s) that were generated by the Ethernet Over MPLS template.

EoMPLS Template Contents

This section provides the contents for the Ethernet Over MPLS Removal template.

```
#if($endpt #eq 1)
{
!
    interface loopback $intf_loop_num
    ip address $PE1_loopback[$i] 255.255.255.255
!
    #repeat ( $MPLS_Vlan_id, $i)
       {
    no vlan $cust_Vlan_num[$i]
!
    no interface VLAN $cust_Vlan_num[$i]

!
     interface $cust_intf_name[$i]
       no switchport
       no switchport trunk vlan $cust_Vlan_num[$i]
       shutdown
!
    }
!
    router ospf 1
    no network $PE1_loopback[$i] 0.0.0.0 area  10
    }

#else
    {
#if($endpt #eq 2)
    {
!
    interface loopback $intf_loop_num
       ip address   $PE2_loopback[$i] 255.255.255.255
!
       #repeat ($MPLS_Vlan_id, $i)
       {
       no vlan $cust_Vlan_num[$i]
!
       no interface VLAN $cust_Vlan_num[$i]

!
    interface $cust_intf_name[$i]
       no switchport
       no switchport trunk vlan $cust_Vlan_num[$i]
```

```
       shutdown
    }
!
    }
!
    router ospf 1
    no network $PE2_loopback[$i] 0.0.0.0 area  10
!
    }
```

Table 10-13 Variable Values for the EoMPLS Removal Template

Variable Name	Value
$cust_intf_name	GigabitEthernet4/2, GigabitEthernet4/3
$cust_vlan_num	25, 60
$endpt	1
$intf_loop_num	0
$MPLS_Vlan_id	50, 1021
$PE1_loopback_addr	143.10.0.1, 143.10.0.1
$PE2_loopback_addr	143.20.0.1, 143.20.0.1

Configlet Generated for the EoMPLS Removal Template

This section provides the configlet generated for the EoMPLS removal template.

```
!EoMPLS Removal Configlet – OSR 1
       interface loopback 0
          ip address 143.10.0.1 255.255.255.255 !
          no vlan 25
!
       no  interface VLAN 25
!
       interface GigabitEthernet4/2
          no switchport
          no switchport trunk vlan 25
          shutdown
!
       no  vlan 60
!
       no interface VLAN 60
!
       interface GigabitEthernet4/3
          no switchport
          no switchport trunk vlan 60
          shutdown
!
    router ospf 1
```

```
    no network 143.10.0.1 0.0.0.0 area 10
!
```

Template Language and Syntax Reference

This section describes the language and syntax conventions used in the VPN Solutions Center template implementation.

Grammar and Syntax

The Extensible Markup Language (XML) template definition section consists of a sequence of IOS command lines, a sequence of control statements and some template comments.

This section describes in detail the template language grammar and syntax. For specifying the syntax of the template language, Cisco used a widely-used notation, called context-free grammars (Bachus-Naur Format). For specifying the semantics of the template language, we use informal descriptions and examples.

Argument

```
argument →
    variable = variable
    | variable = constant
    | variable = array_variable
    | variable = IOS_command
```

Argument List

```
argument_list →
    argument
    | argument_list, argument
```

Array Variable

```
array_variable →
    variable[index]
    | array_variable[index]
```

Constant

```
constant →
    num
    | float
    | quoted_string
```

Expression

```
expression →
    term
    | (expression)
    | expression relational_opr expression
```

```
    | expression logical_opr expression
    | expression additive_opr expression
    | expression multiplicative_opr expression
```

Index

```
index →
    variable
    | num
```

IOS Command Expression

```
IOS_command_expression →
    IOS_command
    | expression
    | IOS_command_expression expression
    | IOS_command_expression IOS_command
```

Repeat Counter Variable

```
repeat_counter_variable →
    $(letter | digit)*
```

Repeat Variable

```
repeat_variable →
    variable
    | num
    | array_variable
```

Statement

```
statement →
    IOS_command_expression
    |IOS_comment
    |Template_comment
    |{statement_list }
    |#if ( expression ) statement
    | #if ( expression ) statement #else statement
    | #repeat ( repeat_variable, repeat_counter_variable ) statement
```

Statement List

```
statement_list →
    statement
    |statement_list statement
```

Template

```
template →
    statement_list
```

Template Variable

```
template_variable →
    variable ( argument_list )
    | variable ( )
```

Term

```
term →
    variable
    | constant
    | array_variable
    | template_variable
```

Variable

```
variable →
    $(letter | digit)*
```

Lexical Conventions

This section summarizes the lexical conventions and the notations used in the VPN Solutions Center template grammar:

Tokens

There are six classes of tokens: IOS commands, variables, keywords, constants, operators and other separators.

Blanks, tabs, new lines and comments (collectively referred to as "white space") as described below are ignored except when they separate tokens. Some white space is required to separate otherwise adjacent keywords and constants.

Template Comments

Template comments are indicated by '**//**'; everything after the comment indicator is treated as a template comment. Template comments are not included in the template configuration file—they are discarded during the template generation process.

Token Variables

A token variable is a sequence of letters and digits. The first character must be a letter '**$**'; it is followed by letters or digits:

letter a to z; A to Z

digit 0 to 9

Upper and lowercase letters are treated as being different. Variables can be up to 32 characters in length.

Keywords

Keywords are reserved and appear in boldface. The following keywords are reserved and cannot be used in any other context:

Table 10-14 Template Keywords

#and	#else	#eq (equal to)	#ge (greater than or equal to)
#gt (greater than)	#if	#le (less than or equal to)	$lt (less than)
#ne (not equal)	#or	#repeat	

Constants

There are several kinds of constants, as described below:

- A Token **num** constant consists of a sequence of digits:

  num digit digit*

- A Token **float** constant consists of three parts: integer, a decimal point, and a fraction. The integer and fraction parts both consist of a sequence of digits:

  float num.num

- A Token **quoted_string** is a sequence of characters surrounded by double quotes, as in "….".

Expressions

The expressions used in VPN Solutions Center templates are as follows:

- IOS commands, variables
- Array_variables
- Template_variables
- Constants
- Expressions in parentheses

The precedence of an expression operator is the same as the order that the expressions are documented in this section (that is, multiplicative operators first, followed by additive operators, and so on), with the highest precedence first.

1. **Multiplicative Operators**

   The operators include * and /, which group left-to-right. The operands of * and / must be of arithmetic type. The binary * operator denotes multiplication. The binary / operator yields the quotient of the division of the first operand by the second.

2. **Additive Operators**

   The operators include + and – , which group left-to-right. The result of the + operator is the sum of the operands. The result of the – operator is the difference of the operands.

3. **Relational Operators**

   The relational operators group left-to-right. The operators are:

   - **#lt** (less than)
   - **#gt** (greater than
   - **#le** (less than or equal to)
   - **#ge** (greater than or equal to)
   - **#eq** (equal to)
   - **#ne** (not equal to)

   All of the relational operators yield 0 if the specified relation is false, and yield 1 if it is true. The type of the result is an integer.

4. **Logical Operators**

   The logical operators group left-to-right. They are as follows:

   - **#or**

     The **#or** operator returns 1 if both its operands compare unequal to zero, and 0 otherwise.

   - **#and**

     The **#and** operator returns 1 if either of its operands compare equal to zero, and 0 otherwise. The operands need to be the same arithmetic type and the result is an integer.

5. **Array Operators**

   The array operators are as follows:

   - One-dimensional operator: [a]
   - Two-dimensional operator: [a] [b]

Statements

Except when described otherwise, statements are executed and output in sequence. Statements are executed for their effect and they do not have values. They fall into the following several groups:

- IOS Command Expression Statement

  Most statements in the template are **IOS command expression** statements, which are Cisco IOS router commands.

- Compound Statement

  A compound statement (also called a “block”) is used so that several statements can be used where one is expected.

- Selection Statements

  The selection statements have two forms. In both forms of the **#if** statement, the expression, which must be an arithmetic type, is evaluated.

  If it compares unequal to 0, the first substatement is executed.

  In the second form, if the expression is 0, the second (**#else**) substatement is executed.

- Iteration Statement

  The iteration (**repeat**) statement specifies a looping operation. Depending on how many values are specified for the **repeat** variable, the **#repeat** statement executes the substatement as many times as there are values defined for the variable.

For example, the following **repeat** statement would cause the **access-list** substatement to repeat as many times as there are values specified for the **ACL** (one-dimensional) variable.

The **$i** specified is the *repeat counter variable*. This repeat counter variable is set to 0 initially; the repeat counter is incremented by 1 for each iteration of the statement.

```
!
#repeat ($ACL, $i)
{
access-list $ACL[$i] permit ip $Source_IP_Address[$i] $Source_Mask[$i]
$Dest_IP_Address[$i] $Dest_Mask[$i] precedence 7
}
```

- The **$i** specified in the repeat statement is the counter for the repeat statement. This counter variable is set to 0 initially; the counter is incremented by 1 for each iteration of the statement.

  For example, the **repeat** statement in the example above would cause the **access-list** substatement to repeat as many times as there are values specified for the **ACL** (one-dimensional) variable. This produces in the template configuration file one line for each iteration.

- The **$i** specified in the variables declared in the substatements indicates that it is a one-dimensional variable. For example, for each value specified for the **Source_IP_Address** variable, the template substitutes one value for **Source_IP_Address**.

About If-Else Statements

The **if-else** statement is used to make decisions. Its syntax is as follows:

#if (*expression*)

Statement-1

#else

Statement-2

where the **#else** and its subsequent statement, *Statement-2*, are optional.

- The **#if** *expression* is evaluated. If the result is **true**, that is if the value of the *expression* is nonzero, *Statement-1* is executed.
- If the result is **false**, that is, if the value of the *expression* is zero and if there is an **else** part, *Statement-2* is executed instead.

Resolving if-else Relationships

Because the **else** part of an **if-else** is optional, an ambiguity may occur when an **else** part is omitted from a nested **if** sequence. This is resolved in the usual way—the **else** is associated with the closest previous **if** statement without an associated **else** statement. In the following example, the **else** goes with the inner **if**, as shown by the indentation.

```
#if ($n $gt 0)
    #if ($a #le $b)
        Statement-1
    #else
        Statement-2
```

If you do not want the **else** to go with the inner **if**, you must use braces to show the proper association. In the following example, the **else** goes with the outer **if**.

```
#if ($n #gt 0)
(
    #if ($a #le $b)
        Statement-1
)
#else
    Statement-2
```

About Subtemplates

All templates can be used by other templates as building blocks. The template that uses other templates is called a *super template*, and the template being used is called a *subtemplate*.

A super template generates a subtemplate by passing values for the variables in the subtemplate. After generating a subtemplate, a super template puts the configlets generated by the subtemplate into the super template.

The values for the variables are passed into a subtemplate by using a C-style function call. The parameters are delimited by commas and all the parameters are treated as a string.

See the following three examples.

Example 1—Simple Case: Pass Each Value, One-by-One

This subtemplate passes each value, one-by-one. In this example, the subtemplate called *SampleTemplate1* declares the following variable:

```
<VarDeclaration>
    <String VarName="sl"/>
    <Integer VarName="bw"/>
</VarDeclaration>
```

To generate this template, enter the following:

```
$SampleTemplate1($sl="abc",$bw=31)
```

Example 2—One Required and One Optional Value

This example presents a subtemplate called *SampleTemplate2* that declares one required variable and one optional variable.

- Case 1:

  ```
  <VarDeclaration>
      <String VarName="sl"/>
      <Integer VarName="bw" declare="optional"/>
  </VarDeclaration>
  ```

  To generate this template, enter the following:

  ```
  $SampleTemplate2($sl="abc")
  ```

- Case 2:

  ```
  <VarDeclaration>
  ```

```
    <String VarName="sl" declare="optional"/>
    <Integer VarName="bw"/>
</VarDeclaration>
```

To generate this template, enter the following:

```
$SampleTemplate2($bw=31)
```

Example 3—Calling Two Subtemplates From the Main Template

This example shows how the main template can call multiple subtemplates.

The syntax for calling subtemplate variables is as follows:

$subtemplateVariable ($variableInSubtemplate=$variableInMainTemplate,...)

Here is an example:

Template_Main

The template body for Template_Main is:

```
$subVar1 ($protocol=$tcp)
$subVar2 ($CE-lo="192.168.1.1," $mgt-prefix="192.168.2.0", $mgt-mask="255.255.255.0")
```

Template_Sub1

The template body for Template_Sub1 is:

```
access-list 104 permit $protocol
```

Template_Sub2

The template body for Template_Sub2 is:

```
access-list 103 permit host $CE-lo $mgt-prefix $mgt-mask
```

Calling Subtemplates

1. Specify the template location for $subVar1 to point to Template_Sub1.
2. Specify the template location for $subVar2 to point to Template_Sub2.

$tcp is a variable defined in Template_Main. So if the value of $tcp is set to "tcp", the configlet output for Template_Main would be as follows:

```
access-list 104 permit tcp
access list 103 permit host 192.168.1.1 192.168.2.0 255.255.255.0
```

Template Built-in Function Calls

This section describes the built-in function calls that the VPN Solutions Center Template Manager provides.

Retrieving a Substring

This template function call retrieves a substring from srcString using delimChar as the separator:

```
#system.substringToDelim (srcString, delimChar, from beginning:0/
from end:1)
```

Examples

`#system.substringToDelim ("abc#def", "#", 0)` returns "abc"

`#system.substringToDelim ("abc#def", "#", 1)` returns "def"

`#system.substringToDelim ("abc#def", "d", 0)` returns "abc#"

`#system.substringToDelim ("abc#def", "d", 1)` returns "ef"

Specifying a Different Comment Character

If you need to download a configuration file to a device other than IOS router, this template function call sets the specified comment character:

```
#system.setCommentChar (commentChar)
```

Example

```
#system.setCommentChar ( : )
```

The Template Manager will use the colon character (**:**) as the comment character.

Retrieving the IP Address From the IpAddrMaskPair String

This template function call retrieves the IP address part of the IpAddrMaskPair string:

```
#system.getAddr (IpAddrMaskPair)
```

Example

#system.getAddr ("10.33.4.5/30") returns "**10.33.4.5**"

Retrieving the IP Mask From the IpAddrMaskPair String

This template function call retrieves the IP mask part of the IpAddrMaskPair string:

```
#system.getMask (IpAddrMaskPair)
```

Example

`#system.getMask ("10.33.4.5/30")` returns "**255.255.255.252**"

Retrieving the IP Reverse Mask From the IpAddrMaskPair String

This template function call retrieves the IP reverse mask part of the IpAddrMaskPair string:

```
#system.getReverseMask (IpAddrMaskPair)
```

Example

`#system.getReverseMask ("10.33.4.5/30")` returns "**0.0.0.3**"

Retrieving the Network Address From the IpAddrMaskPair String

This template function call retrieves the IP network address part of the IpAddrMaskPair string:

```
#system.getNetworkAddr (IpAddrMaskPair)
```

Example

`#system.getNetworkAddr ("10.33.4.5/30")` returns "**10.33.4.4**"

Retrieving the Classful Network Address From the IpAddrMaskPair String

This template function call retrieves the classful network address from the IpAddrMaskPair string:

```
#system.getClassfulNetworkAddr (IpAddrMaskPair)
```

Example

`#system.getClassfulNetworkAddr ("10.33.4.5/30")` returns "**10.0.0.0**"

Template Language Directives

This section describes the template language directives that the VPN Solutions Center Template Manager provides.

Displaying the Error Description in Configlet Output

This template directive displays the error description string in the configlet output.

```
#system.error (error_description_string)
```

Template

```
Hostname myHost
#if (#varA #eq 1)
#system.error ("Some error occurred.")
#else
   ip addr 10.33.4.5 255.255.255.0
```

Output

- If $varA is equal to 1, the configlet output will be as follows:

  ```
  Hostname myHost
  ERROR: Some error occurred.
  ```

- If $varA is not equal to 1, the configlet output will be as follows:

  ```
  Hostname myHost
  ip addr 10.33.4.5 255.255.255.0
  ```

Aborting Configlet Generation

This template directive instructs VPN Solutions Center to abort configlet generation.

```
#system.abort (abort_description_string
```

Example Templates Provided by VPN Solutions Center

The VPN Solutions Center Template Manager provides a robust set of example templates. Table 10-15 provides summary descriptions of each of the example templates.

Table 10-15 Example Templates Provided in VPN Solutions Center

Folder	Template	Description
Examples	AccessList	Demonstrates templates with nested repeat loop and multi-dimension variable.
	AccessList1	Demonstrates the simplest template variable substitution.
	CEWanCOS	Demonstrates if-else statements, repeat statements, mathematical expressions, and one-dimensional variables.
interface	shutdownIf	Shuts down all unsecured interfaces of a given device in the IPsec Service Request.
	noshutdownIf	Brings up all nonsecureded interfaces of a given device in the IPsec Service Request.
Firewall	cust-qos	Permits a) IPsec: Authentication Header (AH), b) Encapsulating Security Payload (ESP), and c) Internet Security Association and Key Management Protocol (ISAKMP) to enter the firewall. Permits a) management traffic to access an IPsec VPN router, b) return management traffic from an IPsec router, c) Internet Control Message Protocol (ICMP) traffic to a VPN router, and d) all return traffic from an outbound session. Denies all other traffic.
	cust-rtr1	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits a) management traffic to access an IPsec VPN router, b) ICMP traffic to access a VPN router, andc) remote site traffic to access a VPN router. Denies access for local network traffic to a VPN router. Permits customer network traffic to access the Internet. Denies other local network traffic to access the Internet.

Table 10-15 Example Templates Provided in VPN Solutions Center (continued)

Folder	Template	Description
	cust-rtr2	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: a) management traffic to access an IPsec VPN router, b) Simple Network Management Protocol (SNMP) traffic to access a VPN router, c) application traffic to access a VPN router, and d) ICMP traffic to access a VPN router. Denies any other IP traffic access to a VPN router. Permits: a) all IP traffic to pass through a VPN router, b) and customer network to access a VPN router. Denies any local network traffic to access a VPN router.
	default-1	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: a) management traffic to access an IPsec VPN router, b) any traffic to originate from a remote site; return Service Level Agreement (SLA), c) Network Time Protocol (NTP), d) domain traffic back to a VPN router; and e) ICMP traffic to access a VPN router. Denies all other access to the router inbound from the inside interface and from the outside interface. Permits all traffic outbound from and inbound to the customer network. Denies all traffic inbound from the inside and outside interfaces.
	default-2	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: a) management traffic to access an IPsec VPN router, b) return User Datagram Protocol (UDP) to a VPN router for tftp and SNMP, c) return SLA, NTP, and domain traffic to the router, and d) ICMP to a VPN router from the management network. Denies all other access to a VPN router inbound from the inside and outside interfaces.
	ipsec-ingress	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; and all traffic to access from the remote site network.

Table 10-15 Example Templates Provided in VPN Solutions Center (continued)

Folder	Template	Description
	negotiated-wan-1	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: a) management traffic to access an IPsec VPN router, b) return UDP traffic to access a VPN router, c) application traffic, SLA, NTP, and domain to return to a VPN router, and d) ICMP traffic to access a VPN router. Denies all access to a VPN router from an outside interface. Permits customer stateful traffic to pass through the firewall. Denies all other traffic from an inside or outside interface.
	negotiated-wan-2	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: a) management traffic to access an IPsec VPN router, b) return UDP traffic to access a VPN router, c) application traffic, SLA, NTP, and domain to return to a VPN router, and d) ICMP traffic to access a VPN router. Denies all access to a VPN router from an outside interface. Permits customer stateful traffic to pass through the firewall. Denies all IP traffic from an inside or outside interface.
	vpnfiltered-1	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: a) management traffic to access an IPsec VPN router, b) return UDP traffic to access a VPN router, c) application traffic, SLA, NTP, and domain to return to a VPN router, and d) ICMP traffic to access a VPN router. Denies any other IP traffic access to a VPN router. Permits ICMP traffic to access a VPN router. Denies any other IP traffic to a VPN router. Permits ICMP traffic to access a customer network. Denies all other IP traffic. Permits: a) IPsec originated from a VPN router, b) stateful traffic from an inside interface back to an outside interface, and c) ICMP traffic from a customer network. Denies all other traffic from the inside interface.

Table 10-15 Example Templates Provided in VPN Solutions Center (continued)

Folder	Template	Description
	vpnfiltered-2	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; return UDP traffic to access a VPN router; application traffic, SLA, NTP, and domain to return to a VPN router; and ICMP traffic to access a VPN router. Denies any other IP traffic access to a VPN router and any IP traffic. Permits stateful traffic back through the firewall router.
	vpnfiltered-dial	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; all IP traffic from a remote network; application traffic, SLA, NTP; and ICMP traffic to access a VPN router from a management network. Denies any other access to a VPN router and any IP traffic. Permits stateful traffic back through the firewall router.
	vpnonly-1	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; return UDP to an IPsec VPN router for tftp and SNMP; application traffic, SLA, NTP, and domain to return to a VPN router; and ICMP traffic to access a VPN router from a management network. Denies all other access to the router inbound to the inside or outside interface. Permits: Extended Services Processor (ESP) outbound from a VPN router; and all services outbound from the customers network. Denies any IP traffic and all traffic inbound from the outside interface.
	vpnonly-2	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; any traffic from a remote site; application traffic, SLA, NTP, and domain to return to a VPN router; and ICMP traffic to access a VPN router from a management network. Denies all other access to the router inbound to the outside interface. Permits: ESP outbound from a VPN router. Denies any IP traffic from the outside interface and all traffic from the inside interface.

Table 10-15 Example Templates Provided in VPN Solutions Center (continued)

Folder	Template	Description
PIX	cust-qos	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; return management traffic from an IPsec router; ICMP traffic to access a VPN router; and return traffic from an outbound session. Denies all other traffic.
	cust-rtr1	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; ICMP traffic to access a VPN router; and remote sites traffic to access a VPN router. Denies local network traffic to access a VPN router. Permits customer network traffic to access the Internet. Denies other local network traffic access to the Internet.
	cust-rtr2	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; SNMP traffic from a management network to access a VPN router; application traffic to access a VPN router; and ICMP traffic to access a VPN router. Denies any other IP traffic access to a VPN router. Permits customer network access to a VPN router. Denies other local network traffic access to a VPN router.
	default-1	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; any traffic to originate from a remote site; SLA, NTP, and domain traffic to return to a VPN router; and ICMP traffic to access a VPN router. Denies all other access to the router inbound from the inside or outside interface. Permits all traffic outbound from and inbound to a customer network. Denies all traffic inbound from an inside or outside interface.

Table 10-15 Example Templates Provided in VPN Solutions Center (continued)

Folder	Template	Description
	default-2	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; return UDP traffic to access a VPN router; SLA, NTP, and domain traffic to return to a VPN router; and ICMP traffic to access a VPN router. Denies all other access to the router inbound from the inside or outside interface.
	ipsec-ingress	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; and all traffic from a remote site network.
	negotiated-wan-1	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; return UDP traffic to access a VPN router; application traffic, SLA, NTP, and domain traffic to return to a VPN router; and ICMP traffic to access a VPN router. Denies all access to a VPN router from an outside interface. Permits stateful traffic back through the firewall router. Denies all other traffic from the inside or outside interface.
	negotiated-wan-2	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; return UDP traffic to access a VPN router; application traffic, SLA, NTP, and domain traffic to return to a VPN router; and ICMP traffic to access a VPN router. Denies outside interface traffic to a VPN router. Denies all IP traffic from the inside or outside interface. Permits stateful traffic back through the firewall router.

Table 10-15 Example Templates Provided in VPN Solutions Center (continued)

Folder	Template	Description
	vpnfiltered-1	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; return UDP traffic to access a VPN router; application traffic, SLA, NTP, and domain traffic to return to a VPN router; and ICMP traffic to access a VPN router. Denies any other IP traffic to a VPN router. Permits ICMP traffic to a VPN router. Denies any other IP traffic to a VPN router. Permits ICMP traffic to a customer network. Denies all other IP traffic. Permits: IPsec traffic originated from a VPN router; stateful traffic from an inside or outside interface; and ICMP traffic from a customer network. Denies all other traffic from the inside interface.
	vpnfiltered-2	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; return UDP traffic to access a VPN router; application traffic, SLA, NTP, and domain traffic to return to a VPN router; and ICMP traffic to access a VPN router. Denies any other IP traffic to a VPN router. Denies any other IP traffic to a VPN router and any IP traffic. Permits stateful traffic back through the firewall router.
	vpnfiltered-dial	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; all IP traffic from a remote network; application traffic, SLA, NTP, and domain traffic to return to a VPN router; and ICMP traffic to access a VPN router. Denies any other IP traffic to a VPN router and any IP traffic. Permits stateful traffic back through the firewall router.

Table 10-15 Example Templates Provided in VPN Solutions Center (continued)

Folder	Template	Description
	vpnonly-1	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; return UDP traffic to access a router for tftp and SNMP; application traffic, SLA, NTP, and domain traffic to return to a VPN router; and ICMP traffic to access a VPN router from a management network. Denies all other access to the router inbound to the inside and outside interfaces. Permits: ESP outbound from a VPN router; and all services outbound from a customer network. Denies any IP traffic and all traffic inbound from the outside interface.
	vpnonly-2	Permits IPsec AH, ESP, and ISAKMP to enter the firewall. Permits: management traffic to access an IPsec VPN router; any traffic from a remote site; application traffic, SLA, NTP, and domain traffic to return to a VPN router; and ICMP traffic to access a VPN router from a management network. Denies all other access to the router inbound to the outside interface. Permits: ESP outbound from a VPN router; and all services outbound from a customer network. Denies any IP traffic from an inside or outside interface.
dslam	RFC1483_routed_prepend	Specifies commands to create the subinterface on the Digital Subscriber Line (DSL) modem interface that uses the Permanent Virtual Path (PVP) command and Request for Comments (RFC) 1483 bridged mode.
	pppoa_vt_append	Creates a Point-to-Point Protocol over Asynchronous Transfer Mode (PPPoA) Virtual Template.
	pppoa_snap_prepend	Creates a subinterface on the DSL interface.
	pppoe_vt_append	Used with VPNSC configlets to map a Point-to-Point Protocol over Ethernet (PPPoEthernet) session on an IP DSL switch into an MPLS VPN. It is appended to the VPN Service Request.
	pppoe_snap_prepend	Used with VPNSC configlets to map a PPPoEthernet session on an IP DSL switch into an MPLS VPN. It is prepended to the VPN Service Request.
	RBE_1483Br_prepend	Creates a subinterface on the IP DSL switch DSL interface and configures it for ATM RBE or RFC 1483 Bridged.

Table 10-15 Example Templates Provided in VPN Solutions Center (continued)

Folder	Template	Description
	pppoa_mux_prepend	Used for DSLAM modems configured for PPPoATM ATM Adaptation Layer 5 Multiplexing Device (AAL5MUX) Encapsulation.
	pppoa_mux_append	Creates a PPPoATM Virtual Template.
VPN 3000/ Routing	Ethernet-RIP	Sets up a Routing Information Protocol (RIP) protocol for a particular Ethernet interface.
	Ethernet-OSPF	Configures an Open Shortest Path First (OSPF) interface of a particular Ethernet interface.
	Static-Routes	Configures static route records.
	Default-Gateway	Sets up a default gateway.
	General-OSPF	Sets up General OSPF parameters.
	Create-OSPFArea	Creates general OSPF Area parameters.
	Modify-OSPFArea	Modifies a particular OSPF area.
	VRRP	Sets up VRRP redundancy
VPN 3000/ Servers	DNS-Server	Sets up Domain Name System (DNS) server parameters.
	FTP-Server	Sets up File Transfer Protocol (FTP) server parameters.
	HTTP-Server	Sets up Hypertext Transfer Protocol (HTTP) server parameters.
	HTTPS-Server	Sets up HTTPS server parameters.
	TFTP-Server	Sets up Trivial File Transfer Protocol (TFTP) server parameters.
	Telnet-Server	Sets up telnet server parameters.
	Telnet-SSL-Server	Sets up telnet over Secure Socket Layer (SSL) parameters.
	SNMP-Server	Sets up SNMP server parameters.
	SNMP-Communities	Sets up an SNMP communities string.
	SSL	Sets up SSL record parameters.
	SSH	Sets up SSH record parameters.
	DHCP-Server-Create	Sets up Dynamic Host Configuration Protocol (DHCP) server parameters.
	DHCP-Server-Modify	Modifies DHCP server parameters.
VPN 3000/ Events	General-Event	Sets up general event parameters.
	FTPLogBackup	Sets up FTP log back up parameters.
	EventClass-Create	Sets up EventClass record parameters.
	EventClass-Modify	Modifies a particular event class.
	TrapDestination-Create	Sets up Trap Destination record parameters.

Table 10-15 Example Templates Provided in VPN Solutions Center (continued)

Folder	Template	Description
	TrapDestination-Modify	Modifies a particular trap destination.
	SyslogServer-Create	Sets up syslog server record parameters.
	SyslogServer-Modify	Modifies syslog server record parameters.
	SMTPServer-Create	Creates an Simple Mail Transfer Protocol (SMTP) server.
	SMTPServer-Modify	Modifies SMTP server parameters.
	EmailRecipient-Create	Creates an e-mail recipient record.
	EmailRecipient-Modify	Modifies an e-mail recipient record.
VPN 3000/ Users	AddUsers	Used to add a list of users to the user database on the VPN 3000 concentrator. This database is used when a Group's authentication is set to Internal. For each user, the user name, password, and the Group to which these users belong must be provided.
	DeleteUsers	Used to remove a list of users from the user database on the VPN 3000 concentrator. This database is used when a Group's authentication is set to Internal.
EOverMPLS	EthernetOverMPLS	Provisions Ethernet over MPLS (EOverMPLS). The end points must be provisioned one at a time.
	Remove_EthernetOMPLS	Removes an EthernetOverMPLS configlet.
UTI		Provisions a Universal Transport Interface (UTI) tunnel between PEs. Each endpoint is provisioned one at a time.
MPLS-\ MultiVRF	MVCE-VRF-Name	Adds the **ip vrf** command to the Multi-VRF CE.
Provider-Facing Interface	NumberedInterface	Adds the **ip vrf forwarding** command for a numbered interface on the Multi-VRF CE.
	UnnumberedInterface	Adds the **ip vrf forwarding** command for an unnumbered interface on the Multi-VRF CE.
	Remove-NumberedInterface	Removes the **ip vrf forwarding** command from a numbered interface on the Multi-VRF CE.
	Remove-Unnumbered Interface	Removes the **ip vrf forwarding** command from an unnumbered interface on the Multi-VRF CE.
Provider-Facing-Routing	BGP-IpNumbered	Provisions **BGP** for an IP numbered provider-facing interface.
	BGP-IpUnnumbered	Provisions **BGP** for an IP unnumbered provider-facing interface.
	OSPF	Provisions **OSPF** on the Multi-VRF CE.
	RIP	Provisions **RIP** on the Multi-VRF CE.
	Static	Provisions a **Static** route on the Multi-VRF CE.
	Remove-BGP-IpUnnumbered	Removes **BGP** provisioning from an IP unnumbered provider-facing interface.

Table 10-15 Example Templates Provided in VPN Solutions Center (continued)

Folder	Template	Description
Customer-Facing Interface	NumberedInterface	Adds the **ip vrf forwarding** command to a customer-facing IP numbered interface.
	UnnumberedInterface	Adds the **ip vrf forwarding** command to a customer-facing IP unnumbered interface.
	FasterEthernetSample	Adds the **ip vrf forwarding** command to a customer-facing Fast Ethernet interface.
	SerialSample	Adds the **ip vrf forwarding** command to a customer-facing Serial interface.
	Remove-Numbered Interface	Removes **ip vrf forwarding** command from a customer-facing IP numbered interface.
	Remove-Unnumbered Interface	Removes **ip vrf forwarding** command from a customer-facing IP unnumbered interface.
Customer-Facing Routing	BGP-IpNumbered	Provisions **BGP** for a customer-facing IP numbered interface.
	BGP-IpUnnumbered	Provisions **BGP** for a customer-facing IP unnumbered interface.
	Static	Provisions a **Static** route for the customer-facing interface.
	RIP	Provisions **RIP** for the customer-facing interface.
	OSPF	Provisions **OSPF** for the customer-facing interface.
	Remove-BGP-IpUnnumbered	Removes **BGP** from customer-facing IP unnumbered interface.
templateEZVpn	ezvpn-client	IOS EZVpn client template for a low-end Cisco router that acts as a VPN client.
CNS	CNS-Agent-Enable	Enables the CNS client on the router to communicate with a CNS server (currently the Cisco IE2100).
Certificate	RSA-Key-Generation	IOS commands to generate the private/public key pair for this router.
	Root-Cert-Import	Imports the root certificate to the router.
	Root-Cert-By-Auth	Authenticates the root certificate server and obtains the root certificate.
	Cert-Enrollment	Enrolls with root certificate server and obtains certificate of this router.

CHAPTER 11

Provisioning Multi-VRF CEs in VPN Solutions Center

What Is a Multi-VRF CE?

The Multi-VRF CE is a feature that was introduced in Cisco IOS release 12.2(4)T. A small subset of the functionality formerly reserved to the PE is added to a Multi-VRF CE router—the PE-like functionality added to a Multi-VRF CE is the ability to have multiple VRFs on the CE router so that different routing decisions can be made. The packets are sent toward the PE as IP packets.

With this feature, a CE can maintain separate VRF tables to extend the privacy and security of an MPLS VPN down to a branch office, rather than just at the PE router node.

A Multi-VRF CE is unlike a CE in that there is no label exchange, no LDP adjacency, and no labeled packet flow between the PE and the CE.

Multi-VRF CE routers use VRF interfaces to form a VLAN-like configuration on the customer side. Each VRF on the Multi-VRF CE router is mapped to a VRF on the PE router.

Figure 11-1 illustrates one method in which a Multi-VRF CE can be used. The Multi-VRF CE router associates a specific VRF by the clients connected to its interfaces and exchanges that information with the PE. Routes are installed in the VRF on the Multi-VRF CE. There also needs to be a routing protocol or a static route that propagates routes from a specific VRF on the Multi-VRF CE to the corresponding VRF on the PE.

Figure 11-1 A Multi-VRF CE in an MPLS VPN Environment

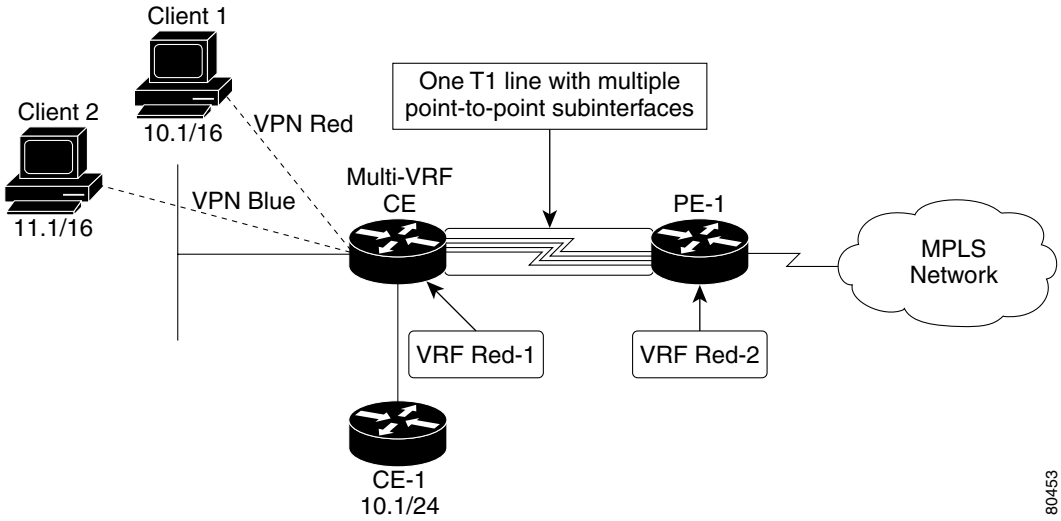

The Multi-VRF CE feature can segment its LAN traffic by placing each client or organization with its own IP address space, either on separate Ethernet interfaces such as CE-1, or through one Fast Ethernet interface segmented into multiple subinterfaces (for Client-1 and Client-2). To differentiate each client, each subinterface contains its own IP address space.

When receiving an outbound customer data packet from a directly attached interface, the Multi-VRF CE router performs a route lookup in the VRF that is associated with that site. The specific VRF is determined by the interface or subinterface over which the data packet is received. Support for multiple forwarding tables makes it easy for the CE router to provide segregation of routing information on a per-VPN basis before the routing information is sent to the PE. The use of a T1 line with multiple point-to-point subinterfaces allows traffic from the Multi-VRF CE router to the PE router to be segmented into each individual VRF.

With Multi-VRF CE configured on the CE router, the data path is as follows from the clients to PE-1 (as shown in Figure 11-1):

1. The Multi-VRF CE learns the VPN Red routes to Client 1 from a subinterface of the Fast Ethernet interface directly attached to Multi-VRF CE.
2. The Multi-VRF CE then installs these routes into VRF Red-1 (the VRF on the Multi-VRF CE).
3. PE-1 learns the VPN Red routes to Client 1 from VRF Red-1 on the Multi-VRF CE and installs the routes into VRF Red-2 (on PE-1).
4. The local VPN Blue routes from Client 2 are not associated with VPN Red and are not imported into VRF Red-1 or VRF Red-2.

Benefits of the Multi-VRF CE Feature

The benefits of the Multi-VRF CE feature are as follows:

- Without the use of cryptographic techniques (IPsec), security on the customer's LAN is equivalent to that supported by existing Layer 2 (ATM or Frame Relay) connections without the use of an additional switch.
- Only one CE router is needed—rather than a multiple CE router solution—thus making provisioning and network management easier.

- Because the Multi-VRF CE has VRF functionality without full PE functionality, there are fewer routing updates to manage.
- Overlapping customer address space.

  VPN customers often manage their own networks and use private address spaces. If customers do not use globally unique IP addresses, the same 32-bit IPv4 address can be used to identify different systems in different VPNs. The result can be routing difficulties because BGP assumes that each IPv4 address it carries is globally unique. To solve this problem, MPLS-VPNs supports a mechanism that converts an IP address that is not unique into globally unique addresses by combining the use of VPN-IPv4 address family with the deployment of Multiprotocol BGP Extensions (MP-BGP).
- No need for NAT (Network Address Translation) to allow support of overlapping IP address space.
- Increases the 32 routing process limits on the PE routers. A Multi-VRF CE could use five different OSPF processes to connect to five different customers in the same site, and then use BGP to propagate the routes to the PE.

Provisioning Procedures

To set up a Multi-VRF CE in VPN Solutions Center, there are two top-level tasks you must complete before you can provision the Multi-VRF CE:

1. Edit the Customer site to set the CE to a Multi-VRF CE (see "Defining a CE as a Multi-VRF CE" section on page 11-4).
2. On the Multi-VRF CE itself, you must apply a set of Multi-VRF CE templates to the following interfaces:
   - The provider-facing interface (see the "Templates Applied to the Provider-Facing Interface" section on page 11-7)
   - The customer-facing interface (see "Templates Applied to the Customer-Facing Interface" section on page 11-7).

For VPNSC 2.2, the Multi-VRF CE feature can be implemented by deploying one service request between the Multi-VRF CE, the regular CE, and a PE. In the scenario shown in Figure 11-2, PE-1 is the PE, CE-4 is the Multi-VRF CE and 2621-CE-6 is a regular CE.

The PE is configured as any other PE. For the multi-VRF CE, two interfaces must be provisioned—the provider-facing interface and the customer-facing interface (see the "Templates Applied to a Multi-VRF CE" section on page 11-6).

Figure 11-2 A Multi-VRF CE in an MPLS VPN Environment

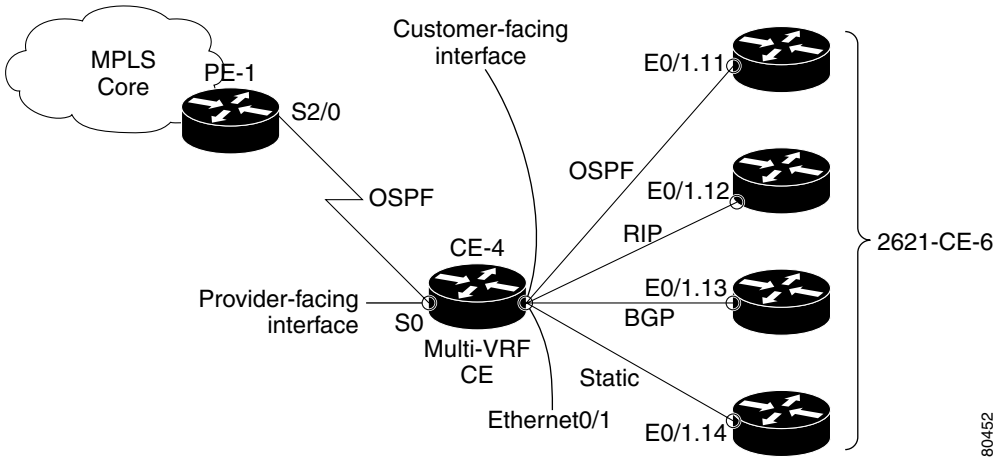

Defining a CE as a Multi-VRF CE

To define a CE as a Multi-VRF CE:

Step 1 In the VPN Console hierarchy view, expand the VPN Customers folder until you can see the sites for the Customer of interest.

Step 2 Select the site where the CE to be designated as the Multi-VRF CE resides, then **right-click**.

The Site menu appears (see Figure 11-3).

Figure 11-3 Opening a Site

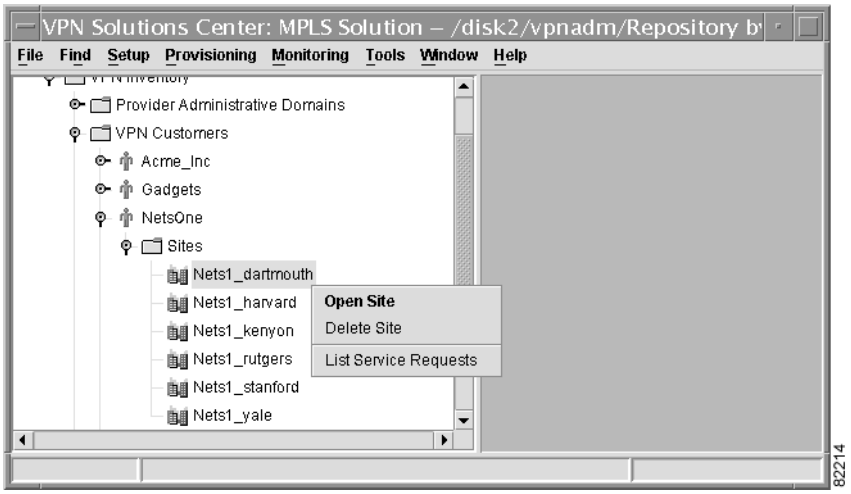

Step 3 From the Site menu, choose **Open Site**.

The Edit Customer Site dialog box is displayed (see Figure 11-4).

Figure 11-4 Editing a Site

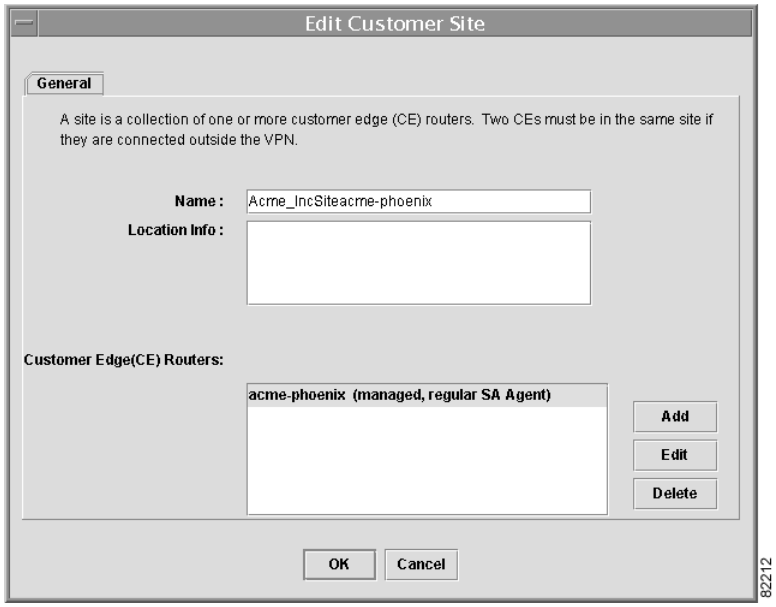

Step 4 From the Edit Customer Site dialog, click **Edit**.

The Edit Customer Edge Routers dialog box appears (see Figure 11-5).

Figure 11-5 Specifying the CE as a Multi-VRF CE

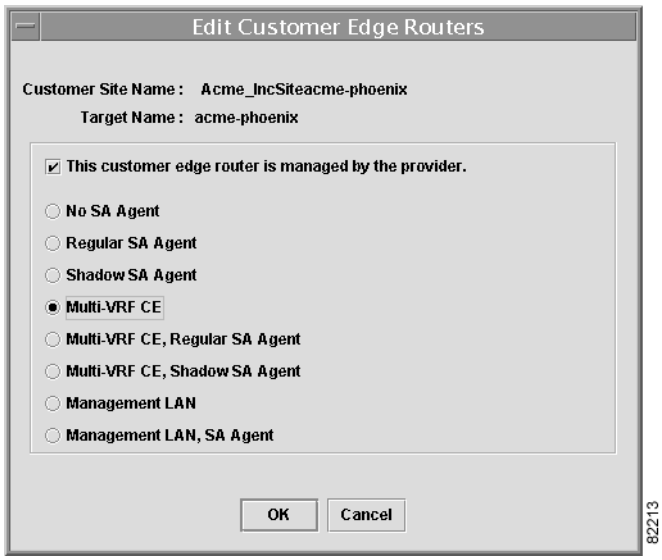

Step 5 Select the appropriate Multi-VRF CE option for the selected CE.

a. *Multi-VRF CE*: Indicates that the CE does not have SA Agent configured on it.

b. *Multi-VRF CE, Regular SA Agent*: Indicates that the CE has a dual function as a CE and router running SA Agent. That is, while functioning as a CE in the VPN, it also monitors traffic response times between CEs in the same VPN.

c. *Multi-VRF CE, Shadow SA Agent*: Indicates that the designated CE is actually a PE in provider space.

Step 6 Click **OK**.

You return to the Edit Customer Site dialog, where the CE is now designated as a Multi-VRF CE.

Step 7 Click **OK**.

Templates Applied to a Multi-VRF CE

To provision a Multi-VRF CE, you must use the templates provided by VPN Solutions Center specifically for a Multi-VRF CE.

Figure 11-6 shows the Multi-VRF CE templates available from the Template Manager:

Figure 11-6 VPNSC Templates Provided for Provisioning the Multi-VRF CE Interfaces

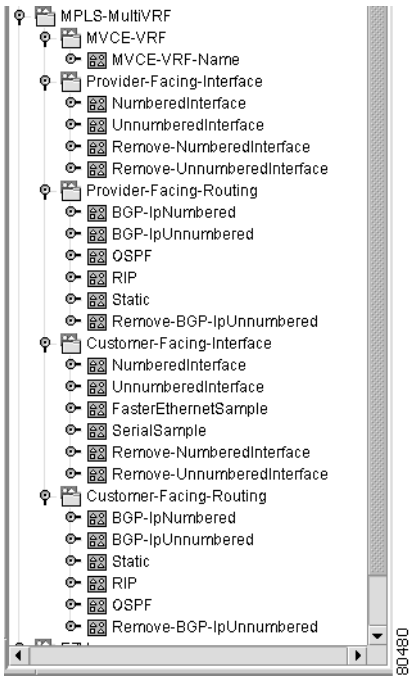

- For details on using the VPNSC Template Manager, see Chapter 10, “Provisioning with the VPN Solutions Center Template Manager.”
- For information on how to integrate templates with VPN Solutions Center configlets, see the “Templates” section on page 5-46.

On the Multi-VRF CE itself, you must apply a set of Multi-VRF CE templates to:

- The provider-facing interface (see the “Templates Applied to the Provider-Facing Interface” section on page 11-7)
- The customer-facing interface (see “Templates Applied to the Customer-Facing Interface” section on page 11-7).

Templates Applied to the Provider-Facing Interface

On a Multi-VRF CE, VPN Solutions Center provisions the provider-facing interface itself. In addition, you must apply three templates to the provider-facing interface:

- *MVCE-VRF template*

  Adds the `ip vrf` command to the provider-facing interface on the Multi-VRF CE.

- *Provider-Facing Interface template*

  Adds the `ip vrf forwarding` command to the provider-facing interface on the Multi-VRF CE.

  – Depending on whether the provider-facing interface is a numbered or unnumbered interface, select the appropriate provider-facing interface template.

  The Template Manager also provides templates to remove the commands added by the provider-facing templates.

- *Provider-Facing Routing template*

  Select the appropriate routing template for the link between the provider-facing interface and the PE.

  The Template Manager also provides templates to remove the BGP with IP unnumbered commands added by the provider-facing routing template.

Templates Applied to the Customer-Facing Interface

On a Multi-VRF CE, VPN Solutions Center provisions the provider-facing interface itself. In addition, you must apply three templates to the provider-facing interface:

- *MVCE-VRF template*

  Adds the `ip vrf` command to the customer-facing interface on the Multi-VRF CE.

- *Customer-Facing Interface template*

  Adds the `ip vrf forwarding` command to the customer-facing interface on the Multi-VRF CE.

  – Depending on whether the customer-facing interface is a numbered or unnumbered interface, select the appropriate customer-facing interface template.

  The Template Manager also provides templates to remove the commands added by the customer-facing templates.

- *Customer-Facing Routing template*

  – Select the appropriate routing template for the link between the customer-facing interface and the CE.

  The Template Manager also provides templates to remove the BGP with IP unnumbered commands added by the customer-facing routing template.

Descriptions of the Multi-VRF CE Templates

Table 11-1 provides a description for each of the Multi-VRF CE templates:

Table 11-1 Multi-VRF CE Templates Provided in VPN Solutions Center

Folder	Template	Description
MPLS-\ MultiVRF	MVCE-VRF-Name	Adds the **ip vrf** command to the Multi-VRF CE.
Provider-Facing Interface	NumberedInterface	Adds the **ip vrf forwarding** command for a numbered interface on the Multi-VRF CE.
	UnnumberedInterface	Adds the **ip vrf forwarding** command for an unnumbered interface on the Multi-VRF CE.
	Remove-NumberedInterface	Removes the **ip vrf forwarding** command from a numbered interface on the Multi-VRF CE.
	Remove-Unnumbered Interface	Removes **ip vrf forwarding** command from an unnumbered interface on the Multi-VRF CE.
Provider-Facing-Routing	BGP-IpNumbered	Provisions **BGP** for an IP numbered provider-facing interface.
	BGP-IpUnnumbered	Provisions **BGP** for an IP unnumbered provider-facing interface.
	OSPF	Provisions **OSPF** on the Multi-VRF CE.
	RIP	Provisions **RIP** on the Multi-VRF CE.
	Static	Provisions a **Static** route on the Multi-VRF CE.
	Remove-BGP-\ IpUnnumbered	Removes **BGP** provisioning from an IP unnumbered provider-facing interface.
Customer-Facing Interface	NumberedInterface	Adds the **ip vrf forwarding** command to a customer-facing IP numbered interface.
	UnnumberedInterface	Adds the **ip vrf forwarding** command to a customer-facing IP unnumbered interface.
	FastEthernetSample	Adds the **ip vrf forwarding** command to a customer-facing Fast Ethernet interface.
	SerialSample	Adds the **ip vrf forwarding** command to a customer-facing Serial interface.
	Remove-Numbered Interface	Removes **ip vrf forwarding** command from a customer-facing IP numbered interface.
	Remove-Unnumbered Interface	Removes **ip vrf forwarding** command from a customer-facing IP unnumbered interface.
Customer-Facing Routing	BGP-IpNumbered	Provisions **BGP** for a customer-facing IP numbered interface.
	BGP-IpUnnumbered	Provisions **BGP** for a customer-facing IP unnumbered interface.
	Static	Provisions a **Static** route for the customer-facing interface.
	RIP	Provisions **RIP** for the customer-facing interface.

Table 11-1 Multi-VRF CE Templates Provided in VPN Solutions Center (continued)

Folder	Template	Description
	OSPF	Provisions **OSPF** for the customer-facing interface.
	Remove-BGP-IpUnnumbered	Removes **BGP** from customer-facing IP unnumbered interface.

Contents of the Multi-VRF CE Templates

This section provides the contents of the Multi-VRF CE templates.

MVCE-VRF

```
ip vrf < MultiCE-vrf-name >
rd < bgp-as-number > : < suffix >
```

Provider-Facing-Interface:

For a Numbered Provider-Facing Interface—NumberedInterface

```
interface < MultiCE-interface-name >
ip vrf forwarding < MultiCE-vrf-name >
ip address < MultiCE-interface-addr > < MultiCE-interface-mask >
```

For an Unnumbered Provider-Facing Interface—UnnumberedInterface:

```
interface < MultiCE-loopback-name >
ip vrf forwarding < MultiCE-vrf-name >
ip address < MultiCE-interface-address > < MultiCE-interface-mask >
interface < MultiCE-interface-name >
ip vrf forwarding < MultiCE-vrf-name >
ip unnumbered < MultiCE-loopback-name >
```

For an IP unnumbered interface, you need to add a static route:

```
ip route vrf < MultiCE-vrf-name > < #system.getAddr($PE-loopback-ip-address) >
                                  < #system.getMask($PE-loopback-mask) >
                                  < MultiCE-interface-name >
```

Provider-Facing-Routing:

For BGP Protocol with an IP Numbered Provider-Facing Interface—BGP-ipNumbered

```
router BGP < autonomous-system-num >
address-family ipv4 vrf < MultiCE-vrf-name >
neighbor < PE-interface-ip-address > remote-as < PE-as-num >
redistribute < protocol-type >
exit-address-family
```

For BGP Protocol with an IP Unnumbered Provider-Facing Interface—BGP-ipUnnumbered

```
router BGP < autonomous-system-num >
address-family ipv4 vrf < MultiCE-vrf-name >
neighbor < PE-interface-ip-address > remote-as < PE-as-num >
neighbor < PE-interface-ip-address > ebgp-multihop
```

```
neighbor < PE-interface-ip-address > update-source < PE-loopback-name >
redistribute < protocol-type >
exit-address-family
```

You must add a static route for an unnumbered interface:

```
ip route vrf < CE-vrf-name > < PE-interface-ip-address > <PE-interface-mask >
                              < MultiCE-interface-name >
```

For OSPF Protocol with IP numbered and unnumbered on Provider-Facing Interface—OSPF

```
router OSPF < process-ID > vrf < MultiCE-vrf-name >
network < #system.getNetworkAddr($CE-interface-subnet) >
       < #system.getReverseMask($CE-reverse-mask) > area < area-num >
```

For RIP on Provider-Facing Interface

```
router RIP
version 2
address-family ipv4 vrf < MultiCE-vrf-name >
network < #system.getNetworkAddr($MultiCE-interface-network-address) >
no auto-summary
exit-address-family
```

For a Static Route on Provider-Facing Interface

```
ip route vrf < CE-vrf-name ><#system.getNetworkAddr($PE-interface-subnet) >
                            < PE-interface-mask > < MultiCE-interface-name >
```

MVCE-VRF

```
ip vrf < CE-vrf-name >
rd < bgp-as-number > : < suffix >
```

Customer-Facing-Interface

For a Numbered Customer-Facing Interface—NumberedInterface

```
interface < Customer-interface-name >
ip vrf forwarding < MultiCE-vrf-name >
ip address < Customer-interface-addr > < Customer-interface-mask >
```

For an Unnumbered Customer-Facing Interface—UnnumberedInterface:

```
interface < MultiCE-loopback-name >
ip vrf forwarding < MultiCE-vrf-name >
ip address < MultiCE-interface-address > < MultiCE-interface-mask >
interface < MultiCE-interface-name >
ip vrf forwarding < MultiCE-vrf-name >
ip unnumbered < MultiCE-loopback-name >
```

For an IP unnumbered interface, you need to add a static route:

```
ip route vrf < CE-vrf-name > < Customer-loopback-ip-address > < Customer-loopback-mask >
                            <MultiCE-interface-name >
```

Customer-Facing-Routing

For BGP Protocol with an IP Numbered Customer-Facing Interface—BGP-ipNumbered

```
router BGP < autonomous-system-num >
address-family ipv4 vrf < MultiCE-vrf-name >
neighbor < Customer-interface-ip-address > remote-as < Customer-AS-num >
redistribute < protocol-type >
exit-address-family
```

For BGP Protocol with an IP Unnumbered Customer-Facing Interface—BGP-ipUnnumbered

```
router BGP < autonomous-system-num >
address-family ipv4 vrf < MultiCE-vrf-name >
neighbor < Customer-interface-ip-address > remote-as < Customer-as-num >
neighbor < Customer-interface-ip-address > ebgp-multihop
neighbor < Customer-interface-ip-address > update-source < Customer-loopback-name >
redistribute < protocol-type >
exit-address-family
```

You must add a static route for an unnumbered interface:

```
ip route vrf < CE-vrf-name > < Cusetomer-interface-ip-address > <Customer-interface-mask >
                             < MultiCE-interface-name >
```

For OSPF Protocol with IP numbered and unnumbered on Customer-Facing Interface—OSPF

```
router OSPF < process-ID > vrf < MultiCE-vrf-name >
network < Customer-interface-subnet > < Customer-reverse-mask > area < area-num >
```

For RIP on Provider-Facing Interface

```
router RIP
version 2
address-family ipv4 vrf < MultiCE-vrf-name >
network < MultiCE-interface-network-address >
no auto-summary
exit-address-family
```

For a Static Route on Customer-Facing Interface

```
ip route vrf < MultiCE-vrf-name >< Customer-interface-ip-address >
                                 < Customer-interface-mask >
                                 < MultiCE-interface-name >
```

CHAPTER 12

Spanning Multiple Autonomous Systems

Overview

The inter-autonomous system for MPLS VPNs feature allows an MPLS VPN to span service providers and autonomous systems. An autonomous system is a single network or group of networks that is controlled by a common system administration group and that uses a single, clearly defined routing protocol.

As VPNs grow, their requirements expand. In some cases, VPNs need to reside on different autonomous systems in different geographic areas. Also, some VPNs need to extend across multiple service providers (overlapping VPNs). Regardless of the complexity and location of the VPNs, the connection between autonomous systems must be seamless to the customer.

The inter-autonomous systems for MPLS VPNs feature provides that seamless integration of autonomous systems and service providers. Separate autonomous systems from different service providers can communicate by exchanging IPv4 network layer reachability information (NLRI) in the form of VPN-IPv4 addresses. The autonomous systems' border edge routers use the Exterior Border Gateway Protocol (EBGP) to exchange that information. An interior gateway protocol (IGP) then distributes the network layer information for VPN-IPv4 prefixes throughout each VPN and each autonomous system. Routing information uses the following protocols:

- Within an autonomous system, routing information is shared using an interior gateway protocol.
- Between autonomous systems, routing information is shared using an Exterior Border Gateway Protocol. An EBGP allows a service provider to set up an interdomain routing system that guarantees the loop-free exchange of routing information between separate autonomous systems.

An MPLS VPN with inter-autonomous system support allows a service provider to provide to customers scalable Layer 3 VPN services, such as web hosting, application hosting, interactive learning, electronic commerce, and telephony service. A VPN service provider supplies a secure, IP-based network that shares resources on one or more physical networks.

The primary function of EBGP is to exchange network reachability information between autonomous systems, including information about the list of autonomous system routes. The autonomous systems use EGBP border edge routers to distribute the routes, which include label switching information. Each border edge router rewrites the next-hop and MPLS labels. See the "Routing Between Autonomous Systems" section on page 12-3 for more information.

Inter-autonomous system configurations supported in an MPLS VPN can include:

- *Interprovider VPN*: MPLS VPNs that include two or more autonomous systems, connected by separate border edge routers. The autonomous systems exchange routes using EBGP. No interior gateway protocol (IGP) or routing information is exchanged between the autonomous systems.

- *BGP Confederations*: MPLS VPNs that divide a single autonomous system into multiple sub-autonomous systems, and classify them as a single, designated confederation. The network recognizes the confederation as a single autonomous system. The peers in the different autonomous systems communicate over EBGP sessions; however, they can exchange route information as if they were IBGP peers.

Supported Edge Device Platforms

The following router platforms are supported at the service provider edge:

- Cisco 3600 series
- Cisco 4500 series
- Cisco 7200 series
- Cisco 7500 series

Benefits

The inter-autonomous system MPLS VPN feature provides the following benefits:

- Allows a VPN to cross more than one service provider backbone

  The inter-autonomous systems for MPLS VPNs feature allows service providers, running separate autonomous systems, to jointly offer MPLS VPN services to the same end customer. A VPN can begin at one customer site and traverse different VPN service provider backbones before arriving at another site of the same customer. Previously, MPLS VPNs could only traverse a single BGP autonomous system service provider backbone. The inter-autonomous system feature allows multiple autonomous systems to form a continuous (and seamless) network between a service provider's customer sites.

- Allows a VPN to exist in different areas

  The inter-autonomous systems for MPLS VPNs feature allows a service provider to create a VPN in different geographic areas. Having all VPN traffic flow through one point (between the areas) allows for better rate control of network traffic between the areas.

- Allows confederations to optimize IBGP meshing

  The inter-autonomous systems feature can make IBGP meshing in an autonomous system more organized and manageable. You can divide an autonomous system into multiple, separate sub-autonomous systems and then classify them into a single confederation (even though the entire VPN backbone appears as a single autonomous system). This capability allows a service provider to offer MPLS VPNs across the confederation because it supports the exchange of labeled VPN-IPv4 network layer reachability information between the sub-autonomous systems that form the confederation.

Routing Between Autonomous Systems

Figure 12-1 illustrates one MPLS VPN consisting of two separate autonomous systems. Each autonomous system operates under different administrative control and runs a different IGP. Service providers exchange routing information through EBGP border edge routers (ASBR1 and ASBR2).

Figure 12-1 EBGP Connection Between Two Autonomous Systems

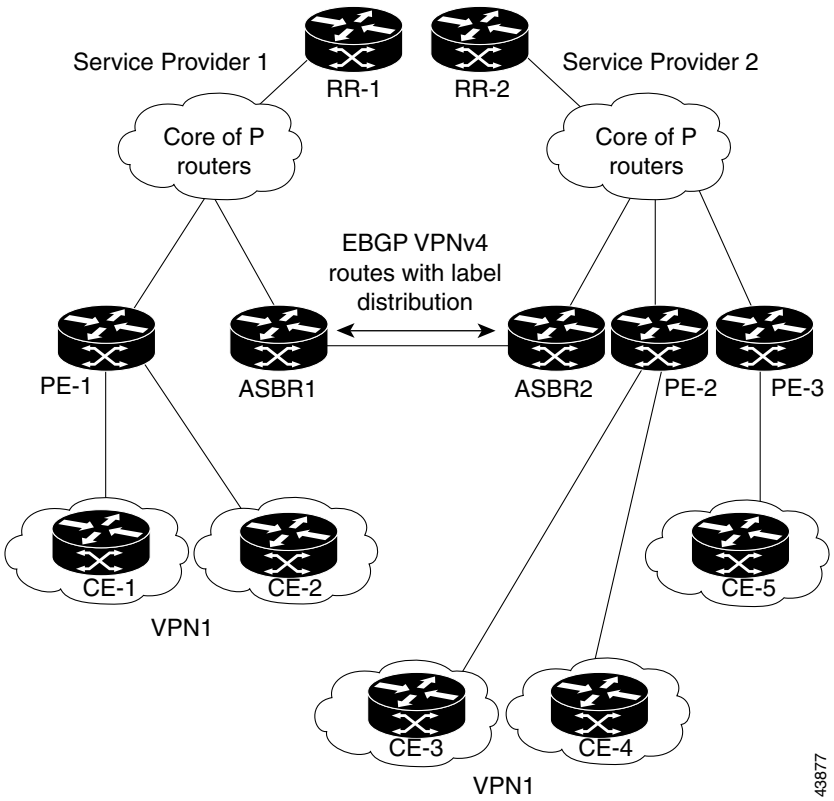

This configuration uses the following process to transmit information:

1. The provider edge router (PE-1) assigns a label for a route before distributing that route. The PE router uses the multiprotocol extensions of a border gateway protocol (BGP) to transmit label mapping information. The PE router distributes the route as a VPN-IPv4 address. The address label and the VPN identifier are encoded as part of the NLRI.
2. The two route reflectors (RR-1 and RR-2) reflect VPN-IPv4 internal routes within the autonomous system. The autonomous systems' border edge routers (ASBR1 and ASBR2) advertise the VPN-IPv4 external routes.
3. The EBGP border edge router (ASBR1) redistributes the route to the next autonomous system, (ASBR2). ASBR1 specifies its own address as the value of the EBGP next hop attribute and assigns a new label. The ASBR1 address ensures the following:
   - The next hop router is always reachable in the service provider (P) backbone network.
   - The label assigned by the distributing router is properly interpreted. The label associated with a route must be assigned by the corresponding next hop router.

4. The EBGP border edge router (ASBR2) redistributes the route in one of the following ways, depending on its configuration:
   - If the IBGP neighbors are configured with the **neighbor next-hop-self** command, ASBR2 changes the next hop address of updates received from the EBGP peer, then forwards it on.
   - If the IBGP neighbors are not configured with the **neighbor next-hop-self** command, the next hop address does not get changed. ASBR2 must propagate a host route for the EBGP peer through the IGP.

     To propagate the EBGP VPN-IPv4 neighbor host route, use the **redistribute connected subnets** command. The EBGP VPN-IPv4 neighbor host route is automatically installed in the routing table when the neighbor comes up. This is essential to establish the label-switched path between PE routers in different autonomous systems.

Exchanging VPN Routing Information

Autonomous systems exchange VPN routing information (routes and labels) to establish connections. To control connections between autonomous systems, the PE routers and EBGP border edge routers maintain a Label Forwarding Information Base (LFIB). The LFIB manages the labels and routes that the PE routers and EBGP border edge routers receive during the exchange of VPN information.

Figure 12-2 illustrates the exchange of VPN route and label information between autonomous systems. The autonomous systems use the following guidelines to exchange VPN routing information:

Routing information includes:

- The destination network (N)
- The next hop field associated with the distributing router
- A local MPLS label (L)

An *RD1: route distinguisher* is part of a destination network address to make the VPN-IPv4 route globally unique in the VPN service provider environment.

The *ASBRs* are configured to change the next hop (next-hop-self) when sending VPN-IPv4 NLRIs to the IBGP neighbors. Therefore, the ASBRs must allocate a new label when they forward the NLRI to the IBGP neighbors.

Figure 12-2 Exchanging Routes and Labels Between Two Autonomous Systems

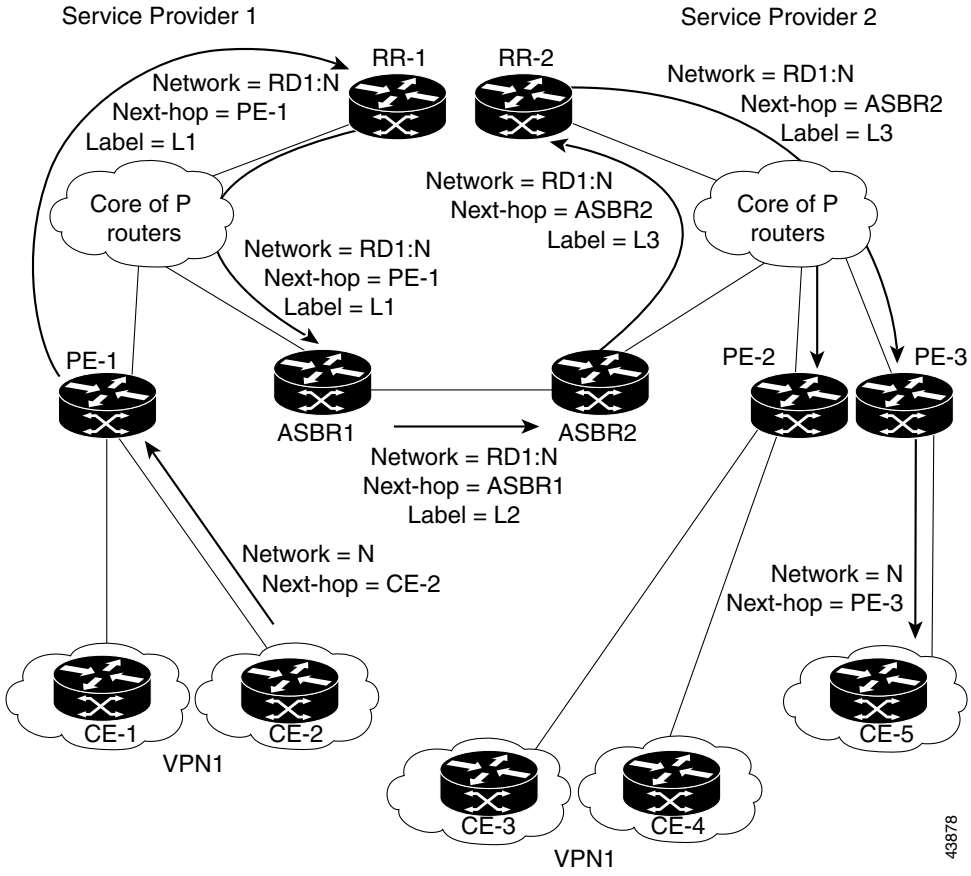

Figure 12-3 illustrates the exchange of VPN route and label information between autonomous systems. The only difference is that ASBR2 is configured with the **redistribute connected** command, which propagates the host routes to all PEs. The **redistribute connected** command is necessary because ASBR2 is not the configured to change the next hop address.

Figure 12-3 Host Routes Propagated to All PEs Between Two Autonomous Systems

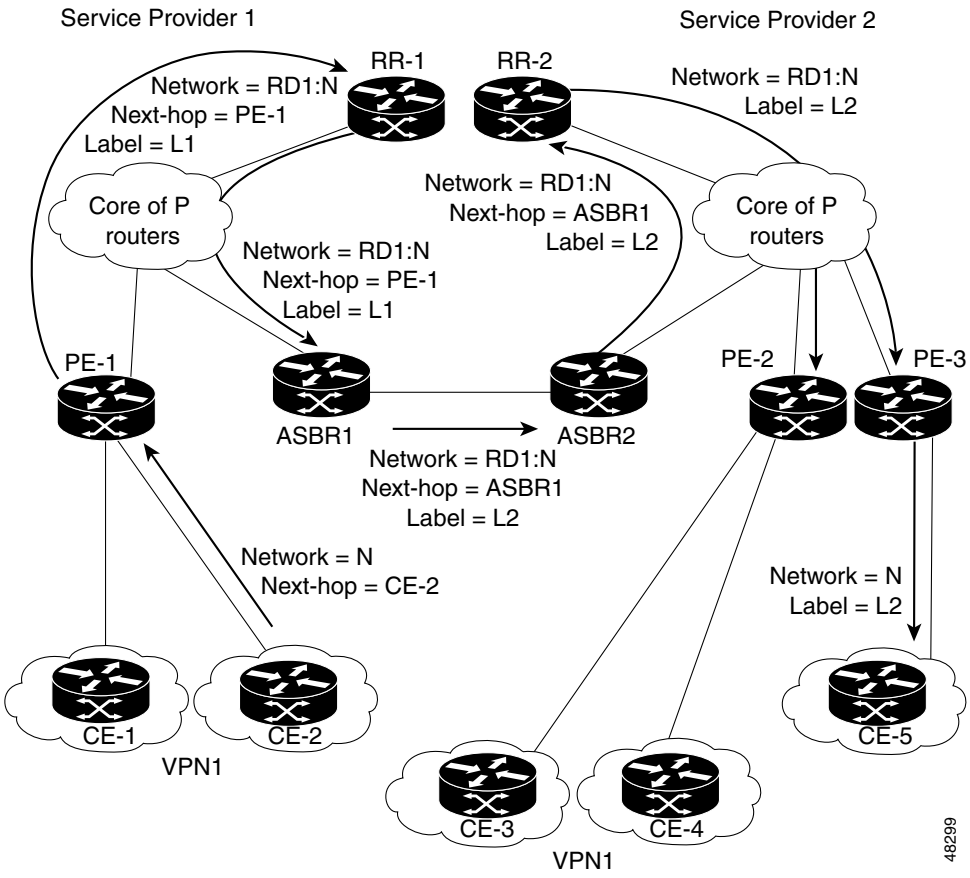

Figure 12-4 illustrates how packets are forwarded between autonomous systems in an interprovider network using the following packet forwarding method:

Packets are forwarded to their destination via MPLS. Packets use the routing information stored in the LFIB of each PE router and EBGP border edge router. The service provider VPN backbone uses dynamic label switching to forward labels.

Each autonomous system uses standard multi-level labeling to forward packets between the edges of the autonomous system routers (for example, from CE-5 to PE-3). Between autonomous systems, only a single level of labeling is used, corresponding to the advertised route.

A data packet carries two levels of labels when traversing the VPN backbone:

- The first label (*IGP route label*) directs the packet to the correct PE router or EBGP border edge router. (For example, the IGP label of ASBR2 points to the ASBR2 border edge router.)
- The second label (*VPN route label*) directs the packet to the appropriate PE router or EBGP border edge router.

Figure 12-4 Forwarding Packets Between Two Autonomous Systems

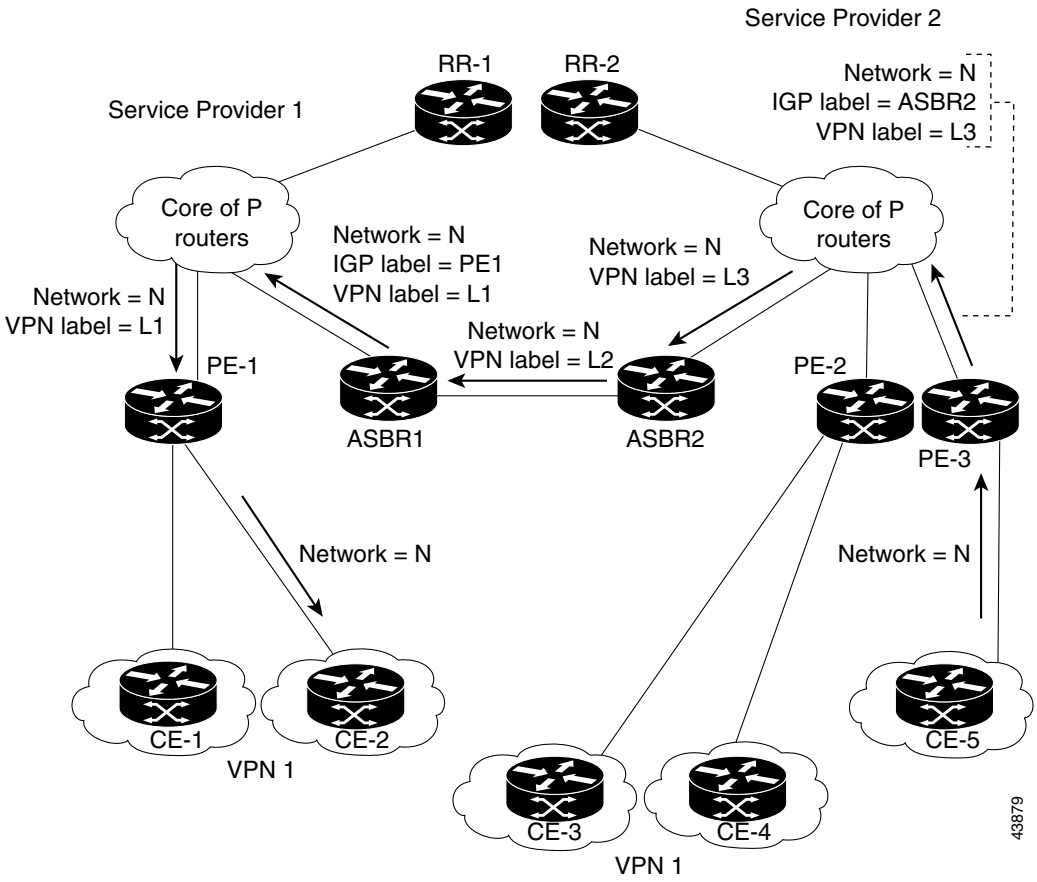

Figure 12-5 illustrates shows the same packet forwarding method, except the EBGP router (ASBR1) forwards the packet without reassigning it a new label.

Figure 12-5 Forwarding Packets Without Reassigning a New Label

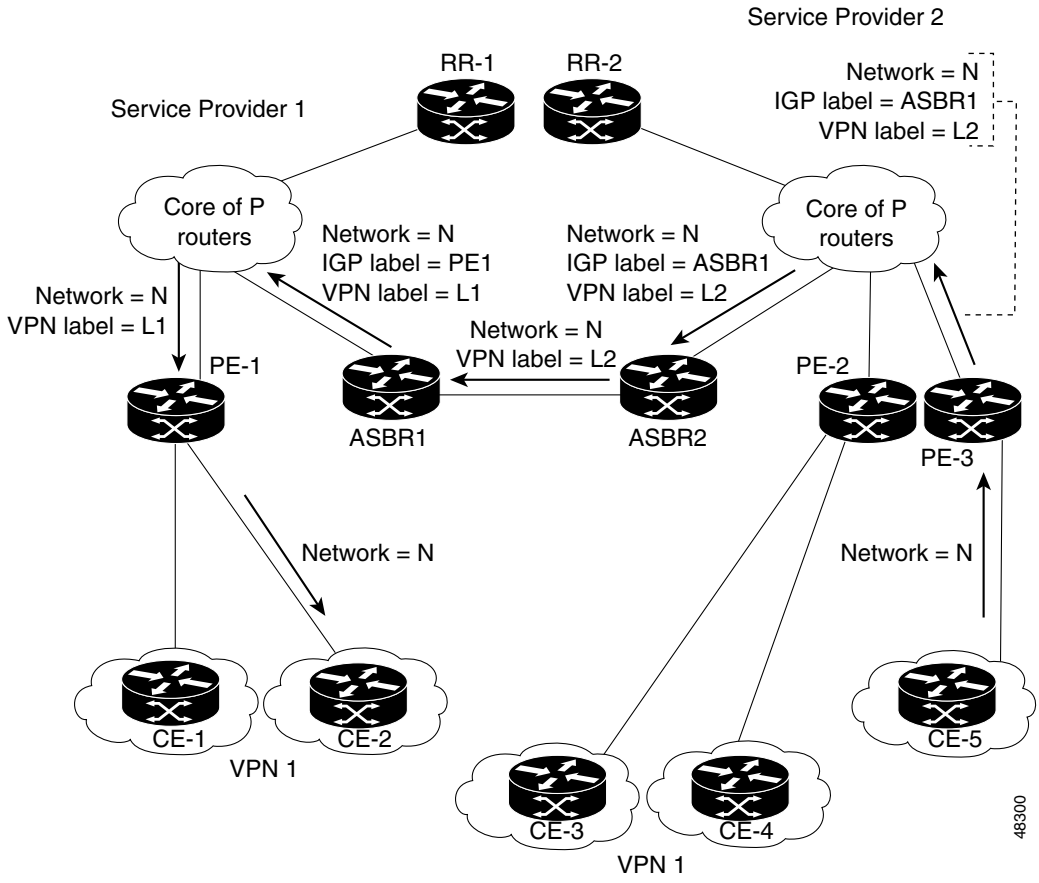

Routing Between Subautonomous Systems in a Confederation

A VPN can span service providers running in separate autonomous systems or between multiple subautonomous systems that have been grouped together to form a confederation.

A confederation reduces the total number of peer devices in an autonomous system. A confederation divides an autonomous system into subautonomous systems and assigns a confederation identifier to the autonomous systems.

In a confederation, each subautonomous system is fully meshed with other subautonomous systems. The subautonomous systems communicate using an IGP, such as Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS). Each subautonomous system also has an EBGP connection to the other subautonomous systems. The confederation EBGP (CEBGP) border edge routers forward next-hop-self addresses between the specified subautonomous systems. The next-hop-self address forces the BGP to use a specified address as the next hop rather than letting the protocol choose the next hop.

You can configure a confederation with separate subautonomous systems in two ways:

- You can configure a router to forward next-hop-self addresses between only the CEBGP border edge routers (both directions). The subautonomous systems (IBGP peers) at the subautonomous system border do not forward the next-hop-self address. Each subautonomous system runs as a single IGP domain. However, the CEBGP border edge router addresses are known in the IGP domains.
- You can configure a router to forward next-hop-self addresses between the CEBGP border edge routers (both directions) and within the IBGP peers at the subautonomous system border. Each subautonomous system runs as a single IGP domain but also forwards next-hop-self addresses between the PE routers in the domain. The CEBGP border edge router addresses are known in the IGP domains.

Figure 12-6 illustrates a typical MPLS VPN confederation configuration. In this confederation configuration:

- The two CEBGP border edge routers exchange VPN-IPv4 addresses with labels between the two subautonomous systems.
- The distributing router changes the next-hop addresses and labels and uses a next-hop-self address.
- IGP-1 and IGP-2 know the addresses of CEBGP-1 and CEBGP-2.

Figure 12-6 EGBP Connection Between Two AS's in a Confederation

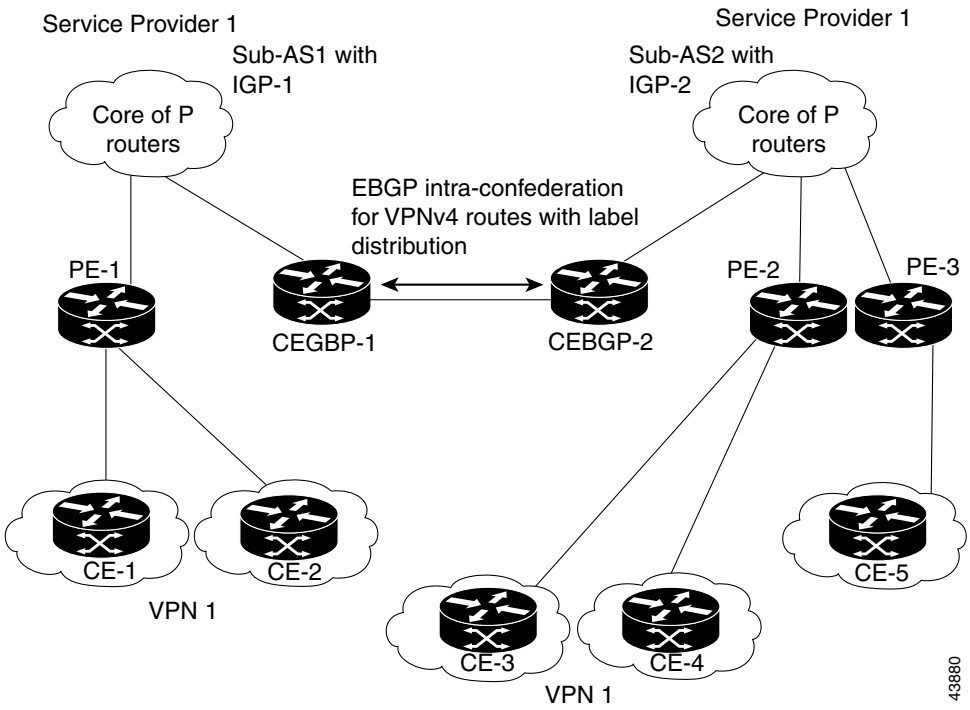

In this confederation configuration:

- CEBGP border edge routers function as neighboring peers between the subautonomous systems. The sub-autonomous systems use EBGP to exchange route information.
- Each CEBGP border edge router (CEBGP-1, CEBGP-2) assigns a label for the route before distributing the route to the next subautonomous system. The CEBGP border edge router distributes the route as a VPN-IPv4 address by using the multiprotocol extensions of BGP. The label and the VPN identifier are encoded as part of the NLRI.

- Each PE and CEBGP border edge router assigns its own label to each VPN-IPv4 address prefix before redistributing the routes. The CEBGP border edge routers exchange VPN-IPv4 addresses with the labels.

  The next-hop-self address is included in the label (as the value of the EBGP next-hop attribute). Within the sub-autonomous systems, the CEBGP border edge router address is distributed throughout the IBGP neighbors and the two CEBGP border edge routers are known to both confederations.

Using VPNSC to Span Multiple Autonomous Systems

As described in the "Exchanging VPN Routing Information" section on page 12-4, autonomous systems exchange VPN routing information (routes and labels) to establish connections. To control connections between autonomous systems, the PE routers and Exterior BGP ASBRs (Autonomous System Boundary Routers) maintain a Label Forwarding Information Base (LFIB). The LFIB manages the labels and routes that the PE routers and EBGP border edge routers receive during the exchange of VPN information.

The ASBRs are configured to change the next hop (next-hop-self) when sending VPN-IPv4 network layer reachability information to their IBGP neighbors. Therefore, the ASBRs must allocate a new label when they forward the NLRI to their IBGP neighbors.

Figure 12-7 shows the example VPNSC network used in this section.

Figure 12-7 Example VPN Network with Two Autonomous Systems

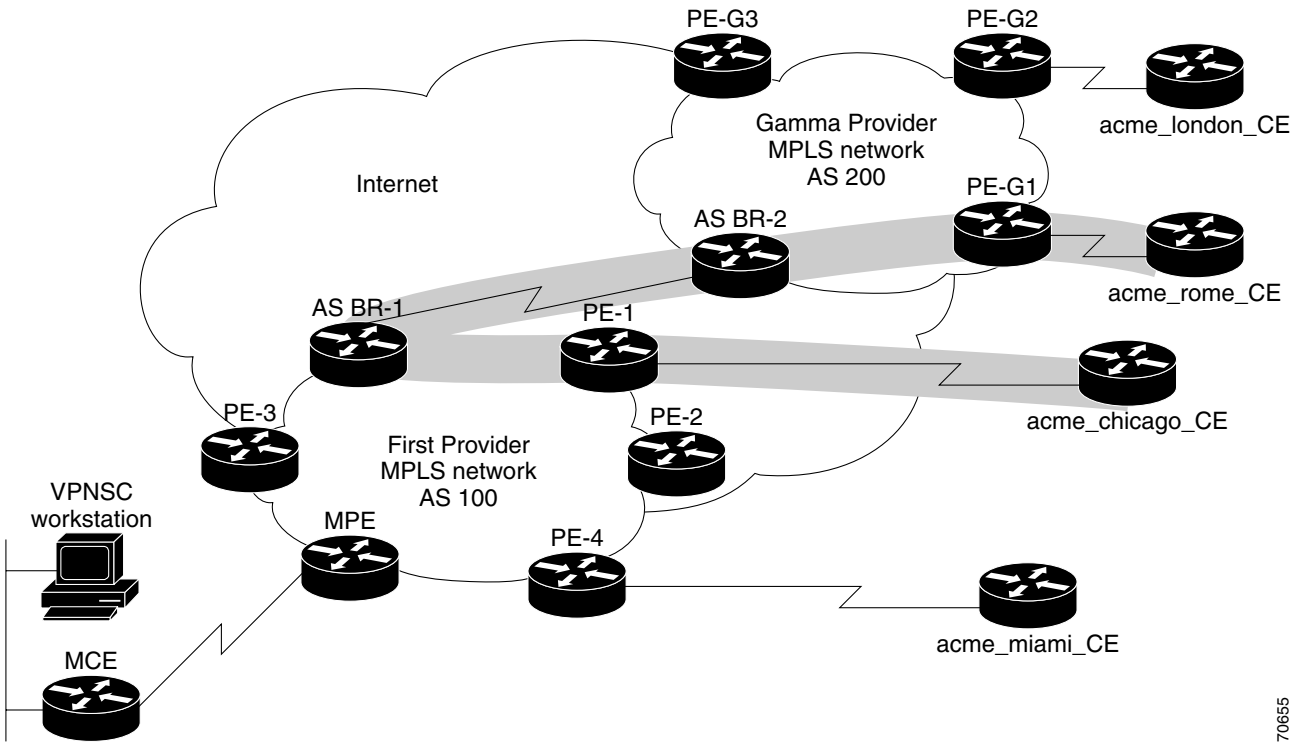

In order for traffic from Acme_Chicago in AS 100 to reach Acme_Rome in AS 200, VPN Solutions Center must provision two links only:

- The link between Acme_Chicago and PE-1
- The link between Acme_Rome and PE-G1

As shown in Figure 12-7 on page 12-10, VPN Solutions Center routes the VPN traffic from PE-1 to ASBR-1, from ASBR-1 to ASBR-2, then from ASBR-2 to PE-G1; finally the traffic is routed to its destination, Acme-Rome.

ASBR-1 and ASBR-2 must run BGP (Border Gateway Protocol). Then iMP-BGP (interior Multiprotocol BGP) handles the routes between PE-1 to ASBR-1 in AS 100 and the routes between PE-2 to ASBR-2 in AS 200. eMP-BGP (exterior Multiprotocol BGP) handles the routes between ASBR-1 and ASBR-2.

Tip The service provider must configure a VPN-IPv4 EBGP session between directly connected Autonomous System Boundary Routers (ASBRs). This is a one-time setup procedure that the service provider must manage. Cisco VPN Solutions Center does not provision the link between the ASBR devices that span autonomous systems.

A VPN-IPv4 address (also referred to as a *VPNv4* address) is the combination of the IPv4 address and the 8-byte route distinguisher (RD). Combining the RD and the IPv4 address makes the IPv4 route globally unique across the MPLS VPN network. BGP considers an IPv4 address as different from another IPv4 address that has the same network and subnet mask when the route distinguishers are different.

In this procedure, we will use the VPN Solutions Center 2.1 service request provisioning user interface:

Provisioning the Links Between Two Autonomous Systems

To use VPN Solutions Center to provision links between two autonomous systems, follow these steps:

Step 1 From the VPN Console, choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor is displayed (see Figure 12-8).

Figure 12-8 MPLS Service Request Editor

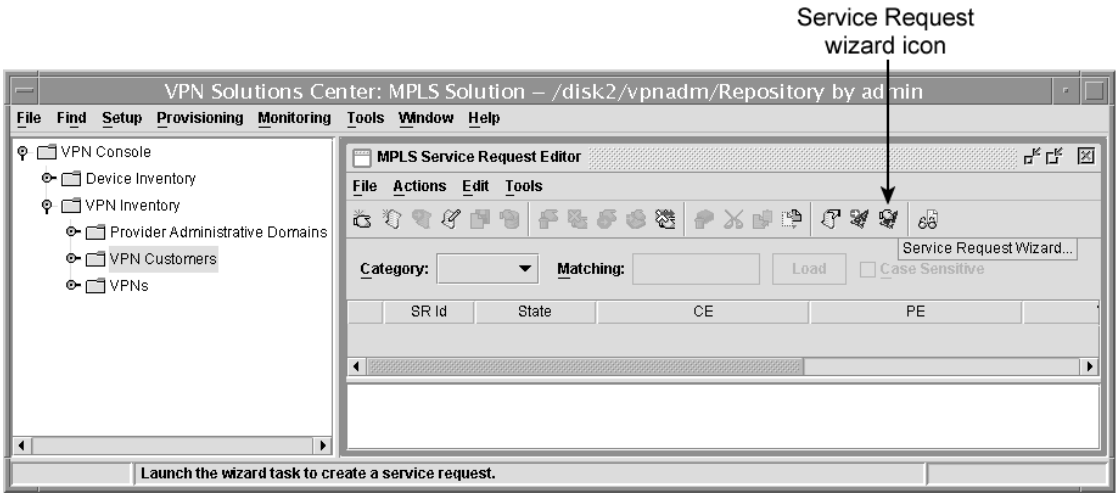

Step 2 To switch to the VPNSC 2.1 service request wizard, click the Service Request wizard icon.

The first—and informational only—screen appears.

Step 3 Click **Next**. The Select CE dialog box appears (see Figure 12-9).

Selecting a Customer Edge Router (CE) in the First AS

Step 4 From the Select CE dialog box, select the customer edge router in the first autonomous system (see Figure 12-9).

Figure 12-9 The Select CE Dialog Box

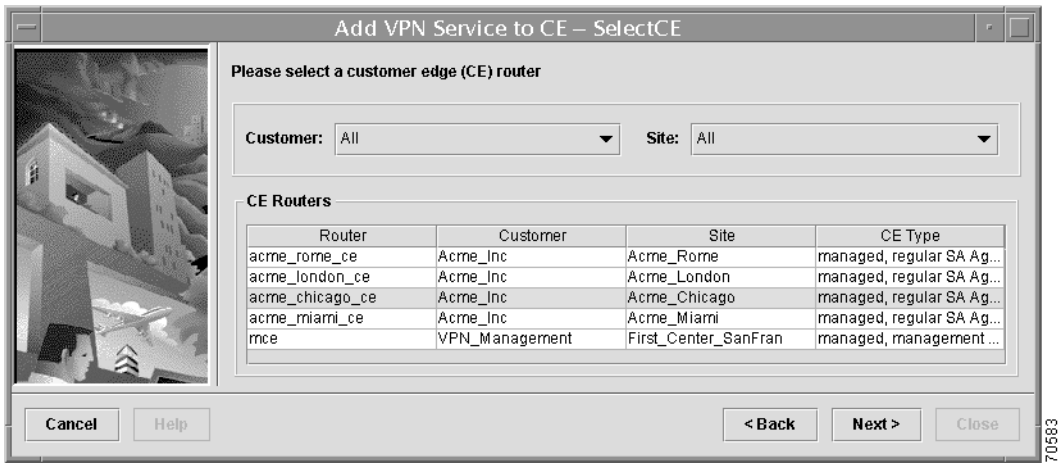

a. *Customer*: From the Customer drop-down list, select the appropriate customer.

b. *Site*: From the Site drop-down list, select the appropriate site.

c. *CE for this link*: From the CE Routers list, select the appropriate CE for this link, then click **Next**.

Selecting the Provider Edge Router (PE)

Step 5 From the Select PE dialog box, select the customer edge router in the first autonomous system (see Figure 12-10).

Figure 12-10 Select PE Dialog Box

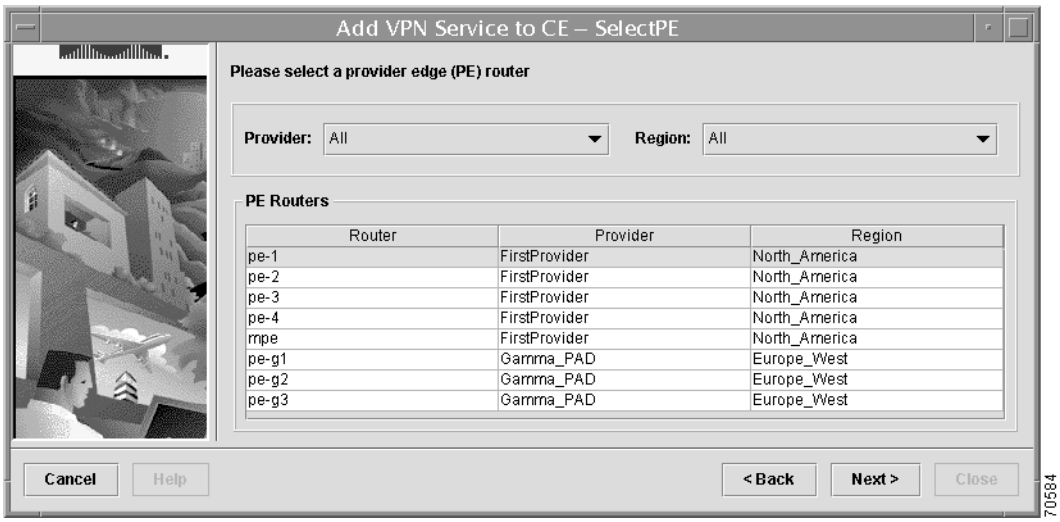

a. From the Provider drop-down list, select the appropriate service provider name.

b. From the Region drop-down list, select the appropriate region.

c. From the PE Routers list, select the provider edge router for this link, then click **Next**.

Selecting the VPN and the Site's Topology

In this step, you select the VPN that the CE belongs to, as well as specifying the CE's role in the VPN.

Tip Please note that when you create the service request for the PE-CE link in the target autonomous system, you must specify the same VPN that you specify here.

The Select VPN dialog box appears (see Figure 12-11).

Figure 12-11 Selecting the VPN

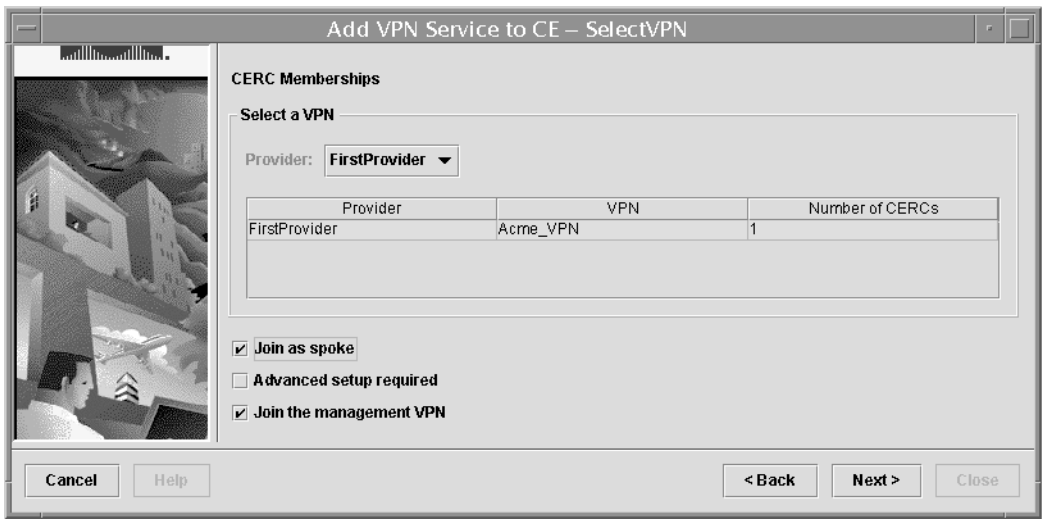

Step 6 Complete the VPN selection options as follows:

a. Select the appropriate Provider from the list (the VPN is associated with the Provider).

b. Select the VPN name.

c. If the selected CE (which represents the customer site) should be a spoke in the VPN, check the **Join as Spoke** check box.

   By default, the **Join as Spoke** check box is not enabled. When this option is not enabled, the selected CE can communicate with all the other sites in the VPN.

d. If you are adding a CE to the *management VPN*, check the **Join the management VPN** check box. For more information, see Chapter 8, "The VPNSC Management Network."

   When you use the VPN Solutions Center software to define a management VPN, the software automatically generates an *export route map* for the management VPN.

e. If you are building a VPN with CEs that are members of multiple VPNs (also referred to as *extranets*), check the **Advanced setup required** check box.

f. When satisfied with your selections, click **Next**.

Choosing the Routing Protocol for the Link

The Select VPN dialog box appears (see Figure 12-12).

Figure 12-12 RIP Protocol Routing Policy Options

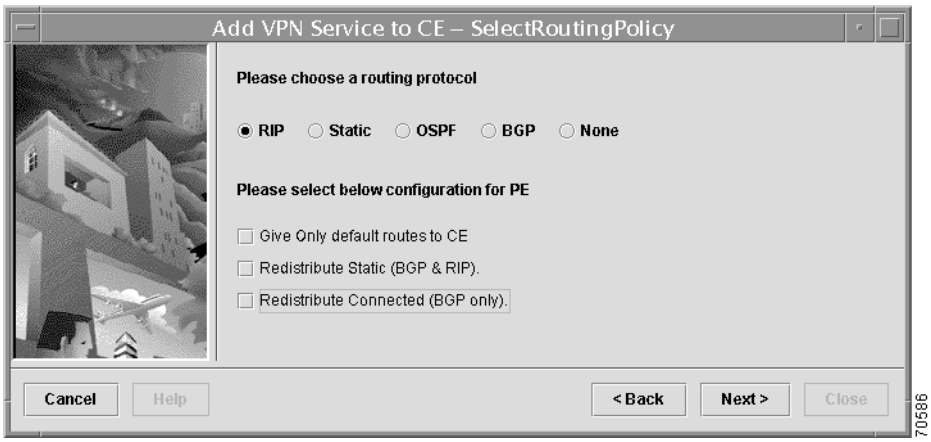

Step 7 Choose the routing protocol for the PE-CE link.

The routing protocol you choose must run on both the PE and the CE.

- You can choose **RIP** (Routing Information Protocol), **Static** (for specifying a static route), **OSPF** (Open Shortest Path First), **BGP** (Border Gateway Protocol), or **None** (to specify parameters for cable service).
- The wizard presents a different sequence of screens and requires different information depending on which protocol you choose.

a. Complete the necessary fields and other information required for the selected routing protocol as described in the router-specific sections below.

b. When satisfied with your selections, click **Next**.

Specifying Redistributed Protocols on the Link

The Redistribute Protocols dialog box appears.

Step 8 Indicate whether you require to redistribute protocols from the selected CE router.

a. If protocol redistribution is not required on this link, click **Next.**

b. If redistribution is required, specify the routing protocols that must be redistributed from the CE by clicking **Add**, then completing the necessary information in the Redistributed Protocols dialog box.

Defining the Interfaces on the PE-CE Link

Step 9 Define the interfaces and their encapsulations on the PE and CE (see Figure 12-13).

Figure 12-13 The Select PE-CE Interfaces Dialog Box

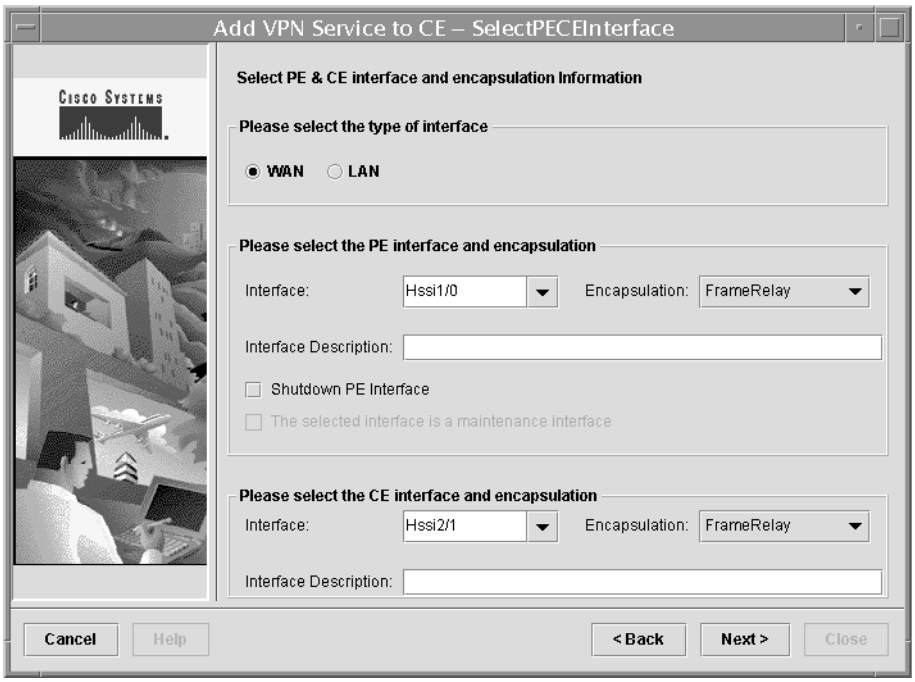

a. Specify the type of interfaces for the PE-CE link—WAN or LAN.

b. Select the PE interface and its encapsulation method from the drop-down lists.

   The interfaces available are determined by the PE's configuration file. The encapsulation methods are determined by which interface you select.

c. If appropriate, enable the **Shutdown PE Interface** and **Maintenance Interface** options.

d. Specify the CE interface and its protocol encapsulation from the drop-down lists, then click **Next**.

Step 10 If you specified serial interfaces for the PE and CE and chose Frame Relay as the encapsulation, specify the Data-Link Connection Identifier (DLCI) numbers for the PE-CE link as shown in Figure 12-14, then click **Next**.

Figure 12-14 Serial Interface Protocol Encapsulation Information

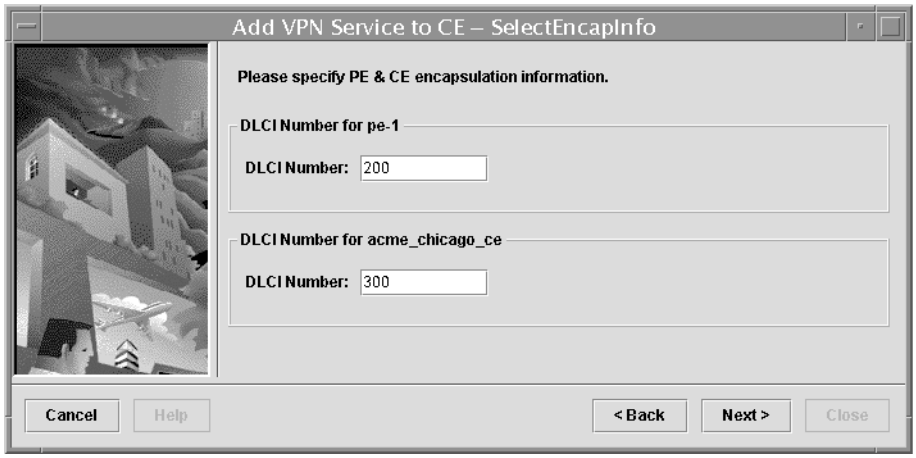

Frame Relay allows multiple logical data streams to be multiplexed onto a single physical link. These logical data streams are called *virtual circuits* and are identified with a Data-Link Connection Identifier (DLCI). A DLC has only local significance in Frame Relay. It can—and generally does—change on each physical link in the Frame Relay network.

Choosing an IP Addressing Scheme

Step 11 Choose an IP addressing scheme for the PE and CE (see Figure 12-15). Then click **Next**.

Figure 12-15 IP Addressing Scheme Dialog Box

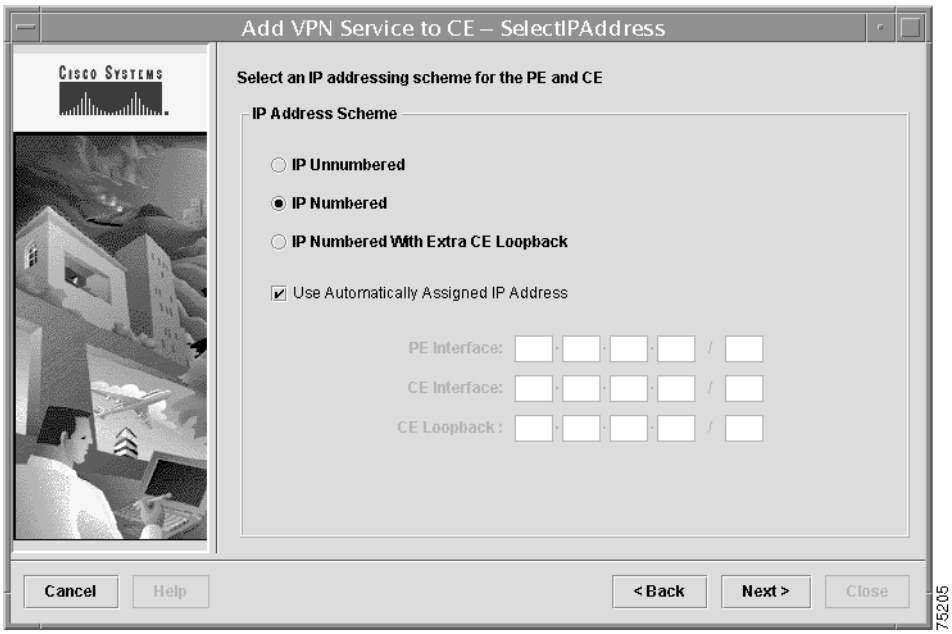

You can choose among four options:

- IP Unnumbered
- IP Numbered

- IP Numbered with Extra CE Loopback
- Use Automatically Assigned IP Address checkbox

When you have finished, the following screen is displayed (see Figure 12-16).

Figure 12-16 Specify VRF Parameters Dialog Box

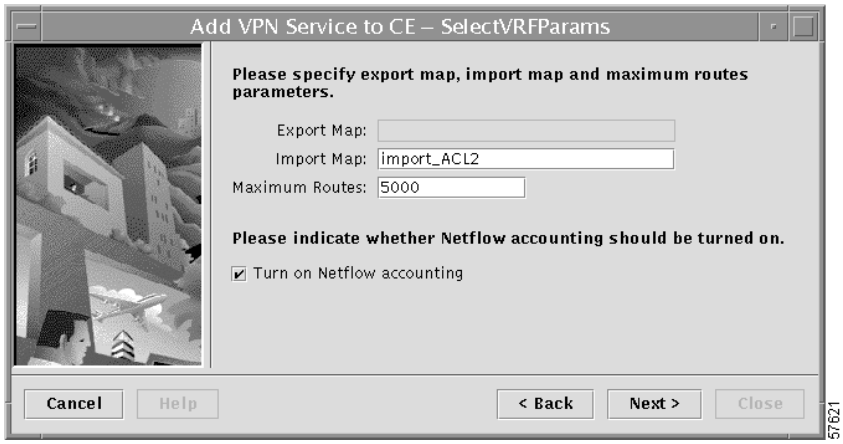

Step 12 Enter the VRF parameters as follows:

a. *Export Map*: If necessary, enter the name of the export map.

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

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 (as shown in Figure 12-16).

b. *Import Map*: Enter the name of the import map.

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

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

d. To enable NetFlow accounting, check the **Turn on NetFlow accounting** checkbox.

e. When you have completed the fields as necessary in the Specify VRF Parameters dialog box, click **Next**.

Selecting a Class of Service Profile for the PE-CE Link

The Class of Service Profile dialog box appears (see Figure 12-17).

Figure 12-17 Selecting a CoS Profile

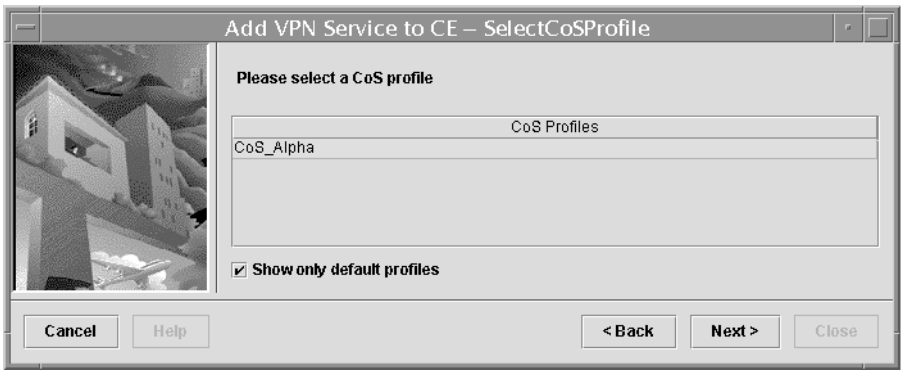

Step 13 If desired, select a Class of Service (CoS) profile to assign to the PE-CE link, then click **Next**.

Class of Service profiles are applied to the Provider Edge Router (PE), but the CoS definition is enforced across the PE-CE link on both the PE and CE.

The next series of dialog boxes let you integrate

Integrating VPN Solutions Center Templates with a Service Request

Step 14 If desired, complete the procedure to integrate a specific template with the current service request.

Completing the Service Request

VPN Solutions Center displays a summary of all the service settings defined for this VPN, including the information on template provisioning for the CE and PE.

Step 15 Verify that the service request information is correct.

a. If the information is not what you intended, click **Back** until you reach the provisioning dialog boxes in question, and edit the information as necessary. The complete the service request wizard as described in the previous sections.

b. When satisfied with the settings, click **Next**.

VPNSC displays the following message:

Your request to "Add VPN Service to CE" has been submitted with ID number n. *This service request can be deployed by using the "Deploy Service Requests" wizard or by using the "Deploy VPN Service" item under the "Provisioning" option of a VPN service request report.*

c. Click **Close**.

You have now queued a service request. It is entered into the VPN Solutions Center Repository and placed in the "Requested" state.

Step 16 Deploy the Management VPN service request.

Step 17 Repeat this procedure for the PE-CE link in the target autonomous system.

Tip When you create the service request for the PE-CE link in the target autonomous system, you must specify the same VPN that you specified in this procedure.

CHAPTER 13

Provisioning MPLS VPNs on the Cisco Series DSL Switch

The Cisco DSL Switch

The Cisco Series DSL Switch product line is the latest generation of DSLAM (Digital Subscriber Line Access Multiplexer) that supports both IP and ATM features. The Cisco Series DSL switch product family includes Cisco 6015, 6160, and the 6260 with the NI-2 controller card.

The product lines are based on a high-performance, non-blocking ATM switching fabric that also encompasses IP Layer 3 services. The ATM switching fabric supports all ATM service classes including CBR, VBR, VBR-nrt, UBR, ABR, and ATM traffic management and shaping capabilities. IP Layer 3 services include MPLS, Layer 2 Tunneling Protocol (L2TP), and subscriber termination—RFC-1483 routed, Routed Bridged Encapsulation (RBE) or RFC 1483 bridged, Point-to-Point Protocol over ATM, and Point-to-Point Protocol over Ethernet (PPPoE).

DSL Switch Deployment Considerations

This section summarizes the current support limitations and restrictions for deploying the Cisco VPN Solutions Center for the Cisco Series DSL switch.

- VPN Solutions Center requires a subinterface to provision the connection between the PE and CE link.

  The IP DSL switch supports subinterfaces on the ATM interface and DSL modem port via the **PVP** command. Configuring subinterfaces is not typically required on the DSL modem port; thus, you may not want to configure extra commands for every DSL interface.

- The IP unnumbered interface is an important feature in DSL configuration. VPNSC supports the IP unnumbered feature with the following limitations:
  - VPNSC supports only IP unnumbered via a loopback interface.
  - VPNSC must manage the loopback interface value and IP address.

    The operator cannot control which loopback value gets assigned to IP unnumbered interfaces, such as the Virtual-Template interface, subscriber ATM interface, and the VRF definition. Without knowing which loopback interface VPNSC assigns, the operator must make an educated guess as to the interface value when creating the Point-to-Point Protocol over ATM (PPPoA) template data files.

  In a future release of VPNSC, the operator will be able to reference which loopback interface is assigned to an IP unnumbered interface.

- VPN Solutions Center supports only IP address management between the PE and CE network.

  VPNSC can automatically assign the IP address and manage the routing process between the PE and CE network. The IP address that VPNSC assigns is from a pool allocated by the operator. This address pool is usually the same one used in the RADIUS and the IOS IP local pool.

- AAA RADIUS—VPN Solutions Center does not support RADIUS authentication, IP address assignment, and VRF definition in the PPPoX to MPLS VPN environment.

  VPNSC requires subscriber authentication, the VRF definition lookup must be localized to the IP DSL switch, and VPNSC must manage the IP address assignment and routing process.

IP DSL Switch Configuration Overview

This section summarizes the role IP DSL switch plays in an MPLS VPN environment and the required steps to configure MPLS VPN on the DSLAM device. The IP DSL switch can only function as a provider edge (PE) router or a label edge router (LER) in an MPLS network; it cannot function as a label switch router (LSR). The current IOS release for DSLAM supports only the tag-switching protocol; it does not support the label distribution protocol. The DSLAM device terminates the customer edge (CE) routers sessions and carries the traffic into the MPLS VPN. The various processes the IP DSL switch use to route CE traffic into the MPLS VPN entail RFC1483 routed, Routed Bridged Encapsulation, Point-to-Point Protocol over ATM, and Point-to-Point Protocol over Ethernet.

To configure the IP DSL switch to become part of an MPLS network, the following steps are necessary:

1. Configure the IP DSL switch as a PE router.
   a. Enable Cisco Express Forwarding (CEF) switching.
   b. Set up the tag-switching protocol.
   c. Set up IGP routing (OSPF or IS-IS) between the PE and P routers.
   d. Set up BGP routing between the PEs.
2. Define the VRF for an MPLS VPN.
3. Configure the DSL modem interface to the customer site.
   a. Assign the DSL profile to the subscriber ATM interface.
   b. Assign IP address on CE and PE interface or use unnumbered interfaces.
   c. Assign the PVC on the CE and PE ATM interface.
   d. Configure access protocol and encapsulation method on the CE and PE ATM interface.
4. Associate the DSL interface to the VRF name.
5. Configure the PE-CE routing process.

Using VPN Solutions Center to Provision an IP DSL Device for MPLS VPNs

As described in the previous section, the first step in setting up the IP DSL switch is to configure the IP DSL switch to participate in the MPLS network. The configuration involves enabling CEF switching, setting up tag-switching, setting up IGP routing and BGP routing between the PE and P routers, including BGP neighbors and VPNv4 BGP neighbors.

After the initial setup of the IP DSL switch, additional DSL configuration is required. Refer to the documentation available on Cisco Connection Online (CCO) for more details on additional DSL configuration. The initial setup is usually a one-time process and can be accomplished through the IOS command line interface (CLI) or by employing an element manager. The configuration information from the initial setup is loaded into VPNSC and used as a basis to configure MPLS VPNs.

Steps 2 through 5 in the "IP DSL Switch Configuration Overview" section on page 13-2 describe how VPN Solutions Center adds sites to an MPLS VPN network. For more information on VPN Solutions Center, refer to the product documentation on CCO at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/rtrmgmt/vpnsc/index.htm

Tasks to Provision IP DSL with VPNSC

The tasks to provision IP DSL with the VPN Solutions Center software are as follows:

1. Import the target (PE and CE routers) configuration files into VPNSC.
   a. Setup or verify the username and password information for all the targets.
   b. Setup or verify the SNMP community information for all the targets.
2. Set up a new Provider Administrative Domain (PAD).

   PADs are the IBGP core of an MPLS/VPN network.
   a. Define the Provider Administrative Domain.
   b. Configure the starting Route Target (RT) and Route Distinguisher (RD) values.
   c. Define a Region and assign a IP address pool to use between PE and CE links.
3. Setup a new customer and customer site.
   a. Define the customer name and customer site(s).
   b. Associate a DSL CPE or CE with the site.
   c. Define the CE/CPE device as an unmanaged device.
4. Define the VRF name for a MPLS VPN.
5. Provision the PE and CE associated with the VRF defined in the previous step.
   a. Select an unmanaged DSL CPE.
   b. Select the corresponding IP DSL switch (PE router).

      The provisioning of the PE and CE link involves using both VPNSC configlets and template files. The template contains configuration commands for the various DSL access protocols and encapsulation types.

      The following section examines the templates for the different access protocols and the encapsulation used in DSL.

   The operator repeats the operations in steps 3, 4, and 5 when he or she provisions VPN service for a customer.

Creating a Service Request for an RFC-1483 Routed Template

Configuring MPLS VPN mapping of RFC-1483 routed sessions is similar to configuring MPLS VPNs on other Cisco MPLS platforms. For general MPLS VPN configuration tasks, examples, and command references, consult the *MPLS Virtual Private Networks and MPLS Virtual Private Network Enhancements* feature modules.

- For details on the procedure for adding a service request for a PE-CE link, see the "Adding a Service for a PE-CE Link" section on page 6-6.
- For details on deploying a service request, see the "Deploying Service Requests" section on page 6-15.

To provision a service request to add an RFC 1483 Routed CE to an MPLS VPN, follow these steps:

Step 1 Completes Steps 1 through 4 as described in the "Tasks to Provision IP DSL with VPNSC" section on page 13-3.

Step 2 Start the VPN Solutions Center in MPLS mode (see the "Starting the VPN Solutions Center Software" section on page 3-1).

Step 3 From the VPN Console, choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor is displayed (see Figure 13-1).

Figure 13-1 MPLS Service Request Editor

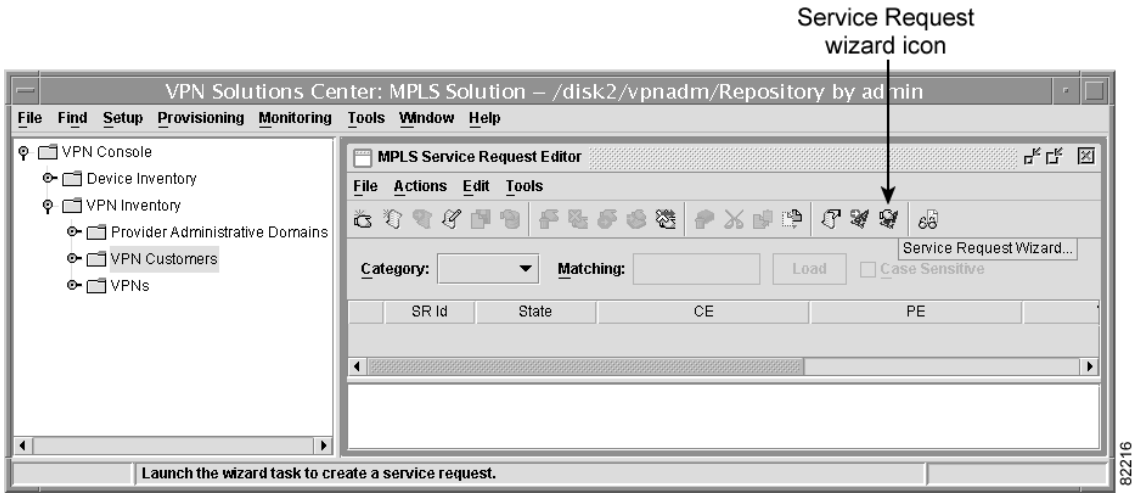

Step 4 To switch to the VPNSC 2.1 service request wizard, click the Service Request wizard icon.

The first—and informational only—screen appears.

Step 5 Click **Next**.

The Select CE dialog box appears.

Step 6 From the Select CE dialog box, select the customer edge router for this link.

Step 7 From the Select PE dialog box, select the provider edge router for this link.

Step 8 From the Select VPN: CERC Memberships dialog box, select the appropriate VPN from the list and specify the VPN topology.

Step 9 From the Select Routing Policy dialog box, specify a static route by choosing the **Static** radio button.

When you complete the Routing Policy wizard and click **Next**, the Redistribution dialog box appears.

If protocol redistribution is not required on this link, click **Next.**

The dialog box shown in Figure 13-2 appears.

Figure 13-2 The Select PE-CE Interfaces Dialog Box

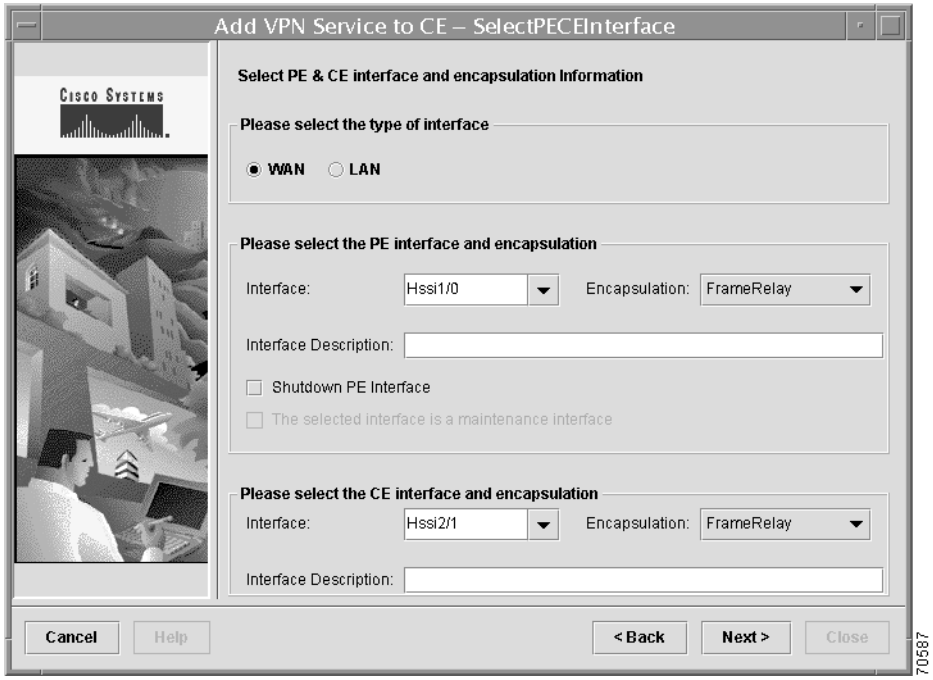

Step 10 Define the interfaces for the PE and CE.

a. Specify the type of interfaces for the PE-CE link: **WAN** or **LAN**.

b. Select the PE interface and its encapsulation method from the drop-down lists.

For both the PE and CE, select the **ATM (AAL5SNAP)** encapsulation method.

c. Specify the CE's DSL interface and its protocol encapsulation from the drop-down lists, then click **Next**.

Step 11 Enter the ATM PVC for the IP DSL switch (PE) and DSL CPE (CE) interfaces.

In most cases, use **subinterface 1**, **VPI 1**, and **VCI 32** for both the PE and CE.

Tip The subinterface and the VPI value must be identical. VPN Solutions Center fetches this value from the Repository and uses the VPI value as the variable in the DSLAM template.

The next dialog box (see Figure 13-3) provides a way to define the IP addressing scheme that is appropriate for this PE-CE link.

Figure 13-3 IP Addressing Scheme Dialog Box

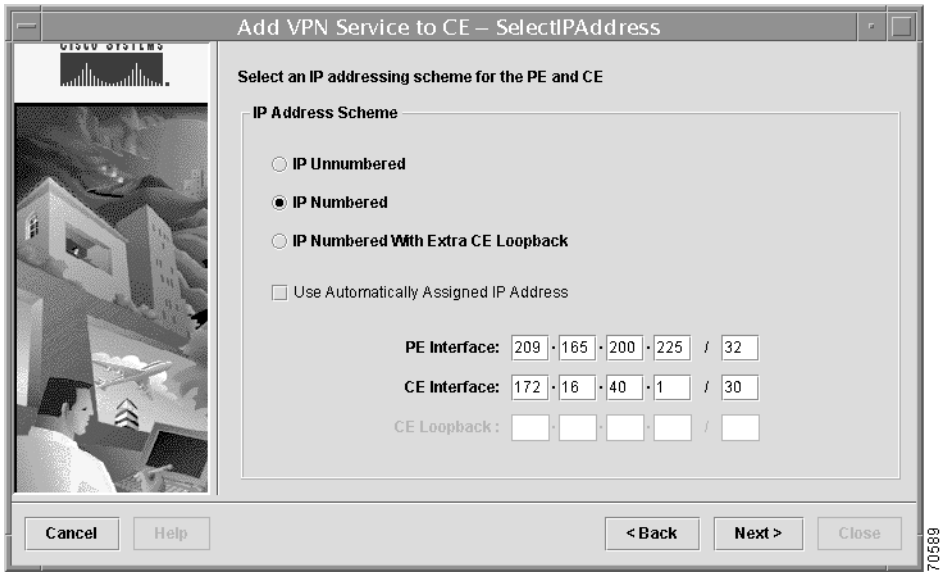

Step 12 Select the **IP Unnumbered** IP addressing scheme for both the PE and CE.

The subscriber ATM interface is IP unnumbered to the loopback interface.

When you choose **IP unnumbered**, VPN Solutions Center automatically creates a loopback interface (unless a loopback interface already exists with the correct attributes).

Step 13 Skip the next two dialog boxes, the *Select VRF Parameters* and the *Class of Service Profiles* dialog boxes, by clicking **Next** at each.

The next part of the service request wizard provides a way to integrate a template with VPN Solutions Center configlets for the PEs and CEs in the RFC-1483R environment. For details on creating and employing VPN Solutions Center templates, see Chapter 10, "Provisioning with the VPN Solutions Center Template Manager."

Step 14 Skip the setup for the CE template dialog boxes by clicking **Next** twice.

This aspect of the configuration must be completed using the command line interface.

The Select PE Data Files dialog box appears (see Figure 13-4).

Figure 13-4 PE Data Files Dialog Box After Selecting a Template Data File

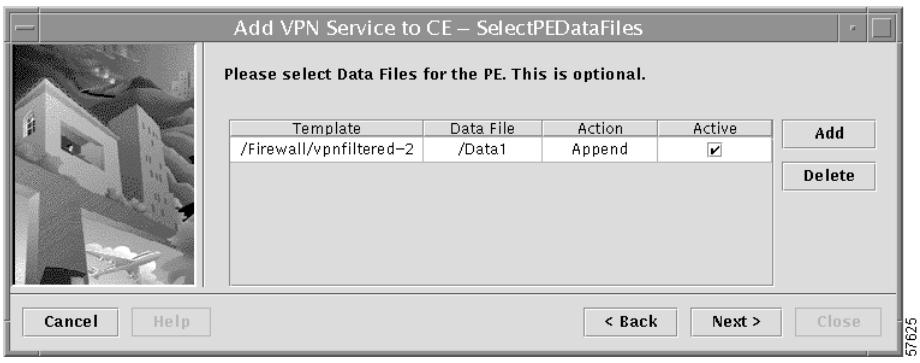

The **Action** column in the dialog box lets you specify where the template configuration file is placed in the VPN Solutions Center configlet—either prepended or appended.

The **Active** column lets you determine whether you want the template configuration file to be merged with the VPN Solutions Center configlet and downloaded to the target device.

Step 15 Specify the prepend template file (**rfc1483_routed_prepend**).

The PE uses one template. The prepend template file (**rfc1483_routed_prepend**) specifies commands to create the subinterface on the DSL modem interface using the **PVP** command. For details on this template, see the "RFC 1483 Routed Template" section on page 13-15.

The final screen in the wizard displays a summary of all the service settings defined for this VPN, including the information on template provisioning for the CE and PE.

Step 16 Verify the service request details and close the wizard.

Deploying the Service Request to the Network

To deploy the service request:

Step 1 From the VPN Console menu, choose **Provisioning > List All Service Requests**.

The All VPN Service Requests Report appears.

Step 2 Check the service request created in the previous section to verify that it is in the Requested state.

Step 3 From the All VPN Service Requests Report, click the **Provisioning** button at the bottom of the screen.

The Provisioning menu appears.

Step 4 From the Provisioning menu, choose **Deploy VPN Service**.

You can then view the VPN deployment status and check for errors from a web browser on the VPNSC workstation.

Step 5 From the VPN Console menu, choose **Provisioning > View Deployment Log**.

A Netscape web browser launches and requires that you to log in (if you haven't already done so).

a. Enter the administrative username and password, then click **OK**.

The default administrative username and password is **admin** (for both).

The Task Logs page is displayed. The task logs are displayed in sets of 10 logs per page. To jump to the next page of task logs, click the **Next** link (in the upper left corner of the task logs table).

The tasks are listed in order of the task start time; the task with the latest start time is listed at the top of list, and the task with the earliest start time is listed at the bottom of the list. The *Status* column gives an indication as to whether the tasks are completed successfully.

Notice the *Logs* column, which provides a **Log** link for every task listed.

b. Click the **Log** link to view the details for each task.

The information in the Status column indicates whether the task were completed successfully.

c. From the Actions frame (in the lower left quadrant), click the **DeployServiceRequest** link.

d. Observe the **Stdout/Stderr | Errors** on the frame on the right to get detailed information about deployment and to view the IOS commands that were downloaded to the router.

e. If there are errors, correct them and repeat the steps in this section ("Deploying the Service Request to the Network"). Most errors are clearly indicated in the deployment log.

Creating a Service Request for a Routed Bridged Encapsulation Template

ATM Routed Bridged Encapsulation (RBE), also known as ATM half-bridging, is the process of routing traffic from a bridged LAN without the use of integrated routing and bridging (IRB). RBE was developed to address the known RFC1483 bridging issues, including broadcast storms and security. Except for the fact that it operates exclusively over ATM, the RBE feature functions identically to half-bridging. Additional scalability, performance, and security can be achieved by using the unique characteristics of xDSL subscribers.

To add a Routed Bridged Encapsulation CE to the MPLS VPN:

Step 1 Completes Steps 1 through 4 as described in the "Tasks to Provision IP DSL with VPNSC" section on page 13-3.

Step 2 Start the VPN Solutions Center in MPLS mode (see the "Starting the VPN Solutions Center Software" section on page 3-1).

Step 3 From the VPN Console, provision a choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor is displayed (see Figure 13-5).

Figure 13-5 MPLS Service Request Editor

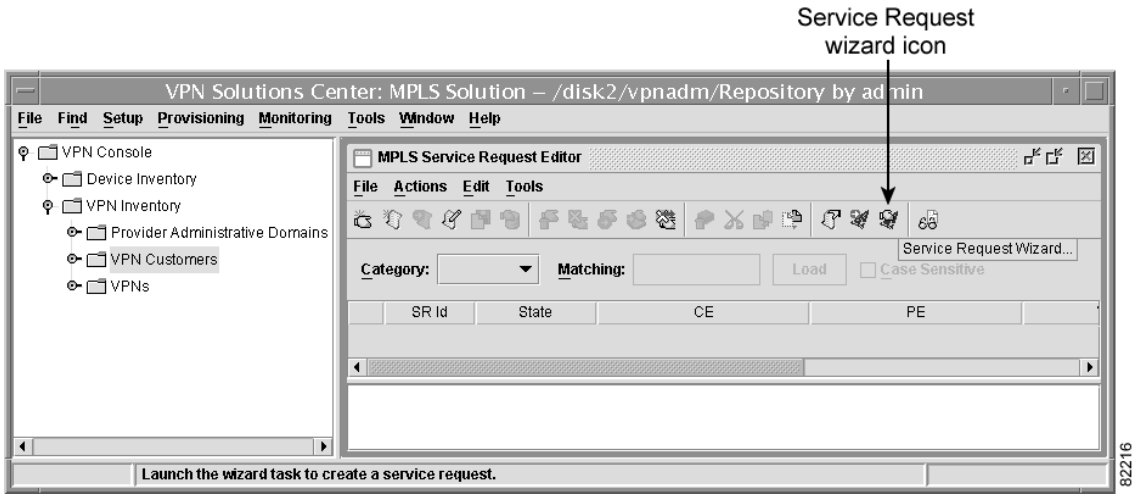

Step 4 To switch to the VPNSC 2.1 service request wizard, click the Service Request wizard icon.

The first—and informational only—screen appears.

Step 5 Click **Next**.

The Select CE dialog box appears.

Step 6 From the Select CE dialog box, select the customer edge router for this link.

Step 7 From the Select PE dialog box, select the provider edge router for this link.

Step 8 From the Select VPN: CERC Memberships dialog box, select the appropriate VPN from the list and specify the VPN topology.

Step 9 From the Select Routing Policy dialog box, specify **None** (because RBE is bridged).

However, a route is needed to get to the hosts connected to the CE's Ethernet interface. VPNSC will automatically add a static route to the CE link.

When you complete the Routing Policy wizard and click **Next**, the Redistribution dialog box appears.

If protocol redistribution is not required on this link, click **Next.**

Step 10 Define the interfaces for the PE and CE.

a. Select the PE interface and its encapsulation method from the drop-down lists.

   For both the PE and CE, select the **ATM (AAL5SNAP)** encapsulation method.

b. Specify the CE's DSL interface and its protocol encapsulation from the drop-down lists, then click **Next**.

Step 11 Enter the ATM PVC for the IP DSL switch (PE) and DSL CPE (CE) interfaces.

In most cases, use **subinterface 1**, **VPI 1**, and **VCI 32** for both the PE and CE.

Tip The subinterface and the VPI value must be identical. VPN Solutions Center fetches this value from the Repository and uses the VPI value as the variable in the DSLAM template.

The next dialog box provides a way to define the IP addressing scheme that is appropriate for this PE-CE link.

Step 12 Select the **IP Unnumbered** IP addressing scheme for both the PE and CE.

The subscriber ATM interface is IP unnumbered to the loopback interface.

For RBE, the CE router does not need an IP address. Therefore, the IP address VPNSC assigns to the CE ATM interface can be used on PC host Ethernet interface. The PC default gateway is set to the PE subscriber ATM interface or loopback IP address.

When you choose **IP unnumbered**, VPN Solutions Center automatically creates a loopback interface (unless a loopback interface already exists with the correct attributes).

Step 13 Skip the next two dialog boxes, the *Select VRF Parameters* and the *Class of Service Profiles* dialog boxes, by clicking **Next** at each.

The next part of the service request wizard provides a way to integrate a template with VPN Solutions Center configlets for the PEs and CEs in the RFC-1483 Bridged (RBE) environment. For details on creating and employing VPN Solutions Center templates, see Chapter 10, "Provisioning with the VPN Solutions Center Template Manager."

Step 14 Skip the setup for the CE template dialog boxes by clicking **Next** twice.

(This aspect of the configuration will be completed using the CLI.)

The Select PE Data Files dialog box appears.

Step 15 Specify the prepend template file (**RBE_1483Br_prepend**).

The PE uses one template. The prepend template file (**RBE_1483Br_prepend**) specifies commands to create the subinterface on the DSL modem interface using the **PVP** command and RFC 1483 bridged mode.

The final screen in the wizard displays a summary of all the service settings defined for this VPN, including the information on template provisioning for the CE and PE.

Step 16 Verify the service request details and close the wizard.

Step 17 Deploy the service request as described in the "Deploying the Service Request to the Network" section on page 13-7.

Creating a Service Request for the Point-to-Point Protocol over ATM Template

In a Point-to-Point Protocol over ATM (PPPoA) to an MPLS VPN environment, the IP DSL switch creates virtual access interface that maps to a specific VRF. One physical virtual access interface is created for each DSL subscriber. The PPP session terminates into an IP connection. The configuration tasks for enabling MPLS VPN mapping of PPPoA sessions is similar to the dial access environment.

The PE router uses two templates. The *append template file* specifies commands to create a subinterface on the DSL interface. The *prepend template file* specifies commands to create a PPPoA virtual template.

To provision a service request to add a Point-to-Point Protocol over ATM CE to the VPN, follow these steps:

Step 1 Completes Steps 1 through 4 as described in the "Tasks to Provision IP DSL with VPNSC" section on page 13-3.

Step 2 Start the VPN Solutions Center in MPLS mode (see the "Starting the VPN Solutions Center Software" section on page 3-1).

Step 3 From the VPN Console, provision a choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor is displayed (see Figure 13-6).

Figure 13-6 MPLS Service Request Editor

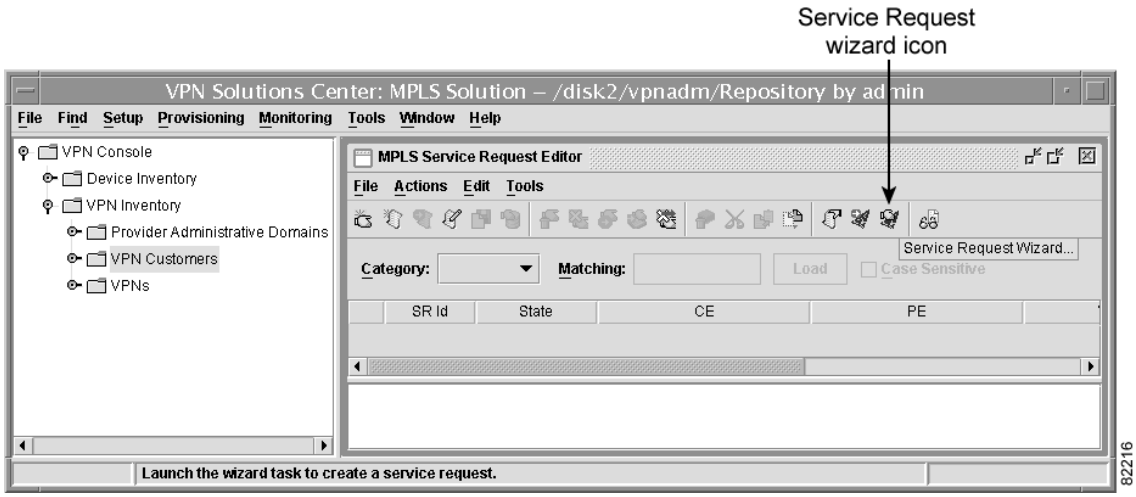

Step 4 To switch to the VPNSC 2.1 service request wizard, click the Service Request wizard icon.

The first—and informational only—screen appears.

Step 5 Click **Next**.

The Select CE dialog box appears.

Step 6 From the Select CE dialog box, select the customer edge router for this link.

Step 7 From the Select PE dialog box, select the provider edge router for this link.

Step 8 From the Select VPN: CERC Memberships dialog box, select the appropriate VPN from the list and specify the VPN topology.

Step 9 From the Select Routing Policy dialog box, specify a static route by choosing the **Static** radio button.

When you complete the Routing Policy wizard and click **Next**, the Redistribution dialog box appears.

If protocol redistribution is not required on this link, click **Next.**

Step 10 Define the interfaces for the PE and CE.

a. Select the PE interface and its encapsulation method from the drop-down lists.

For both the PE and CE, select the **ATM (AAL5SNAP)** encapsulation method.

b. Specify the CE's DSL interface and its protocol encapsulation from the drop-down lists, then click **Next**.

Step 11 Enter the ATM PVC for the IP DSL switch (PE) and DSL CPE (CE) interfaces.

In most cases, use **subinterface 1**, **VPI 1**, and **VCI 32** for both the PE and CE.

Tip The subinterface and the VPI value must be identical. VPN Solutions Center fetches this value from the Repository and uses the VPI value as the variable in the DSLAM template.

The next dialog box provides a way to define the IP addressing scheme that is appropriate for this PE-CE link.

Step 12 Select the **IP Unnumbered** IP addressing scheme for both the PE and CE.

The subscriber ATM interface is IP unnumbered to the loopback interface.

When you choose **IP unnumbered**, VPN Solutions Center automatically creates a loopback interface (unless a loopback interface already exists with the correct attributes).

Step 13 Skip the next two dialog boxes, the *Select VRF Parameters* and the *Class of Service Profiles* dialog boxes, by clicking **Next** at each.

The next part of the service request wizard provides a way to integrate a template with VPN Solutions Center configlets for the PEs and CEs in the ATM (PPPoA) environment. For details on creating and employing VPN Solutions Center templates, see Chapter 10, "Provisioning with the VPN Solutions Center Template Manager."

Step 14 Skip the setup for the CE template dialog boxes by clicking **Next** twice.

(This aspect of the configuration will be completed using the CLI.)

The Select PE Data Files dialog box appears.

Step 15 Specify the appropriate template files.

The PE router uses two templates. The *prepend template file* specifies commands to create a subinterface on the DSL interface. The *append template file* specifies commands to create a PPPoA virtual template.

- If you are provisioning PPPoA MUX, specify these templates:
  - *pppoa_mux_append*
  - *pppoa_mux_prepend*
- If you are provisioning PPPoA SNAP, specify these templates:
  - *pppoa_vt_append*
  - *pppoa_snap_prepend*

The final screen in the wizard displays a summary of all the service settings defined for this VPN, including the information on template provisioning for the CE and PE.

Step 16 Verify the service request details and close the wizard.

Step 17 Deploy the service request as described in the "Deploying the Service Request to the Network" section on page 13-7.

Creating a Service Request for a Point-to-Point Protocol over Ethernet Template

In a Point-to-Point Protocol over Ethernet (PPPoE) to MPLS VPN environment, the IP DSL switch creates virtual access interface that maps to a specific VRF. One physical virtual access interface is created for each DSL subscriber. The PPP session terminates into an IP connection. The configuration tasks for enabling MPLS VPN mapping of PPPoE sessions is similar to the dial access environment.

To provision a service request to add a Point-to-Point Protocol over Ethernet CE to the VPN, follow these steps:

Step 1 Completes Steps 1 through 4 as described in the "Tasks to Provision IP DSL with VPNSC" section on page 13-3.

Step 2 Start the VPN Solutions Center in MPLS mode (see the "Starting the VPN Solutions Center Software" section on page 3-1).

Step 3 From the VPN Console, choose **Provisioning > Add VPN Service to CE**.

The MPLS Service Request Editor is displayed (see Figure 13-7).

Figure 13-7 MPLS Service Request Editor

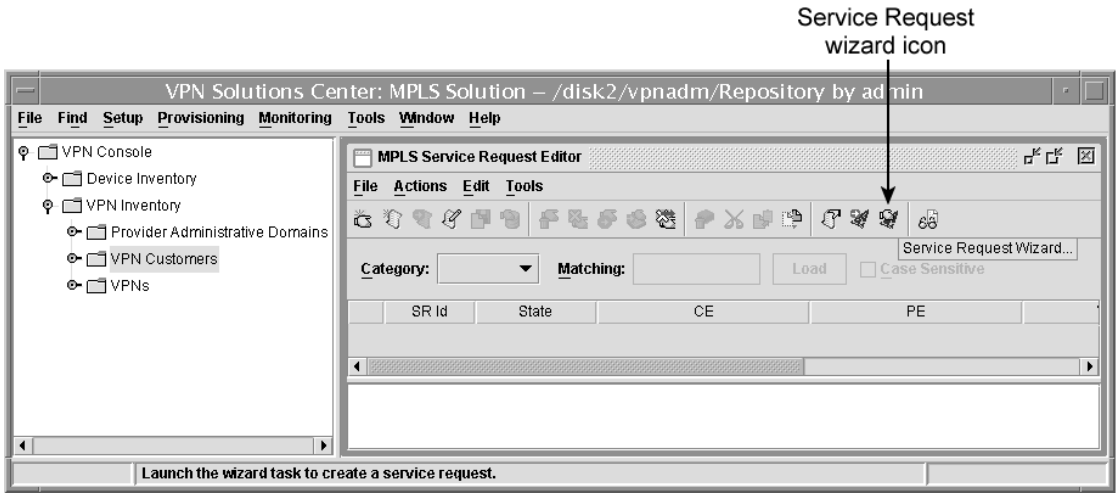

Step 4 To switch to the VPNSC 2.1 service request wizard, click the Service Request wizard icon.

The first—and informational only—screen appears.

Step 5 Click **Next**.

The Select CE dialog box appears.

Step 6 From the Select CE dialog box, select the customer edge router for this link.

Step 7 From the Select PE dialog box, select the provider edge router for this link.

Step 8 From the Select VPN: CERC Memberships dialog box, select the appropriate VPN from the list and specify the VPN topology.

Step 9 From the Select Routing Policy dialog box, specify a static route by choosing the **Static** radio button.

When you complete the Routing Policy wizard and click **Next**, the Redistribution dialog box appears.

If protocol redistribution is not required on this link, click **Next.**

Step 10 Define the interfaces for the PE and CE.

a. Select the PE interface and its encapsulation method from the drop-down lists.

For both the PE and CE, select the **ATM (AAL5SNAP)** encapsulation method.

b. Specify the CE's DSL interface and its protocol encapsulation from the drop-down lists, then click **Next**.

Step 11 Enter the ATM PVC for the IP DSL switch (PE) and DSL CPE (CE) interfaces.

Tip The subinterface and the VPI value must be identical. VPN Solutions Center fetches this value from the Repository and uses the VPI value as the variable in the DSLAM template.

The next dialog box provides a way to define the IP addressing scheme that is appropriate for this PE-CE link.

Step 12 Select the **IP Unnumbered** IP addressing scheme for both the PE and CE.

The subscriber ATM interface is IP unnumbered to the loopback interface.

When you choose **IP unnumbered**, VPN Solutions Center automatically creates a loopback interface (unless a loopback interface already exists with the correct attributes).

Step 13 Skip the next two dialog boxes, the *Select VRF Parameters* and the *Class of Service Profiles* dialog boxes, by clicking **Next** at each.

The next part of the service request wizard provides a way to integrate a template with VPN Solutions Center configlets for the PEs and CEs in the PPPoE environment. For details on creating and employing VPN Solutions Center templates, see Chapter 10, "Provisioning with the VPN Solutions Center Template Manager."

Step 14 Skip the setup for the CE template dialog boxes by clicking **Next** twice.

(This aspect of the configuration will be completed using the command-line interface.)

The Select PE Data Files dialog box appears.

Step 15 Specify the appropriate template files.

The PE router uses the following two templates:

- The append template file (*pppoe_vt_append*) specifies commands to create a PPPoA virtual template.
- The prepend template file (*pppoe_snap_prepend*) specifies commands to create a subinterface on the DSL interface.

The final screen in the wizard displays a summary of all the service settings defined for this VPN, including the information on template provisioning for the CE and PE.

Step 16 Verify the service request details and close the wizard.

Step 17 Deploy the service request as described in the "Deploying the Service Request to the Network" section on page 13-7.

Removing Template-Generated Statements From a Configuration File

When you decommission a service request, statements added to a configuration file through the Template Manager are not automatically removed from the file. To do this, you must create a template to remove template configuration statements. This type of template is called a *negate template*.

A negate template can handle the removal of configuration statements in one of two ways:

- Remove the IOS statements from the template that generates the statements in the PE router's configuration file.
- Manually log in to the PE router and delete the statements added by the template.

Tip If a PPoE or PPoA template was used to create a virtual template interface on the DSLAM provider edge router, you must remove the virtual template configuration prior to deploying the negate template that removes the service request. Otherwise, VPN Solutions Center does not completely remove the VRF commands and the loopback interface (because VPNSC sees that the VRF is applied to the virtual interface and therefore does not modify the VRF).

VPNSC Templates for the IP DSL Switch

This section examines the various templates used to configure the IP DSL switch for MPLS VPNs. The VPN Solutions Center Template Manager consists of the *template file*, *template data file,* and *template configuration file*.

- The *template file* is a file created by the Template Manager that stores a VPN Solutions Center template definition.
- The *template data file* is a text file (in XML format) that stores variable values to generate the template file.
- A *template configuration file* is an IOS configuration file that stores the Cisco IOS commands created by the Template Manager.

The template file and data file merge to form the template configuration file. The IP DSL switch uses the template file and the template data file for various DSL interface and encapsulation types. VPN Solutions Center creates the initial VPNSC configlet. The VPNSC configlet is a set of IOS commands for configuring MPLS VPNs. VPNSC downloads both the configlet and the template data file to form the set of commands used to configure the router (PE).

If the template data file is downloaded before the VPNSC configlet, it is considered *prepended*. On the other hand, if the template data file is downloaded after the VPNSC configlet, it is considered *appended*. The distinction between "prepend" and "append" is important because certain IOS commands must be issued before or after other IOS commands. For DSL, different types of access protocol and encapsulation need to be downloaded before or after the VPNS configlet.

The DSL templates are in XML format. The VPN Solutions Center Template Manager creates these files and stores them as flat files in the VPNSC Repository. The VPNSC installation utility places the files by default in */opt/vpnadm/Repository/Templates/ROOTDIR/*.

RFC 1483 Routed Template

The template for RFC 1483 routed issues commands for creating an ATM subinterface on the DSL modem interface or subscriber ATM interface. It is prepended to the VPNSC service request. The RFC 1483 template contents are as follows:

```
<Template>
<VarDeclaration>
<String VarName="interface_to_modem" />
<String VarName="subinterface_number" />
</VarDeclaration>

<SubTemplateDeclaration>
</SubTemplateDeclaration>

<MainTemplateDef>
interface #system.substringToDelim ($interface_to_modem, ".", 0)
atm pvp $subinterface_number
</MainTemplateDef>
</Template>
```

Every template file must contain a template data file. The template data file for RFC 1483 routed stores values for the subscriber ATM interface and the DSL modem subinterface.

Table 13-1 summarizes the template variables and descriptions.

Table 13-1 RFC 1483 Routed Template Variables

Variable Names	Description
$interface_to_modem	References the DSL modem subinterface. The value is the same as the VPI value on the ATM DSL modem interface.
$subinterface_number	Stores the VP tunnel value. A VP tunnel is a required to create a subinterface on a DSL modem interface or an ATM interface on the DSLAM device.

Cisco IOS Commands

The Cisco IOS commands are as follows:

```
C6260-2 (config)# interface ATM 1/2 ! DSL modem interface
C6260-2 (config-it)#atm pvp 1 ! ATM VP 1
```

Route Bridged Encapsulation (RBE-RFC 1483) Template

The template for RBE-RFC 1483 half-bridged issues commands for creating an ATM subinterface and specifies RFC 1483 bridged on the DSL modem interface. It is prepended to the VPNSC service request.

The RBE-RFC 1483 template contents are as follows:

```
<Template>
<VarDeclaration>
<String VarName="dsl_modem_interface" />
<String VarName="vci" />
<String VarName="vpi" />
<String VarName="subinterface_number" />
<String VarName="slash" />
<String VarName="dsl_modem_subinterface" />
</VarDeclaration>
<SubTemplateDeclaration>
</SubTemplateDeclaration>
<MainTemplateDef>
interface #system.substringToDelim ($dsl_modem_interface, ".", 0)
atm pvp $subinterface_number
interface $dsl_modem_subinterface
atm route-bridged ip
pvc $vpi+$slash+$vci
</MainTemplateDef>
</Template>
```

Cisco IOS Commands

The IOS commands for the RBE-RFC 1483 template are as follows:

```
interface ATM1/1
no ip address
atm clock INTERNAL
no atm ilmi-keepalive
atm pvp 1
!
interface ATM1/1.1 point-to-point
description ATM1/1.1 atm pvc vpi=1 vci=32 : Provisioned By VPNSC: Service
Request Id# = 71
ip vrf forwarding V19:CISCO
ip unnumbered Loopback1
atm route-bridged ip
no atm ilmi-keepalive
pvc 1/32
encapsulation aal5snap
```

Table 13-2 describes the variables used in the RBE-RFC 1483 template.

Table 13-2 RBE Subscriber ATM Interface Variables

Variable Names	Description
$dsl_modem_interface	References the DSL modem subinterface. The value is the same as the VPI value on the ATM DSL modem interface.
$dsl_modem_subinterface	References the DSL modem subinterface. The value is the same as the VPI value on the ATM DSL modem interface.
$subinterface_number	Stores the VP tunnel value. A VP tunnel is a required to create a subinterface on a DSL modem interface or an ATM interface on the DSLAM device.
$vpi	Stores the ATM PVI value for the subscriber ATM interface.
$vci	Stores the ATM VCI value for the subscriber ATM interface.
$slash	Holds the "/" character used in the `ATM slot#/port#` command.

Point-to-Point Protocol over ATM (PPPoA) Templates

The templates for the Point-to-Point protocol (PPP) over ATM issue commands for creating a Virtual Template interface, ATM subinterfaces; the templates specify either AAL5SNAP or AAL5MUX encapsulation on the DSL modem interface. The PPP over ATM templates are:

- pppoa_vt_append
- pppoa_snap_prepend
- ppoa_mux_append
- pppoa_mux_prepend

The pppoa_vt_append and the ppoa_mux_append templates are appended to the VPN service request during deployment. It is required only once—when the operator initializes provisioning of the PPPoA subscriber or the CE router. All subsequent PPPoA CEs connecting to the same VPN can share the same virtual template.

Both the pppoa_snap_prepend and pppoa_mux_prepend templates are prepended to the VPNSC service request during deployment.

Cisco IOS Commands

The IOS commands associated with these templates are as follows:

```
c6260-2#sh run | begin interface Virtual-Template1
interface Virtual-Template1 !pppoa_vt_append template
ip vrf forwarding V17:ford
ip unnumbered Loopback3
no peer default ip address
ppp authentication pap chap
c6260-2#sh run | begin interface ATM1/2
interface ATM1/2 !pppoa_snap_prepend template
no ip address
atm clock INTERNAL
no atm ilmi-keepalive
atm pvp 1
!
interface ATM1/2.1 point-to-point
description ATM1/2.1 atm pvc vpi=1 vci=32 : Provisioned By VPNSC: Service
Request Id# = 66
ip vrf forwarding V17:ford
ip unnumbered Loopback3
no atm ilmi-keepalive
pvc 1/32
encapsulation aal5snap
protocol ppp Virtual-Template1
```

The pppoa_vt_append Template Contents

The pppoa_vt_append template contents are as follows:

```
<Template>
<VarDeclaration>
<String VarName="loopback_int" />
<String VarName="Virtual_template" />
<String VarName="vrf_name" />
</VarDeclaration>

<SubTemplateDeclaration>
</SubTemplateDeclaration>

<MainTemplateDef>
interface $Virtual_template
    ip vrf forwarding $vrf_name
    ip unnumbered $loopback_int
    no peer default ip address
    ppp authentication pap chap
</MainTemplateDef>
</Template>
```

Table 13-3 describes the variables used in the pppoa_vt_append template.

Table 13-3 pppoa_vt_append Template Variables

Variable Names	Description
$loopback_int	Stores the loopback interface used in PPPoA virtual template. This value must match the VPNSC CE router or subscriber ATM interface used in an unnumbered configuration.
$Virtual_template	Stores the PPPoA virtual-template interface value.
$vrf_name	Stores the value of VRF name associated with the subscriber ATM interface. The value is reference to VPNSC internal variable, $MPLSvrfName.

The pppoa_snap_prepend Template Contents

The pppoa_vt_prepend template contents are as follows:

```
<Template>
<VarDeclaration>
<String VarName="interface_to_modem" />
<String VarName="vci" default="32" />
<String VarName="vpi" default="1" />
<String VarName="subinterface_number" />
<String VarName="Virtual_template" />
<String VarName="slash" default="/" />
<String VarName="interface_subscriber_port" />
</VarDeclaration>
<SubTemplateDeclaration>
</SubTemplateDeclaration>
<MainTemplateDef>
interface #system.substringToDelim ($interface_to_modem, ".", 0)
atm pvp $subinterface_number
interface $interface_subscriber_port
pvc $vpi+$slash+$vci
encapsulation aal5snap
protocol ppp $Virtual_template
</MainTemplateDef>
</Template>
```

Table 13-4 describes the variables in the pppoa_snap_prepend template.

Table 13-4 ppoa_snap_prepend Template Variables

Variable Names	Description
$interface_subscriber_port	The interface name for the MPLS PE.
$interface_to_modem	References the DSL modem subinterface. The value is the same as the VPI value on the ATM DSL modem interface.
$slash	Holds the "/" character used in the `ATM` **slot#/port#** command.
$subinterface_number	Stores the VP tunnel value. A VP tunnel is a required to create a subinterface on a DSL modem interface or an ATM interface on the DSLAM device.
Virtual_template	Stores the PPPoA virtual-template interface value.
$vpi	Stores the ATM PVI value for the subscriber ATM interface.
$vci	Stores the ATM VCI value for the subscriber ATM interface.

The pppoa_mux_append Template Contents

The pppoa_mux_append template contents are as follows:

```
<Template>
<VarDeclaration>
<String VarName="interface_subscriber_port" />
<String VarName="vrf_name" />
<String VarName="slash" />
<String VarName="vci" />
<String VarName="loopback_int" />
<String VarName="vpi" />
<String VarName="Virtual_template" />
</VarDeclaration>

<SubTemplateDeclaration>
</SubTemplateDeclaration>

<MainTemplateDef>
interface $interface_subscriber_port
pvc $vpi+$slash+$vci
encapsulation aal5mux ppp $Virtual_template

interface $Virtual_template
 ip vrf forwarding $vrf_name
 ip unnumbered $loopback_int
 no peer default ip address
 ppp authentication pap chap
</MainTemplateDef>
</Template>
```

Table 13-5 describes the variables in the pppoa_mux_append template.

Table 13-5 ppoa_mux_append Template Variables

Variable Names	Description
$interface_subscriber_port	The interface name for the MPLS PE.
$loopback_int	Stores the loopback interface used in PPPoA virtual template. This value must match the VPNSC CE router or subscriber ATM interface used in an unnumbered configuration.
$slash	Holds the "/" character used in the `ATM` **slot#/port#** command.
$vci	Stores the ATM VCI value for the subscriber ATM interface.
$Virtual_template	Stores the PPPoA virtual-template interface value.
$vpi	Stores the ATM PVI value for the subscriber ATM interface.
$vrf_name	Stores the value of VRF name associated with the subscriber ATM interface. The vrf_name value refers to the VPNSC internal variable, **$MPLSvrfName**.

The pppoa_mux_prepend Template Contents

The pppoa_mux_prepend template contents are as follows:

```
<Template>
<VarDeclaration>
<String VarName="interface_to_modem" />
<String VarName="vci" default="32" />
<String VarName="vpi" default="1" />
<String VarName="subinterface_number" />
<String VarName="Virtual_template" />
<String VarName="slash" default="/" />
<String VarName="interface_subscriber_port" />
</VarDeclaration>
<SubTemplateDeclaration>
</SubTemplateDeclaration>
<MainTemplateDef>
interface #system.substringToDelim ($interface_to_modem, ".", 0)
atm pvp $subinterface_number
interface $interface_subscriber_port
pvc $vpi+$slash+$vci
encapsulation aal5mux ppp $Virtual_template
</MainTemplateDef>
</Template>
```

Table 13-6 describes the variables in the pppoa_mux_prepend template.

Table 13-6 pppoa_mux_prepend Template Variables

Variable Names	Description
$interface_subscriber_port	The name of the interface for the MPLS PE.
$interface_to_modem	References the DSL modem subinterface. The value is the same as the VPI value on the ATM DSL modem interface.
$slash	Holds the "/" character used in the `ATM slot#/port#` command.
$subinterface_number	Stores the VP tunnel value. A VP tunnel is a required to create a subinterface on a DSL modem interface or an ATM interface on the DSLAM device.
$vci	Stores the ATM VCI value for the subscriber ATM interface.
$vpi	Stores the ATM PVI value for the subscriber ATM interface.

Point-to-Point Protocol over Ethernet (PPPoE) Templates

The templates for PPP over Ethernet issue commands to enable VPDN, virtual template, and the ATM subinterface and specify AAL5SNAP encapsulation on the DSL modem interface. The PPP over Ethernet templates are:

- pppoe_vt_append
- pppoe_snap_prepend

The pppoe_vt_append template is appended to the VPNSC configlet. It is required only once when the operator initializes provisioning the PPPoE subscriber or the CE router. All subsequent PPPoE CE routers connecting to the same VRF can share the same virtual template.

The pppoe_snap_prepend is prepended to the VPNSC MPLS configlet during configuration. If the operator is provisioning the PPPoE subscriber for the first time, the pppoe_snap_prepend template must be used in conjunction with pppoe_vt_append.

The template variables for PPP over Ethernet are the same as those for PPP over ATM.

The pppoe_vt_append Template Contents

The pppoe_vt_append template is appended to the VPNSC configlet. It is required only once when the operator initializes provisioning the PPP over Ethernet subscriber or the CE router. All subsequent PPPoE CE routers connecting to the same VRF can share the same virtual template.

The pppoe_vt_append template contents are as follows:

```
<Template>
<VarDeclaration>
<String VarName="loopback_int" />
<String VarName="Virtual_template" />
<String VarName="vrf_name" />
</VarDeclaration>
<SubTemplateDeclaration>
</SubTemplateDeclaration>
<MainTemplateDef>
vpdn enable
vpdn-group 1
accept-dialin
protocol pppoe
$Virtual_template
pppoe limit per-mac 101
pppoe limit per-vc 102
interface $Virtual_template
ip vrf forwarding $vrf_name
ip mtu 1492
ip unnumbered $loopback_int
no ip directed-broadcast
ip mroute-cache
no peer default ip address
ppp authentication pap chap
</MainTemplateDef>
</Template>
```

Table 13-7 describes the variables in the pppoe_vt_append template.

Table 13-7 pppoe_vt_append Template Variables

Variable Names	Description
$loopback_int	Stores the loopback interface used in PPPoA virtual template. This value must match the VPNSC CE router or subscriber ATM interface used in an unnumbered configuration.
$Virtual_template	Stores the PPPoA virtual-template interface value.
$vrf_name	Stores the value of VRF name associated with the subscriber ATM interface. The value is reference to VPNSC internal variable, $MPLSvrfName.

The pppoe_snap_prepend Template Contents

The pppoe_snap_prepend template is prepended to the VPNSC MPLS configlet during configuration. If the operator is provisioning the PPPoE subscriber for the first time, this template must be used in conjunction with the pppoe_vt_append template.

The pppoe_snap_prepend template contents are as follows:

```
<Template>
<VarDeclaration>
<String VarName="interface_to_modem" />
<String VarName="vci" default="32" />
<String VarName="vpi" default="1" />
<String VarName="subinterface_number" />
<String VarName="slash" default="/" />
<String VarName="interface_subscriber_port" />
</VarDeclaration>
<SubTemplateDeclaration>
</SubTemplateDeclaration>
<MainTemplateDef>
interface #system.substringToDelim ($interface_to_modem, ".", 0)
atm pvp $subinterface_number
interface $interface_subscriber_port
pvc $vpi+$slash+$vci
encapsulation aal5snap
protocol pppoe
</MainTemplateDef>
</Template>
```

Table 13-8 describes the variables in the pppoa_snap_prepend template.

Table 13-8 ppoa_snap_prepend Template Variables

Variable Names	Description
$interface_subscriber_port	The name of the interface for the MPLS PE.
$interface_to_modem	References the DSL modem subinterface. The value is the same as the VPI value on the ATM DSL modem interface.
$slash	Holds the "/" character used in the `ATM slot#/port#` command.
$subinterface_number	Stores the VP tunnel value. A VP tunnel is a required to create a subinterface on a DSL modem interface or an ATM interface on the DSLAM device.
$vci	Stores the ATM VCI value for the subscriber ATM interface.
$vpi	Stores the ATM PVI value for the subscriber ATM interface.

Cisco IOS Commands

The Cisco IOS commands for the pppoe_vt_append and pppoe_snap_prepend templates are as follows:

```
c6260-2#sh run | begin vpdn
vpdn enable
!
vpdn-group 1
accept-dialin
protocol pppoe
virtual-template 2
pppoe limit per-mac 101
pppoe limit per-vc 102
!
interface ATM1/1
no ip address
atm clock INTERNAL
no atm ilmi-keepalive
atm pvp 1
!
interface ATM1/1.1 point-to-point
description ATM1/1.1 atm pvc vpi=1 vci=32 : Provisioned By VPNSC: Service
Request Id# = 68
ip vrf forwarding V17:ford
ip unnumbered Loopback3
no atm ilmi-keepalive
pvc 1/32
encapsulation aal5snap
protocol pppoe
!
c6260-2#sh run | begin interface Virtual-Template2
interface Virtual-Template2
ip vrf forwarding V17:ford
ip unnumbered Loopback3
ip mtu 1492
ip mroute-cache
no peer default ip address
ppp authentication pap chap
```

CHAPTER 14

Repository Administration

This chapter discusses Repository administration for VPN Solutions Center 2.1. It contains the following topics:

Converting a VPN Solutions Center 1.x Repository to 2.x Format

The VPNSC 1.x Repository is not compatible with the VPNSC 2.x Repository format. Therefore, before you can run VPNSC 2.1, you must convert your existing 1.x Repository to the 2.x format.

When you convert the 1.x Repository to 2.1 format, the Repository Conversion utility creates a directory for the converted Repository database files called *2.xRep* and places it under the existing Repository directory. For example, if the 1.x Repository currently resides at */opt/vpnadm/Repository,* the Repository Conversion utility creates the new directory as: */opt/vpnadm/Repository/2.xRep*.

Note The existing 1.x Repository is *not* altered during the conversion process. The database conversion utility copies the existing database files into the */2.xRep* directory, then converts the files to the 2.x format.

To convert a VPN Solutions Center 1.x Repository to VPNSC 2.x format, follow these steps.

Step 1 Install the VPN Solutions Center 2.1 software as described in the *Cisco VPN Solutions Center Installation Guide*.

Step 2 From the VPN Solutions Center workstation, open a terminal window.

Step 3 Log in as the owner of the VPN Solutions Center software (*vpnadm).*

```
su - vpnadm
```

Or if you are logging in remotely, enter this command:

```
rlogin VPNSC_hostname -l vpnadm
```

Step 4 If the Watch Dog is running, go to the terminal window where it's running and shut down the Watch Dog utility with this command:

```
stopwd -y
```

Step 5 Go to the directory where VPN Solutions Center is installed. For example:

```
cd /opt/vpnadm/vpn/
```

Step 6 Issue the following command to source the vpnadm user environment as required for your shell.

C-shell: `source vpnenv.csh`

K-shell: `. vpnenv.sh`

Step 7 Change directory to the directory where the VPNSC 1.x Repository resides. For example:

```
cd /opt/vpnadm/vpn/Repository
```

Step 8 Enter the following at the command line to convert the existing 1.x database files:

```
convert1.xRep.sh
```

The Repository Conversion utility creates a new */2.xRep* directory under the existing Repository directory, copies the VPNSC 1.x Repository files to the new directory, and converts the files to VPNSC 2.1 format.

If you did not specify the */2.xRep* directory for the Repository when you installed the VPN Solutions Center 2.1 software, you must update the *rep.list* file with the new path to the 2.1 Repository, as explained in the next step.

Updating the Pointer to VPNSC Repository

Step 9 Change directory to the location of the *rep.list* file. For example:

```
cd /opt/vpnadm/vpn/etc
```

Step 10 Open the *rep.list* file with a text editor.

The following statement is displayed in the file:

```
DEFAULT "<Repository_path>" "Default Solaris repository"
```

Step 11 Update the path for the default Solaris repository to the new path for the VPN Solutions Center 2.1 Repository, then save the *rep.list* file and exit.

Note When you modify the path to the VPNSC Repository, be sure to keep the "DEFAULT" and "Default Solaris repository" strings intact. Altering or removing these text strings in the *rep.list* file prevents a number of key servers from launching.

Step 12 From the same */etc* directory, open the *Template.properties* file.

The following statements are displayed in the file:

```
TemplateRootDir=/<Repository_path>/Templates/ROOTDIR
TemplateRootDir=/<Repository_path>/Templates/TemplateTable.xml
```

Step 13 Update the Repository paths in the **TemplateRootDir** and **TemplateIndexXML** statements to the new path for the VPN Solutions Center 2.1 Repository, then save the *Template.properties* file and exit.

You can now start the VPN Solutions Center software as described in Chapter 3, "Starting and Stopping the VPN Solutions Center Software."

Converting a 2.0 Repository to 2.x Format

The conversion of a VPNSC 2.0 Repository to version 2.1 or 2.2 format requires very little updating from the version 2.0 database. But since minor changes are necessary, this conversion utility is provided.

When you convert the 2.0 Repository to 2.x format, the Repository conversion utility copies the appropriate files into the existing Repository and updates the *rep.list* file with the path to the updated Repository.

To convert a VPN Solutions Center 2.0 Repository to VPNSC 2.x format, follow these steps.

Step 1 Install the VPN Solutions Center 2.1 software as described in the *Cisco VPN Solutions Center Installation Guide*.

Step 2 From the VPN Solutions Center workstation, open a terminal window.

Step 3 Log in as the owner of the VPN Solutions Center software (*vpnadm).*

```
su - vpnadm
```

When logged in, you are placed in the */opt/vpnadm* directory.

Or if you are logging in remotely, enter this command:

```
rlogin VPNSC_hostname -l vpnadm
```

Step 4 If the Watch Dog is running, go to the terminal window where it's running and shut down the Watch Dog utility with this command:

```
stopwd -y
```

Step 5 Go to the directory one level above the VPN Solutions Center Repository directory. For example:

```
cd /opt/vpnadm/vpn
```

Step 6 Issue one of the following commands to source the vpnadm user environment as required for your shell.

C-shell: `source vpnenv.csh`

K-shell: `. vpnenv.sh`

Step 7 Change directory to the directory where the VPNSC 2.0 Repository resides. For example:

```
cd /opt/vpnadm/vpn/Repository
```

Step 8 Enter the following at the command line to convert the existing 2.0 database files:

```
convert2.0Rep.sh
```

The Repository Conversion utility converts the VPNSC 2.0 Repository files to VPNSC 2.1 format.

You can now start the VPN Solutions Center software.

Migrating Users and Passwords from VPNSC 1.x to 2.x

In VPN Solutions Center, the usernames and passwords are kept in the *passwordFile* file.

- In the 1.x version of VPNSC, the *passwordFile* file is kept in the <*installation_directory*>*/ vpn/etc* directory.
- In the 2.x version of VPNSC, the *passwordFile* file is kept in the <*installation_directory*>*/ Repository/users* directory.

To migrate the existing 1.x users and password information to VPN Solutions Center 2.x, copy the old *passwordFile* to the new location for VPNSC 2.x.

Backing Up the Repository

Backing up the Repository is managed through the Repository Management tools, which you can access through any Web browser.

The Backup tool backs up the VPN Solutions Center Repository, which includes all the database files, collected raw data sets, generated baseline reports, logs, service objects, and configlets, to a local directory on the VPN Solutions Center machine. The backup options include **Tar** (which stands for "tape archiver," even though tape is rarely the backup medium used these days) or **Tar and compress**.

To back up the Repository, follow these steps:

Step 1 From the VPN Console menu, choose **File > Repository Administration**.

The Cisco VPN Solutions Center Repository Management Tool window appears (see Figure 14-1).

Figure 14-1 VPN Solutions Center Repository Management Tool

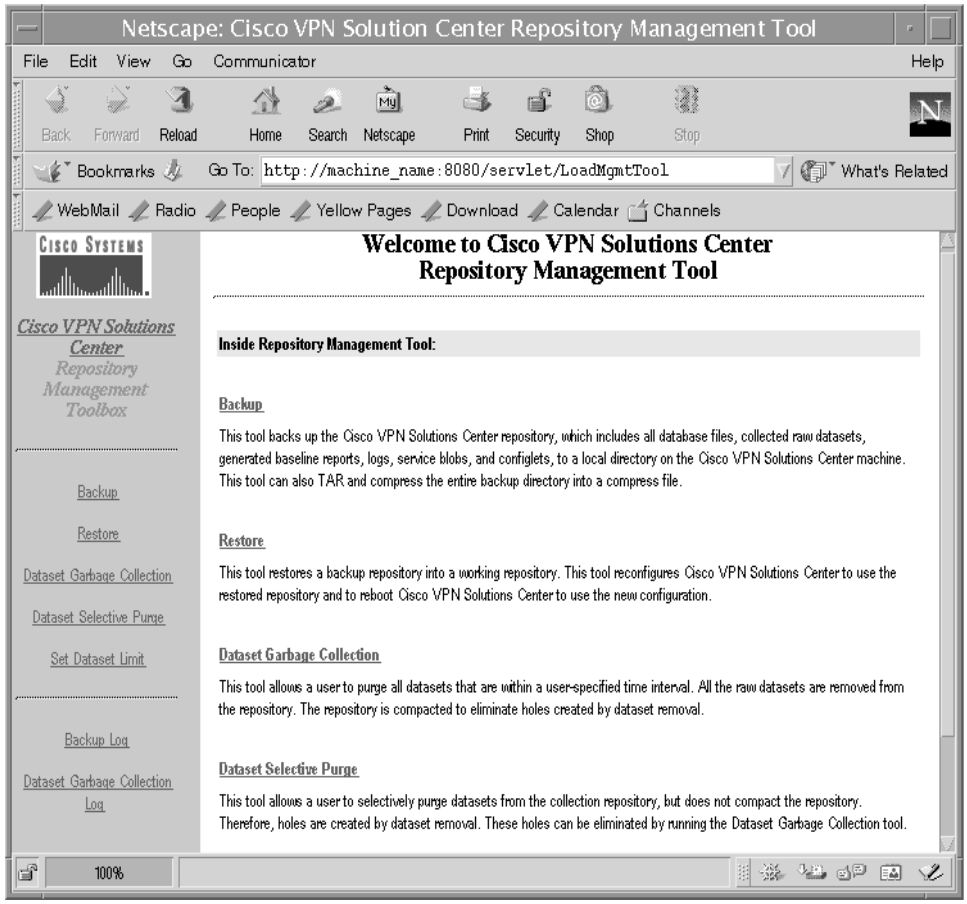

Step 2 Click **Backup**.

The Database Backup dialog box appears (see Figure 14-2).

Figure 14-2 Repository Backup Options Dialog Box

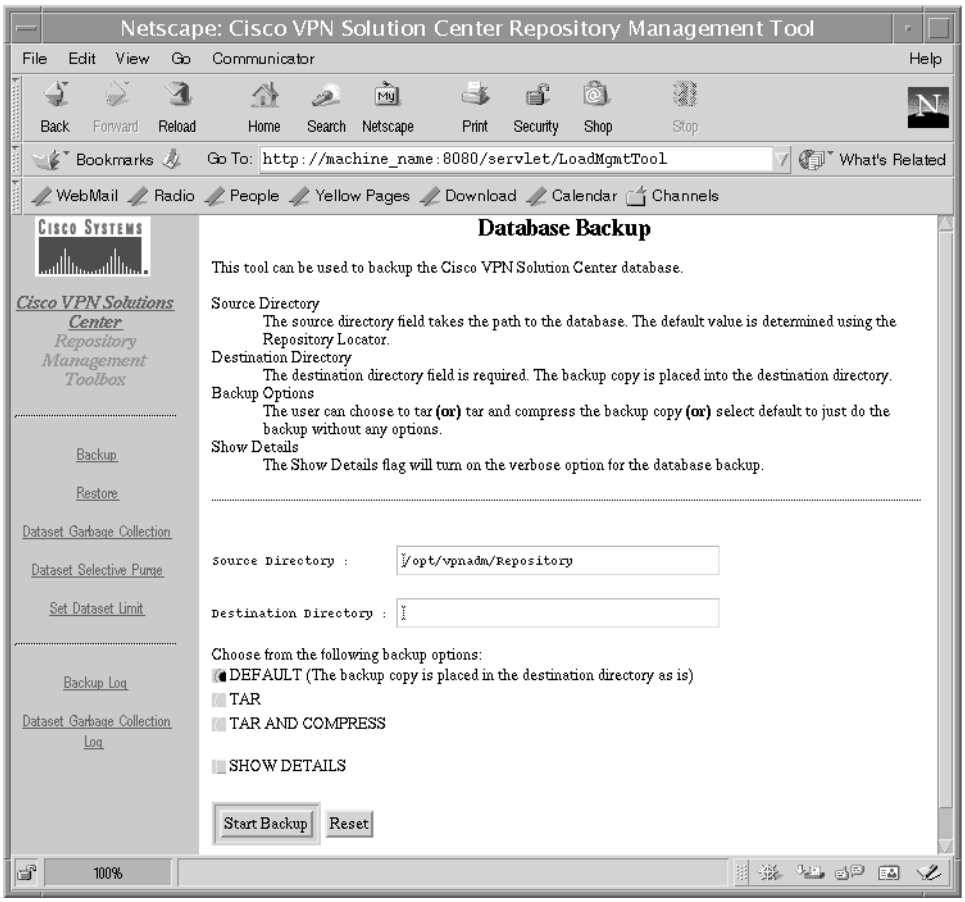

Step 3 *Source Directory*: Enter the path name for the Repository you want to back up.

The *Source Directory* field is required. This field defaults to the directory of the currently used Repository. If you choose a different Repository to back up, in this field place the full path name to the directory of the Repository that you want to back up.

Step 4 *Destination Directory*: Enter the full path name to the directory where you want to copy the Repository files.

Step 5 Determine the method you want for Repository backup by choosing one of the following:

- **Default**

  Choose this option if you want to back up the Repository and leave everything as is and copy the files to the destination directory.

- **Tar**

  Choose this option if you want to back up the Repository and copy it to a TAR file (so-called "tape archive") in the destination directory.

- **Tar and Compress**

  Choose this option if you want to back up the Repository and copy it to a compressed TAR file in the destination directory.

Step 6 If you want to turn on the verbose option when backing up the database, check the **Show Details** check box. This option gives you detailed progress information.

Step 7 Once you have completed the fields, buttons, and boxes in the Database Backup dialog box, click **Start Backup**.

Note *Optional*: To return the fields and other settings on the Database Backup dialog box to their default values, choose **Reset**.

Using the Database Backup Utility

You can use the Database Backup (dbBackup) utility to back up a VPN Solutions Center Repository from the command line. Run the **dbBackup** command from the *$ECSP_HOME/bin/solaris* directory.

The Database Backup utility also allows you to back up the Repository with a third-party backup program (by specifying the **-p** <*prog_name*> parameter), as described below. You do not have to specify a third-party program. If you do not specify the *prog_name*> parameter, the Database Backup utility uses its own program to backup the Repository.

When you execute the **dbBackup** command, the Database Backup utility does the following:

1. Locks the Repository.
2. If specified on the command line, the Database Backup utility executes the specified backup program. Otherwise, the utility runs its own backup program.
3. When finished, the Database Backup utility unlocks the Repository.

To run the Database Backup utility, follow these steps:

Step 1 From a terminal window on the VPNSC Workstation, log in as a valid VPNSC user.

Step 2 Execute the command as follows:

dbBackup -db <*db_path*> [**-dest** <*dest_dir*>] [**-tar**|**-compress**] [**-v**] [**-p** <*prog_name*>] [**-help**]

Syntax

The syntax for the **dbBackup** command is as follows:

dbBackup -db <*db_path*> [**-dest** <*dest_dir*>] [**-tar**|**-compress**] [**-v**] [**-p** <*prog_name*>] [**-help**]

-db <*dbpath*>	The path to the Repository you wish to back up.
-dest <*destDir*>	The existing destination directory into which to place the backed up Repository. If this parameter is not specified, the default destination directory is */tmp*.
-tar	Tars the backed up Repository into the specified destination directory.
-compress	Tars and compresses the backed up Repository into the specified destination directory.
-v	Specifies verbose log output.

-p <*prog_name*>	Specifies the third-party program <*prog_name*> to perform the Repository backup. The <*prog_name*> parameter must contain the complete command line parameters for the specified program.
-help	Displays the **dbBackup** syntax information

Using the Recover Tool Utility

The Recover Tool utility is intended for use when a Repository has been lost or corrupted. In this circumstance, it may not be possible to entirely recover the lost Repository from the most recent good backup. Thus, you would need to use the Recover Tool utility to recreate the Repository by recovering the Repository events that occurred between the time of the last good backup and the Repository loss. This utility plays back events recorded in the journal files in order to recreate a Repository.

To use the Recover Tool utility, follow these steps:

Step 1 If the VPN Console is running, shut it down by choosing **File > Exit**.

Step 2 Log in as a valid VPNSC user.

Step 3 From the terminal window where the Watch Dog is running, stop the journaling process with this command:

```
wdclient stopwd
```

Step 4 Change directory to the *$ECSP_HOME/bin/solaris* directory.

Step 5 Run the Recover Tool utility once for each set of journal files (which are organized by Repository directory) in timestamp order, starting with the earliest set of journal files.

The **recovertool** command has the following syntax:

execjava.sh netsys.repository.journal.recovertool.Main <*journal_dir*> <*rep_dir*>

<*journal_dir*>	The path to the directory containing the journal files *dir.jnl* and *vi.jnl*.
<*rep_dir*>	The complete path of the destination directory that contains the Repository on which the journal files are to be played back.

Step 6 Start the VPN Solutions Center software.

Restoring the Repository

The Cisco VPNSC Repository Management Tool provides a way to restore a backup Repository into a working Repository. The tool automatically reconfigures the VPN Solutions Center software to use the restored Repository and restarts the VPNSC Watch Dog servers so that they use the restored Repository.

To restore a backup Repository, follow these steps:

Step 1 From the VPN Console menu, choose **File > Repository Administration**.

As shown in Figure 14-1 on page 14-5, the Cisco VPN Solutions Center Repository Management Tool page appears.

Step 2 Shut down the VPN Console by choosing **File > Exit**.

Shutting down the VPN Console prevents other processes from writing to the Repository during the restoration process.

Step 3 Choose **Restore**.

The Database Restore dialog box appears (see Figure 14-3).

Figure 14-3 The Database Restore Dialog Box

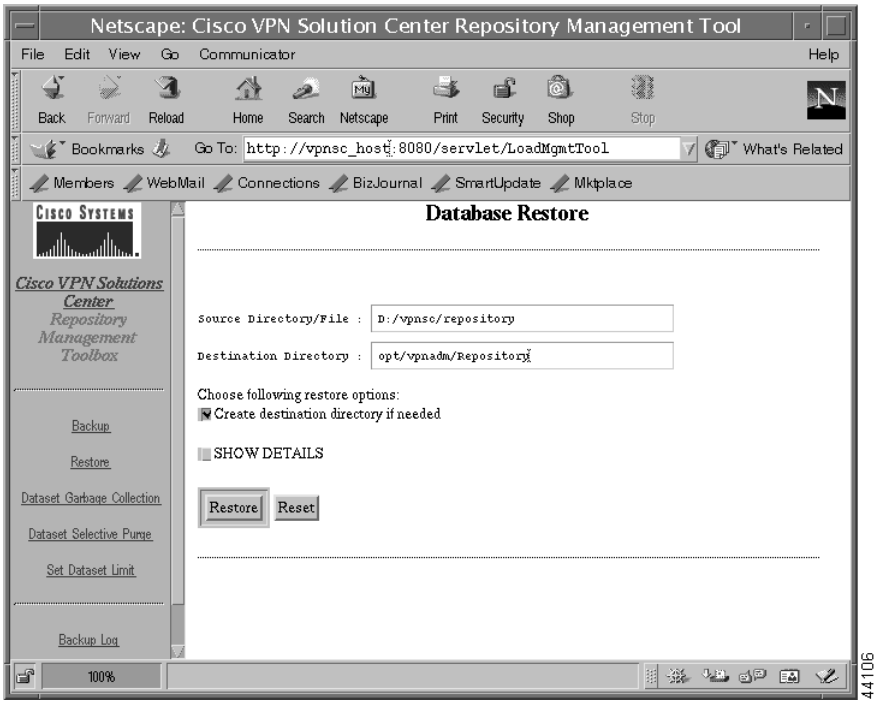

Step 4 *Source Directory/File*: Enter the directory (or path to a filename) from which you are restoring the Repository.

Step 5 *Destination Directory*: Enter the drive and path to the directory where you want to restore the Repository.

Caution The destination directory for the restoration operation must be empty. If you restore a Repository over an existing Repository, valuable data may be lost and service requests may not function correctly.

Step 6 Set the following Restore operation options if desired:

- If you want VPN Solutions Center to create the destination directory, check the **Create destination directory if needed** check box.
- To provide detailed progress information regarding the Restore operation, check the **Show Details** check box.
- If you wish to erase the values you have entered so far, click **Reset**.

Step 7 When you have completed the fields, click **Restore**.

You receive the prompt, "Do you really want to continue?"

Step 8 To proceed with the Repository restoration operation, click **OK**.

The next page provides the following information:

```
The Database Restore Process has started!
The Database backup file is path/filename.
Restore to directory pathname.
The execution string is dbRestore source_path ~dest destination_path -c
Click here to look at the Status.
```

Step 9 To check the status of the Restore operation, choose the **Status** link.

The Cisco VPNSC Repository Management Tool displays the Status page.

Using the Database Restore Utility

The Database Restore (dbRestore) utility restores a Repository from an existing backup file. Run the **dbRestore** command from the *$ECSP_HOME/bin/solaris* directory.

The syntax for the **dbRestore** command is as follows:

dbRestore <*backup_file*> **-dest** <*dest_Dir*> [**-c**] [**-v**] [**-help**]

<*backup_file*>	The path to the backup file or directory that you wish to restore from.
-dest <*dest_Dir*>	The destination directory (complete path) into which the restored Repository should be placed.
-c	Creates the destination directory, as specified, if it does not exist.
-v	Supplies verbose log output.
-help	Displays the **dbRestore** syntax information.

Using a Third-Party Application to Restore a Repository and Journal Files

To use a third party program to restore a Repository and journal files, follow these steps:

Step 1 If the VPN Console is running, shut it down by choosing **File > Exit**.

Shutting down the VPN Console removes any locks on the Repository.

Step 2 If the Watch Dog user interface (wdgui) is running, close it by selecting the window, right-click, then choose **Close** from the menu.

Step 3 From the window where Watch Dog was launched, close the Watch Dog by issuing this command:

```
stopwd -y
```

Step 4 Uncompress and untar the Repository as follows:

a. **uncompress** <*BackupDirectoryPath*> **Repository**.*DD_MM_YYYY_HH_MM_SS*.**tar**

b. **cd** <*RepositoryPath*>

c. **tar xvf** <*BackupDirectoryPath*>**Repository**.*DD_MM_YYYY_HH_MM_SS*.**tar**

Step 5 Uncompress and untar the journal as follows:

a. **uncompress** <*BackupDirectoryPath*> **journal**.*DD_MM_YYYY_HH_MM_SS*.**tar**

b. **cd** <*RepositoryPath*>

c. **tar xvf** <*BackupDirectoryPath*>**journal**.*DD_MM_YYYY_HH_MM_SS*.**tar**

Step 6 Restart the Watch Dog and the VPN Console.

For instructions, see the "Starting the Watch Dog and the VPN Console" section on page 3-2.

Using the VPNSC Repository Import/Export Utility

VPN Solutions Center provides a utility that allows you to populate the Repository by importing information from an XML file or exporting an existing Repository to an XML file. You can run this utility either on the VPN Solution Center workstation or on a remote machine.

Importing the Repository From an XML File

Using the VPN Inventory Repository (*VpnInvImport)* utility, you can populate the VPN Solutions Center VPN Inventory Repository by reading information from a specified XML file.

The */tmp/VpnInvImport.log* file contains information about each element that was created. In addition, this Repository utility writes output information to **stdout** for the Inventory Elements that are created, as well as the elements that already exist. If an error occurs while populating the Repository, the program exits and a message is sent to **stderr**.

Before Running the Import Utility

Complete the following steps before implementing the VPN Inventory Repository utility commands:

Step 1 Make sure the VPN Solutions Center Watch Dog is running.

Step 2 Log in as the **vpnadm** administrative user.

Step 3 In the $ECSP_HOME directory, source the files as follows.

C-shell: `source vpnenv.csh`

K-shell: `. ./vpnenv.sh`

Note Perform Steps 4 and 5 when running VPN Solutions Center as a user *other* than the VPN Solutions Center administrative user (vpnadm).

Step 4 If you choose to run the utility on a remote machine rather than on the same machine that runs VPN Solutions Center, be sure that the permissions are set correctly on the Naming Server (NS) and the VPN Inventory Server (VpnInvServer) is running on the VPN Solutions Center workstation.

For example, to allow a user (<*username*>) to run this utility on a remote machine, enter the following commands on the VPN Solutions Center workstation:

- **chmodit NS i+** *username*
- **chmodit VpnInvServer i+** *username*

Step 5 To allow the user to launch either of these two servers, issue this command:

chmodit l+ *username*

Step 6 If the environment variable is not already set, enter the following command:

setenv IT_DAEMON_PORT 1570

Running the VPN Inventory Import Utility

To run the VPN Inventory Import utility, follow these steps:

Step 1 Be sure you have implemented the prerequisites as specified in the previous section.

Step 2 Log in as the **vpnadm** user.

Step 3 From the VPN Solutions Center workstation, issue the following command:

VpnInvImport <*filename*> [**-u** <*username*>] [**-p** <*password*>] [**-f** <*password_file*>] [<*hostname*>]

Note If the format of the input file (the Document Type Definition) is not known, run the VPN Inventory Export utility as described in the "Exporting a Repository to an XML File" section on page 14-13. All values are case sensitive. Therefore, be sure the Boolean variables are all specified in lowercase.

Also be sure that the ampersand character (**&**) is replaced by **&** before using XML import format. This is required because the character **&** is a reserved character in XML.

<*filename*>	The name of the XML file from which to read data.
-u <*username*>	The name of the valid user on the VPN Solutions Center workstation. If this optional parameter is not specified, the default for the valid user name is **admin**.
-p <*password*>	The password for the valid user specified by the **-u** option above or the password for the default user **admin**. If this optional parameter is not specified, the default for the password is **admin**.

-f <*password_file*>	The file during VPN Solutions Center installation against which the <*username*> and <*password*>, specified above, is validated. If this optional parameter is not specified, the default for the password_file is the password_file created during VPN Solutions Center installation.
<*hostname*>	The name of the machine where the Naming Server is running, for example, **abc.company.com**, where **abc** is the host name. When this optional parameter is not specified, the default is **localhost.**

To get help on the Import utility parameters, enter this command:

```
VpnInvImport -h
```

Exporting a Repository to an XML File

You can export the VPN Solution Center Repository into an XML file using the VPN Inventory Export (*VPNInvExport)* utility. You can run this utility either on the VPN Solution Center workstation or on a remote machine.

Before You Run the Export Utility

Before you run the VPN Inventory Export utility, set the permissions on both the Naming Server (NS) and the VpnInvServer running on the VPN Solution Center workstation so that the VPNInvExport utility running on the remote machine can invoke them.

For example, to allow a user to use the utility, you would issue the following commands on the VPN Solution Center workstation:

Step 1 Make sure the VPN Solutions Center Watch Dog is running

Step 2 Log in as the **vpnadm** user.

Step 3 In the $ECSP_HOME directory, source the files as follows.

C-shell: `source vpnenv.csh`

K-shell: `. ./vpnenv.sh`

Step 4 Issue the following commands:

- **chmodit NS i+** *username*
- **chmodit VpnInvServer i+** *username*

Step 5 To allow the user to launch either of these two servers, issue this command:

```
chmodit l+ username
```

For the Repository export utility to work, you must set the environment variable *IT_DAEMON_PORT.*

Step 6 If the environment variable is not already set, issue this command:

```
setenv IT_DAEMON_PORT 1570
```

Running the VPN Inventory Export Utility

Be sure you have implemented the prerequisites, as specified in the previous section.

From the VPN Solutions Center workstation, logged on as the **vpnadm** user, issue the following command:

VpnInvExport *<filename>* [**-u** *<username>*] [**-p** *<password>*] [**-f** *<password_file>*] [*<hostname>*]

<filename>	The XML file to which the Repository is to be exported.
-u *<username>*	The name of the valid user on the VPN Solutions Center workstation. If you do not specify this optional parameter, the default for the user name is **admin**.
-p *<password>*	The password for the valid user specified by the **-u** option above or the password for the default user **admin**. If you do not specify this optional parameter, the default for the password is **admin**.
-f *<password_file>*	The file during VPN Solutions Center installation against which the *username* and *password*, specified above, are validated. If you do not specify this optional parameter, the default for the password file is the password file created during VPN Solutions Center installation.
<hostname>	The name of the machine where the Naming Server for VPN Solution Center is running. For example, **abc.company.com**, where **abc** is the host name. When you do not specify this optional parameter, the default is **localhost.**

To get Help on the Export utility parameters, issue this command:

```
VpnInvExport -h
```

About Journaling and the Journal Files

VPN Solutions Center contains four Repository databases—one each for the VPN Inventory, Directory, Task, and Collection Repositories.

- *VPN Inventory Repository.* Contains the provisioning service model and is the center of the service provisioning operations.
- *Directory Repository.* Stores information entered by the operator from the VPN Console GUI. For example, the Directory Repository stores information such as device names, IP addresses, passwords, and so on.
- *Task Repository.* Stores persistent task definitions, as well as task runtime information. For example, a task is created in the Task Repository when a service request is launched.
- *Collection Repository.* Keeps the information collected from the network, such as router configurations.

Each of these Repository databases has a corresponding journal file. For instance, messages for the VPN Inventory Repository are recorded in the *vi.jnl* file and messages for the Directory Repository are recorded in the *dir.jnl* file (see Table 14-1).

Table 14-1 Journal Files and Corresponding Repository

Journal Names	Repository Events
vi.jnl	VPN Inventory Repository events
dir.jnl	Directory Repository events
task.jnl	Task Repository events
col.jnl	Collection Repository events

Each of the journal files hold data corresponding to database actions such as create, delete, and update operations. The journal files are text files and you can view the contents through any text editor. Whenever a create, update, or delete event occurs on any of these Repositories, VPN Solutions Center sends a Tibco event with the event data as payload.

The journaling tool is a subscriber with Tibco for these messages, and whenever a message arrives, it records them into the appropriate journal file. Messages are delineated by a '<Message subject...>' and </Message>. The subject indicates the Repository for which the message is intended and also contains the type of action taken and the particular Repository object on which the action was taken.

As soon as a Repository backup is completed, VPN Solutions Center creates a new set of journal files.

Journaling starts automatically when the Watch Dog program is started. It stops when the Watch Dog program is terminated. Thus, you can force new journal files to be created by stopping and then restarting the Watch Dog.

To start journaling manually, issue this command:

```
wdclient start journal
```

To stop journaling manually, issue this command:

```
wdclient stop journal
```

For instructions on how to use a third party program to backup the Repository and journal files, see the "Using the Database Backup Utility" section on page 14-7.

Also refer to the "Using a Third-Party Application to Restore a Repository and Journal Files" section on page 14-10.

Specifying the Duration Between Journal Backups

By default, VPNSC copies the journal files into a subdirectory within the journal directory every seven days. You can configure how often VPNSC copies the journal files into a subdirectory in the *csm.properties* file, using one or both of the following properties:

- *netsys.watchdog.server.journal.archiving*

  The number specified in the *journal.archiving* property indicates the number of *days* between occurrences when VPNSC creates an archive directory for journal files. Only this archiving process creates a new directory and places *all the journal files*—both current and "aged" files—into the new journal subdirectory. A new (and empty) set of journal files are produced and placed in the current journal directory. The default setting for journal archiving is once every **seven** days.

- *netsys.watchdog.server.journal.aging*

  The number specified in the *journal.aging* property allows you to determine the number of *hours* between journal backups. However, note that this property goes into effect only when the specified number of hours has elapsed *and* the journal file exceeds 1.2 MB in size.

Note Both of these properties can be in effect at the same time. That is, you can set both the *journal.aging* parameter and the *journal.archiving* parameter, thus combining short-term and long-term journal backups.

Time Stamps and Journal Files

The journal subdirectory name includes a timestamp suffix. For example, the following journal file:

<*RepositoryPath*>/journal/journal.15_08_2002_10_30_00

contains the journal saved on August 15, 2002 at 10:30 A.M.

When the size of the journal files exceeds 1.2 Mb, VPN Solutions Center renames the current journal files (by appending a time-stamp string to the name), then VPNSC starts new journal files.

When you make a backup using the VPN Solutions Center VPN Console, the backup filename also has a time-stamp string suffix.

Thus, the time-stamps correlate backups and journal files. For example, if the journal backup filename is Repository.08_05_2002_12_15_00 (indicating a backup made on August 5, 2002 at 12:15:00), then the journal files contain data on events that occurred *after* the backup made on that date. For example, dir.jnl.08_05_2002_12_25_00 would be a valid journal file to use with the backup.

CHAPTER 15

VPNSC: MPLS Solution Troubleshooting Guide

This chapter describes how to recognize and troubleshoot problems you might encounter when deploying VPN Solutions Center: MPLS Solution VPN services.

General Topics

1. **Question:** What are the service deployment states? What do they mean?

   Answer: Table 15-1 describes the VPN service request deployment states.

Table 15-1 Summary of VPN Solutions Center Service Request States

Service Request Type	Description
Broken	While the router is correctly configured, the service is unavailable (due to a broken cable or Layer 2 problem, for example). A service request moves to Broken if the Auditor finds the routing and forwarding tables for this service, but they do not match the service intent.
Closed	A service request moves to Closed if the service request should no longer be used during the provisioning or auditing process. A service request moves to the Closed state only upon a successful audit of a remove request. VPNSC: MPLS Solution does not remove a service request from the database to allow for extended auditing. Only a specific administrator action results in service requests being removed.
Deployed	A service request moves to Deployed if the configlet commands have been verified as found in the router configuration file. Deployed indicates that the configuration file on the router matches the information specified in the VPNSC service request.
Failed Audit	The Failed Audit state indicates that VPNSC downloaded the configlet to the router successfully, but the service request did not pass the audit. Therefore, the service did not move to either the Functional or Deployed state. The Failed Audit state is initiated from the Pending state. Once a service request is deployed successfully, it cannot reenter the Failed Audit state (except when the service request is redeployed).

Table 15-1 Summary of VPN Solutions Center Service Request States (continued)

Service Request Type	Description
Failed Deploy	After provisioning occurred, the service request failed to download the configlets to the router. A service request moves to Failed Deploy if an error was detected during the deployment process by the Telnet Gateway Server (TGS). If TGS is not being used to download configlets, and VPNSC is simply exporting configlets to a directory, there is no way to distinguish between a service request in the Failed Deploy and Pending states. There are two causes for Failed Deploy status: • TGS reports to the Provisioning Driver that the download failed (lost connection, faulty password, etc.). • The object could not establish configuration-level verification of intent. If the configlets are exported to a directory, the service request cannot move into a Failed Deploy state.
Functional	A service request moves to Functional when the Auditor finds the VPN routing and forwarding tables (VRF) for this service and they match with the service intent. This state requires that both the configuration file audit and the routing audit are successful.
Invalid	Indicates that the service request information is incorrect in some way. A service request moves to Invalid if the request was either internally inconsistent or not consistent with the rest of the existing network/router configurations (for example, no more interfaces were available on the router). The Provisioning Driver cannot generate configlets to service this request.
Lost	A service request moves to Lost when the Auditor cannot find a configuration-level verification of intent in the router configuration files. The service request was deployed, but now some or all router configuration information is missing. A service request can move to the Lost state *only* when the service request had been Deployed or Functional.
Pending	A service request moves to Pending when the Provisioning Driver determines that the request looks consistent and was able to generate the required configlets for this request. Pending indicates that the service request has generated the configlets and the configlets are successfully downloaded to the routers. The Auditor regards pending service requests as new requests and begins the audit. If the service has been freshly provisioned and not yet audited, it is not an error (pending audit). However, if an audit is done and the service is still pending, it is in an error state.
Requested	If the service is newly entered and not yet deployed, it is not an error. However, if a Deploy is done and it remains Requested, the service is in an error state.

Table 15-2 on page 15-3 and Table 15-3 on page 15-4 show the state transition paths for VPN Solutions Center service requests. The beginning state of a service request is listed in the first column; the states that service requests transition to are displayed in the heading row.

For example, to use Table 15-2 to trace the state of a Pending service request to Functional, find "**Pending**" in the first column and move to your right until you find "**Functional**" in the heading. You can see that for a service request to move from Pending to Functional, a successful routing audit must take place.

Table 15-2 shows the service request transitions from *Requested* to *Lost*.

Table 15-2 State Transition Paths for VPNSC Service Requests (Part 1)

Service Request States	Requested	Pending	Failed Audit	Deployed	Functional	Lost
Requested	No transition to Requested	Successful service request deployment	No transition to Failed Audit	No transition to Deployed	No transition to Functional	No transition to Lost
Pending	No transition to Requested	—Successful service request deployment —Audit with error	Audit is not successful	Audit is successful	Routing audit is successful	No transition to Lost
Failed Audit	No transition to Requested	Successful service request redeployment	No transition to Failed Audit	Audit is successful	Routing audit is successful	No transition to Lost
Deployed	No transition to Requested	Successful service request redeployment	No transition to Failed Audit	Audit is successful	Routing audit is successful	Audit found error
Functional	No transition to Requested	Successful service request redeployment	No transition to Failed Audit	No transition to Deployed	Routing audit is successful	Audit found error
Lost	No transition to Requested	Successful service request redeployment	No transition to Failed Audit	Audit is successful	Routing audit is successful	Audit found error
Broken	No transition to Requested	Successful service request redeployment	No transition to Failed Audit	No transition to Deployed	Routing audit is successful	Audit found error
Invalid	No transition to Requested	Successful service request redeployment	Redeployment caused service request error	No transition to Deployed	No transition to Functional	No transition to Lost
Failed Deploy	No transition to Requested	Successful service request redeployment	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Deployed	No transition to Functional	No transition to Lost
Closed	No transition to Requested	No transition to Pending	No transition to Failed Audit	No transition to Deployed	No transition to Functional	No transition to Lost

Table 15-3 shows the service request transitions from *Broken* to *Closed*.

Table 15-3 State Transition Paths for VPNSC Service Requests (Part 2)

Service Request States	Broken	Invalid	Failed Deploy	Closed
Requested	No transition to Broken	Deploy Service Request error	Deployment failed	No transition to Closed
Pending	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	Removal of the service request is successful
Failed Audit	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Deployed	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Functional	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Lost	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Broken	Route audit is not successful. Configlet is correct.	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Invalid	No transition to Broken	Redeployment caused service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Failed Deploy	No transition to Broken	Redeploy service request error	Redeployment service request failed. Configlet cannot be downloaded.	No transition to Closed
Closed	No transition to Broken	No transition to Invalid	No transition to Failed Deploy	No transition to Closed

2. **Question:** Which of the error states are due to provisioning and which are due to auditing?

   Answer: After provisioning, Requested, Invalid, and Failed Deploy are due to error conditions in discovered during the provisioning process. After auditing, Pending, Failed Audit, Lost, and Broken are due to error conditions discovered during the auditing process.

3. **Question:** I executed an *Add VPN Service to CE* followed by a *Deploy Service Requests*, then I selected *Generate Audit Reports*. However, the All VPN Service Requests Report indicates that the service request is not in either a Deployed or Functional state. Where do I look?

   Answer: If the service request is in the Requested, Invalid, or the Failed Deploy state, refer to the "Provisioning Problems" section on page 15-5. However, if the service request is stuck in Failed Audit, refer to the "Auditing Problems" section on page 15-9.

Provisioning Problems

The VPNSC: MPLS Solution provisioning system has the following functions:

- Function 1: Collect the PE router configuration files (PE-upload)
- Function 2: Collect the CE router configuration files (CE-upload)
- Function 3: Provisioning
- Function 4: Write the changed configuration information to the PE (PE-DownLoad)
- Function 5: Write the changed configuration information to the CE (CE-DownLoad)

Functions 1, 2, 4, and 5 are executed by a server called the Telnet Gateway Server (TGS).

Errors in functions 1 or 2 lead to functions 3, 4, and 5 being skipped.

The provisioning engine has a model of the router. This router model is modified as necessary to introduce attributes to support the service request.

1. **Question:** What is the flow of the provisioning operation?

   Answer:

   The program that runs all these functions is called *Provisioning Driver*. The Provisioning Driver calls the TGS server to perform functions 1 and 2 (collecting PE and CE configuration files). After the VPN Solutions Center uploads the fresh configuration files from the router, it provisions the service request.

   After a successful operation to update the configlet with the necessary changes, the Provisioning Driver calls the Telnet Gateway Server to download the new configlets to the routers (that is, functions 4 and 5).

   If functions 1 or 2 (collecting PE or CE configuration files) fail, the other functions are skipped, and the service request stays in the Requested state.

   If function 3 (Provisioning) fails, the service request becomes Invalid.

   If function 3 succeeds, but functions 4 or 5 (writing the changed configuration information) fail, the service request moves to the Failed Deploy state.

2. **Question:** Where can I see how the provisioning functions performed in my audit?

   Answer: The first place to look at is in the Task Logs:

   a. From the VPN Console, choose **Tools > Task Logs**.

      The browser opens and displays the VPN Solutions Center Task Logs, as shown in Figure 15-1.

Figure 15-1 VPNSC: MPLS Solution Task Logs Browser

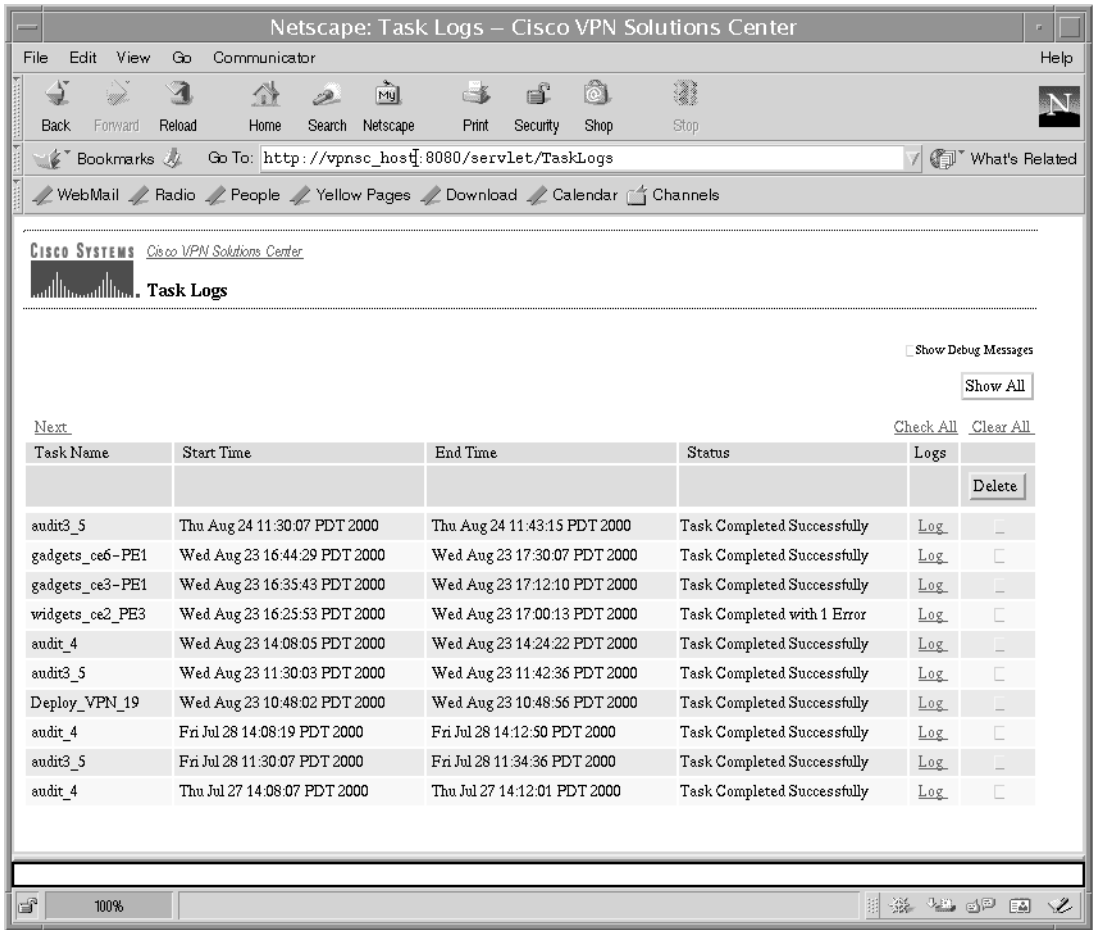

b. Choose the task that was run for this deployment.

The task name is the name you assigned. The tasks are listed in reverse chronological order (with the latest one first).

c. Click the **Log** link (in the rightmost column).

Summary information appears in the left pane (see Figure 15-2).

Figure 15-2 Task Log Summary Information and Action Report

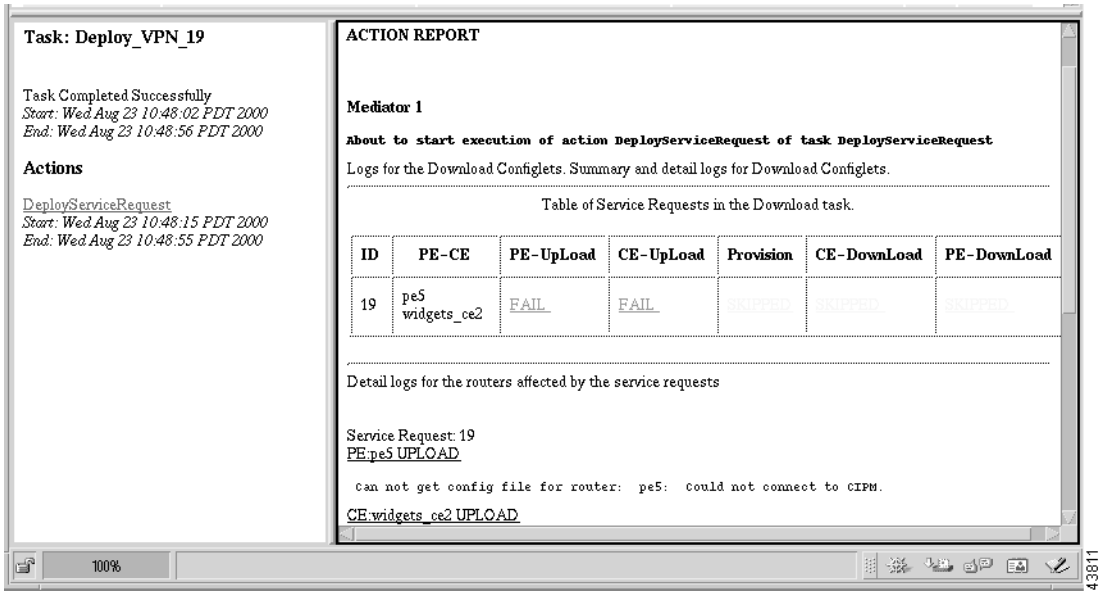

d. To see the Action Report, click the link under the **Actions** heading.

The Action Report appears, as shown in the right pane of Figure 15-2.

3. **Question:** My service request is stuck in the Requested state. Where should I go to look for errors?

Answer: In the Task Log Summary Table, look at the PE-UpLoad and CE-UpLoad information for that service request. One or both of them should say "Fail." This makes the rest of them "skipped." Thus, the service request remains in the Requested state.

4. **Question:** I do not see the task in the Task Log. What happened?

The typical reasons for this are as follows:

- The Task Scheduler has crashed and it is disabled.
- The Task Scheduler is malfunctioning
- The Task Scheduler is "disabled-dependent."

In each of these cases, the scheduled job does not run and no task log is produced.

a. Use the Watch Dog (wdgui) to check the state of the Task Scheduler.

b. If the Task Scheduler is disabled, issue the following command:

```
wdclient start scheduler
```

c. If the Task Scheduler is "disabled-dependent," some of the dependent servers (such as the lock_manager and the EventServiceServer monitor poller) did not start. Start them by issuing the command:

```
wdclient start server_name
```

If the dependent servers are started, then Task Scheduler starts automatically.

5. **Question:** I see in the wdlog that the Task Scheduler is up and running. Yet my scheduled task is not running. What's wrong?

Answer: When the Task Scheduler starts, it automatically picks up pending requests. However, there are cases in which it may not. Watch the wdgui messages for the Task Scheduler for a few minutes. If the Task Scheduler does not start after a few minutes, schedule the task again.

If the Task Scheduler is malfunctioning, do the following:

a. Check the wdgui information for the Task Scheduler. Does this show any abnormality?

b. If not, go to the Task dialog box and issue a refresh.

c. Does the task still show up as active?

d. Note the time it is supposed to start.

   Has the time already passed (as shown by the system clock where the Task Scheduler is running)?

e. In these cases, delete the earlier task and reschedule. After a few minutes, check the wdgui Task Scheduler information for any activity or messages.

f. If there are any abnormal error messages in the Task Scheduler, read the messages and take the recommended corrective action.

g. However, if there are no messages, or the messages are not understandable, delete the old task from the task dialog box and restart the Task Scheduler by the command: wdclient restart Task Scheduler.

h. When the Task Scheduler starts, reschedule the command.

i. If the problem persists, stop the Watch Dog.

j. Wait for two minutes, then restart the Watch Dog with the **startwd** command.

Do not forget to delete the active task before restarting Watch Dog. Wait until the servers stabilize before you reschedule the task.

6. **Question:** My service request is in the Invalid state. How do I correct the problem?

   Answer: When a service request is in moved to Invalid, it is because the request cannot be serviced. Either something asked for in the service request cannot be serviced or there was an internal error.

   a. Bring up the Task Log Summary table at the Task Logs browser page by choose **Tools > Task Logs.**

      You can also access this information from your browser by entering this URL:

      http://*vpnsc_hos*t:8080/servlet/TaskLogs

      Refer to the Questions 2 and 3 in this section for additional information.

   b. In the Task Log Summary table, click **Fail** (under **Provisioning**).

      This takes you to the description of what error condition was noticed during the provisioning. In some cases, this may give a more detailed error message.

      Another place to find internal error information is from the VPN Console.

   a. From the VPN Console, choose **Provisioning > List all Service Requests**.

      The All VPN Service Requests Report appears. Notice that the service request is Invalid.

   b. Click **Request Details**.

   c. Observe the Last State Change Comment, which displays the reason the service request is in the Invalid state. In some cases, this may include a summarized error message.

   d. Read the error message and note the incorrect value(s) entered.

   e. From the All VPN Service Request Report, click **Provisioning** and choose **Modify VPN Service** (see the Modifying an Existing Service Request, page 6-25.

   f. Correct the errors in the modified request and redeploy the service request.

7. **Question:** My service request is in the Failed Deploy state. How should I address the problem?

   Answer: Failed Deploy indicates that there is an error while downloading the changed configlets back to the router (refer to Table 15-1 on page 15-1 for more information).

   The procedure described for an Invalid request (see Question 8) pertains here as well. However, in the Task Logs Summary table, look at the PE-Download or CE-Download information. That is where the error is. Take the link from there.

   The cause of the problem could be either one of two possibilities: 1) while the configuration changes were downloading, the link to the router(s) was dropped or 2) a configuration command that was sent to a router invoked a warning or error message.

   a. First, read the error message and try to understand it. Was it a communication error?
   b. If so, Telnet to the router from the VPN Solutions Center workstation. Get the communication to work first.
   c. Redeploy the service request by choosing **Provisioning > List All Service Requests**.
   d. From the All VPN Service Requests Report, click **Provisioning**, and choose **Deploy VPN Services**.

      If the error was due to a command that generated a warning or error, the router may have rejected the command because the version of the Cisco IOS on the router does not support it. If the problem persists, contact the Cisco Technical Assistance Center and provide a) the Command Rejected information, and b) the **show version** output from the router.

Auditing Problems

Both the *Audit new service requests* and *Audit existing service requests* audit types provide two tests: first, a **configuration test** and then a **routing test** (if specified by checking *Use VPN routing information during audits* in the task window). If the configuration test fails, the routing test is not performed.

- The **configuration test** checks to see whether the router configuration files match the information specified in the VPNSC service request. VPNSC: MPLS Solution takes the configuration files and builds a "model router" from the configuration information. The configuration test builds a software model of the router and audits it. The configuration test cannot determine whether the service is actually running.
- The **routing test** ("audit routing") actually performs three tests: 1) It checks for the presence of routes toward the CE; 2) Checks for the presence of routes away from the CE; and 3) If the CE is managed using the management VPN technique, the routing test checks for the presence of a route to the Management CE (MCE) in the management VPN. This audit is based on the routing information found on the PE, and the routing test audit details are placed under the PE.

  All these tests are done for the VRF in which the service request belongs. To test the presence of routes toward the CE, VPNSC looks for a route toward the CE's provider-facing IP address.

  The test for routes away from the CE looks for routes toward the "other side of the VPN." This test checks the remote connectivity status. If the service request is in a management VPN, the final test checks for a route to the management CE (MCE).

If the service request fails an Audit New Service Request, it moves to the Failed Audit state. If the service request passes this audit and audit routing is enabled (by checking *Use VPN routing information during audits* in the task window), VPNSC: MPLS Solution tests whether there is adequate VRF routing information for audit routing. If the routing audit does not have the VRF routing information or the *Use VPN routing information during audits* option is not checked, the service request remains in the state it was in during the configuration audit.

If the routing test passes, the service request moves to Functional. If the routing test fails, the service request moves to Broken.

The starting state for an Audit Existing Service Request can be Pending, Deployed, Failed Audit, Lost, Functional, or Broken (refer to Table 15-1 on page 15-1). In each case, the deployed and routing tests are performed in order.

If the configuration test fails, the service request moves to (or remains in) the Lost state. If the routing test fails, the service request moves to (or remains in) the Broken state. If the configuration test succeeds, the service request moves to (or remains in) the Deployed state; if the routing test succeeds, the service request moves to (or remains in) the Functional state.

1. **Question:** My service request is stuck in Failed Audit or Lost. How do I find out what went wrong?

   Answer:

   a. For a service request stuck in Failed Audit or Lost, generate a service request audit report by choosing **Provisioning > List All Service Requests**.

   b. From the All VPN Service Requests report, select the service request of interest, then click the **Request Details** button.

   c. From this window, click the **Audit Details** button.

      The Audit Details report gives you the audit trace for both the PE and CE; this report is organized into two parts—one for the PE and other for the CE. Any errors found are highlighted in yellow.

   d. Read the error information.

      The problem is usually missing or wrong configuration information on the router. Has this configuration command been generated in the configlet?

   e. To see the configlet, first click **Back**, then click the **Configlets** button.

      VPN Solutions Center displays the current configlet.

   f. See if the command is present, and if present, whether it is correct.

   g. If you feel that the command that was generated is incorrect and should be changed, contact Cisco Technical Assistance Center for help.

   h. If the command is correct, check the configuration command line in the configlet (that is, the line that is apparently missing or incorrect as discovered by the audit) and see if the command line is present in the router.

   i. If the command line is removed or absent or incorrectly changed in the router, redeploy the service request to get the corrections sent to the router.

   j. However, if the configlet seems correct and the configuration commands are present in the router, call Cisco Technical Assistance Center for assistance.

2. **Question:** How do I move my service request to the Functional state?

   Answer: To move a service request to Functional, run an audit. A successful audit moves the service request to Deployed.

   a. In the hierarchy pane, open the **VPNs** folder.

   b. Select the VPN of interest and **right-click**.

   c. From the menu, choose **Audit Service Requests**.

   d. From the Audit VPN dialog box, choose the following:

   – **All Service request(s)**

      This audits all the service requests for the current VPN only.

- To the prompt, "Do you want to perform a just-in-time (jit) collection...," choose **Yes**.
- To the prompt, "Do you want audit routing?" choose **Yes**.
- Click **OK**.

e. Select the **Schedule** tab and complete the fields in the Schedule dialog box to schedule the audit, then click **OK**.

3. **Question:** I checked the "Use VPN routing info during audits" option, but I receive the message, "No VPN routing information found."

   Answer: You must collect VPN routing info to use it. Schedule data collection by choosing **Monitoring > Collect VPN Routing Information**.

4. **Question:** My service request is the Broken state. How do I address the problem?

   Answer: While the router is correctly configured, the service is unavailable (due to a broken cable or Layer 2 problem, for example). A service request moves to Broken if the Auditor finds the routing and forwarding tables for the router, but none of the routing information in the service request is present in the router's VRF routing table.

   The routing test ("audit routing") actually performs three tests: 1) It checks for the presence of routes toward the CE; 2) Checks for the presence of routes away from the CE; and 3) If the CE is managed using the management VPN technique, the routing test checks for the presence of a route to the Management CE (MCE) in the management VPN.

   If any of the routing tests fail, the service request is set to Broken. Since this audit is based on the routing information found on the PE, the routing test audit details are placed under the PE. Because the routing test is a Layer 3 operational test, VPNSC cannot know why things have gone wrong at that layer—it could be a broken cable, lost Layer 2 connectivity, and so on.

5. **Question:** I have a VPN with two CEs. One PE-CE Layer 2 connection is down and in a Broken or Lost state. Why does the other PE-CE pair's service request also move to Broken?

   Answer: The routing test checks whether the service request provides a routing level connection of the site to the VPN. For a VPN to exist, there must be at least two sites. Hence, in a VPN with two sites, if connectivity of one of the sites go down, the VPN no longer exists. Therefore, the connectivity of the other site to the VPN also fails—there are no routes across the provider's core network.

6. **Question:** I just created a VPN and added my first site to it. Why does the service request move to Broken?

   Answer: Until you add at least two sites to the VPN, the VPN is not complete. The first site cannot participate in a VPN because the VPN does not yet exist. Thus, the service request moves to Broken. The service request moves to Functional when the second site is connected to the VPN, assuming core network connectivity exists and is functional.

7. **Question:** I see an extra row in the Audit Reports: *Repository Erro*r. What does this mean?

   Answer: A read/write operation to the Repository failed. Report this error and send your current Repository to Cisco Technical Assistance Center.

Cisco VPN Solutions Center Configuration File Examples

This chapter provides several examples of configuration files generated by VPN Solutions Center: MPLS Solution, Release 2.2. The IP addresses and network device names included in these examples are generic and are not intended to be used in your network.

Tip When using these configuration file examples in live networks, be sure to substitute appropriate IP addresses for the sample addresses used in these examples.

The following configuration file examples are included in this appendix:

CEs Configured as Hubs in the VPN

This configuration file provides an example of CEs configured as hubs in the VPN. In this example, a unique route distinguisher (RD) value is provisioned for each VRF.

```
!!
!! Topology:
!!
!! CE1---PE==PE1---CE2
!
!! --------------------------------
!! Provider Edge router PE is a member of the Blue VPN without
!! Management VPN connectivity.
!! CE1 is provisioned as a hub in the Blue VPN.
!
! Hostname: PE
!
! Version 12.0
!
!! Provisioned routing forwarding instance for Blue VPN—vrf V9:blue
!! Route target 200:5 is used for hub-to-hub routing connectivity.
!! Route-target 200:6 is used for spoke routing connectivity.
!
ip vrf V6:blue
rd 200:6
route-target import 200:5
route-target import 200:6
route-target export 200:5
!
!! The subinterface on the PE faces the CE. The address is from the VPNSC
!! IP address Pool.
!
interface Serial2/3.333 point-to-point
description Serial2/3.333 fr dlci=333 : Provisioned by VPNSC: Service Request Id# = 14
ip vrf forwarding V6:blue
ip address 209.165.201.17 255.255.255.252
frame-relay interface-dlci 333
no shutdown
!
!! The routing protocol for the PE-to-CE link is RIP.
!! Definition for a RIP routing instance for VRF Blue.
!! Routes from the IBGP core that are associated with route-targets 200:5 or 200:6
!! are redistributed into RIP.
!
router rip
address-family ipv4 vrf V6:blue
redistribute bgp 200 metric transparent
network 209.165.201.0
exit-address-family
no auto-summary
version  2
!
!! Definition for the core-facing IBGP routing protocol routing instance for VRF Blue
!! VRF blue RIP routes are redistributed into the IBGP core.
```

```
!! Exported RIP routes are associated with route target 200:5.
!
router bgp 200
address-family ipv4 vrf V6:blue
redistribute rip
exit-address-family
!
!! --------------------------------
!! Customer Edge router CE1 is provisioned as a hub in the Blue VPN.
!
! Hostname: CE1
!
! Version 12.0
!
interface Serial0
encapsulation frame-relay
!
!! The subinterface on the CE is facing the PE. The IP address is from the VPNSC Pool.
!
interface Serial0.333 point-to-point
description Serial0.333 fr dlci=333 : Provisioned By VPNSC: Service Request Id# = 14
ip address 209.165.201.21 255.255.255.252
frame-relay interface-dlci 333
no shutdown
!
!! The routing protocol for the PE-to-CE1 link is RIP.
!
router rip
network 209.165.201.0
no auto-summary
version  2
!
!! --------------------------------
!! Provider Edge router PE1 is a member of the Blue VPN without
!! Management VPN connectivity.
!
! Hostname: PE1
!
! Version 12.0
!
!! Provisioned routing forwarding instance for Blue VPN—vrf V9:blue
!! Route target 200:5 is used for hub-to-hub routing connectivity.
!! Route-target 200:6 is used for spoke routing connectivity.
!
ip vrf V9:blue
rd 200:9
route-target import 200:5
route-target import 200:6
route-target export 200:5
!
!! The subinterface on the PE is facing the CE. The IP address is from the VPNSC Pool.
!
interface Serial2/0.334 point-to-point
```

```
description Serial2/0.334 fr dlci=334 : Provisioned by VPNSC: Service Request Id# = 15
ip vrf forwarding V9:blue
ip address 209.165.201.21 255.255.255.252
frame-relay interface-dlci 334
no shutdown
!
!! The routing protocol for the PE-to-CE link is RIP.
!! Definition for a RIP routing instance for VRF Blue.
!! Routes associated with route-targets from the BGP core that are associated
!! with route-targets 200:5 or 200:6 are redistributed into RIP.
!
router rip
address-family ipv4 vrf V9:blue
redistribute bgp 200 metric transparent
network 209.165.201.0
exit-address-family
no auto-summary
version  2
!
!! Definition for the core-facing IBGP routing protocol routing instance for VRF Blue.
!! VRF Blue RIP routes are redistributed into the IBGP core.
!! Exported RIP routes are associated with route target 200:5.
!
router bgp 200
address-family ipv4 vrf V9:blue
redistribute rip
exit-address-family
!
!! --------------------------------
!! Customer Edge router CE2 is provisioned as a hub in the Blue VPN.
!
! Hostname: CE2
!
! Version 12.0
!
!! The subinterface on the CE is facing the PE. The IP address is from the VPNSC Pool.
!
interface Serial0.334 point-to-point
description Serial0.334 fr dlci=334 : Provisioned by VPNSC: Service Request Id# = 15
ip address 209.165.201.22 255.255.255.252
frame-relay interface-dlci 334
no shutdown
!
!! The routing protocol for the PE1-to-CE2 link is RIP.
!
router rip
network 209.165.201.0
no auto-summary
version  2
```

Sample Hub-and-Spoke Topology

This configuration file shows a sample hub-and-spoke topology with three CEs. CE1 is a hub in the VPN; CE2 and CE3 are spokes in the same VPN. An *-s* appended to the VRF name indicates that the VRF is associated with spoke connectivity. The VRF naming and the RD/RT allocation would not change if one or more PEs are employed.

```
!! Topology:
!!
!! CE1---PE---CE2
!!        |
!! CE3----
!!
!! This configuration would not change if the CEs were attached to the same
!! or different PEs.
!
!! --------------------------------
!! Provider Edge router: the PE is a member of the Blue VPN without
!! Management VPN connectivity.
!! CE1 is provisioned as a hub; CE2 and CE3 are provisioned as spokes in the Blue VPN.
!
!Hostname: PE
!
! Version 12.0
!
!! Provisioned routing forwarding instance for Blue VPN—vrf V6:blue
!! for CE1 hub connectivity.
!! Route target 200:5 is used for hub-to-hub routing connectivity.
!! Route-target 200:6 is used for spoke routing connectivity.
!
ip vrf V6:blue
rd 200:6
route-target import 200:5
route-target import 200:6
route-target export 200:5
!
!! Provisioned routing forwarding instance for Blue VPN—vrf V7:blue-s
!! for CE2 spoke connectivity.
!! The "-s" appended to the VRF name indicates that this VRF is associated with
!! spoke connectivity.
!! Route target 200:5 is used for hub routing connectivity.
!! Route-target 200:6 is used for spoke routing connectivity.
!
ip vrf V7:blue-s
rd 200:7
route-target import 200:5
route-target export 200:6
!
!! Provisioned routing forwarding instance for Blue VPN—vrf V8:blue-s
!! for CE3 spoke connectivity.
!! The "-s" indicates that this VRF is associated with spoke connectivity.
!! Route target 200:5 is used for hub routing connectivity.
!! Route-target 200:6 is used for spoke routing connectivity.
!
```

```
ip vrf V8:blue-s
rd 200:8
route-target import 200:5
route-target export 200:6
!
!! The subinterface on the PE faces CE1; the address is from the VPNSC IP address Pool.
!
interface Serial2/0.122 point-to-point
description Serial2/0.122 fr dlci=122 : Provisioned by VPNSC: Service Request Id# = 11
ip vrf forwarding V6:blue
ip address 209.165.201.1 255.255.255.252
frame-relay interface-dlci 122
no shutdown
!
!! The subinterface on the PE faces CE2; the address is from the VPNSC IP address pool.
!
interface Serial2/1.123 point-to-point
description Serial2/1.123 fr dlci=123 : Provisioned by VPNSC: Service Request Id# = 12
ip vrf forwarding V7:blue-s
ip address 209.165.201.5 255.255.255.252
frame-relay interface-dlci 123
no shutdown
!
!! The subinterface on the PE faces CE3; the address is from the VPNSC IP address pool.
!
interface Serial2/2.124 point-to-point
description Serial2/2.124 fr dlci=124 : Provisioned by VPNSC: Service Request Id# = 13
ip vrf forwarding V8:blue-s
ip address 209.165.201.9 255.255.255.252
frame-relay interface-dlci 124
no shutdown
!
!! The routing protocol is RIP on the PE-CE link.
!!
router rip
!
!! Definition for RIP routing instance for VPN Blue.
!! Routes from the IBGP core that are associated with route-targets 200:5 or 200:6
!! are redistributed into RIP.
!! Provides hub VRF definition.
!
address-family ipv4 vrf V6:blue
redistribute bgp 200 metric transparent
network 209.165.201.0
exit-address-family
!
!! Definition for RIP routing instance for VRF Blue-s (spoke)
!! Routes from the IBGP core that are associated with route-targets 200:5
!! are redistributed into RIP.
```

```
!!
address-family ipv4 vrf V7:blue-s
redistribute bgp 200 metric transparent
network 209.165.201.0
exit-address-family
!
!! Definition for RIP routing instance for VRF Blue-s (spoke)
!! Routes from the IBGP core that are associated with route-targets 200:5
!! are redistributed into RIP.
!!
address-family ipv4 vrf V8:blue-s
redistribute bgp 200 metric transparent
network 209.165.201.0
exit-address-family
!
no auto-summary
version  2
!
!! Definition for the core-facing IBGP routing protocol routing instance for VRF Blue.
!! VRF Blue RIP routes are redistributed into the IBGP core.
!
router bgp 200
!
!! Exported RIP routes are associated with route target 200:5.
!
address-family ipv4 vrf V6:blue
redistribute rip
exit-address-family
!
!! Exported RIP routes are associated with route target 200:6.
!
address-family ipv4 vrf V7:blue-s
redistribute rip
exit-address-family
!
!! Exported RIP routes are associated with route target 200:6.
!
address-family ipv4 vrf V8:blue-s
redistribute rip
exit-address-family
!
!! --------------------------------
!! Customer Edge router CE1 is provisioned as a hub in the Blue VPN.
!
! Hostname: CE1
!
! Version 12.0
!
!! The CE subinterface faces the PE; the address is from the VPNSC IP address pool.
!
interface Serial0
encapsulation frame-relay
!
```

```
interface Serial0.122 point-to-point
description Serial0.122 fr dlci=122 : Provisioned by VPNSC: Service Request Id# = 11
ip address 209.165.201.2 255.255.255.252
frame-relay interface-dlci 122
no shutdown
!
!! The routing protocol for the PE-to-CE1 link is RIP.
!! Provides optional redistribution of the customer routing protocol EIGRP into the VPN.
!
router rip
network 209.165.201.0
redistribute eigrp 11 metric 1
no auto-summary
version  2
!
router eigrp 11
redistribute rip metric 1544 2000 255 1 1500
!
!! --------------------------------
!! Customer Edge router CE2 is provisioned as a spoke in the Blue VPN.
!
! Hostname: CE2
!
! Version 12.0
!
!! The CE subinterface faces the PE; the address is from the VPNSC IP address pool.
!
interface Serial0.123 point-to-point
description Serial0.123 fr dlci=123 : Provisioned by VPNSC: Service Request Id# = 12
ip address 209.165.201.6 255.255.255.252
frame-relay interface-dlci 123
no shutdown
!
!! The routing protocol for the PE-to-CE2 link is RIP.
!
router rip
network 209.165.201.0
no auto-summary
version  2
!
!! --------------------------------
!! Customer Edge router CE3 is provisioned as a spoke in the Blue VPN.
!!
! Hostname: CE3
!
! Version 12.0
!
!! The subinterface on the CE is facing the PE, the IP address is from the VPNSC Pool.
!
interface Serial0.124 point-to-point
description Serial0.124 fr dlci=124 : Provisioned By VPNSC: Service Request Id# = 13
ip address 209.165.201.6 255.255.255.224
frame-relay interface-dlci 124
```

```
no shutdown
!
!! The routing protocol for the PE-to-CE3 link is RIP.
!
router rip
network 209.165.201.0
no auto-summary
version  2
```

Management VPN Configuration Example

This configuration file provides an example of provisioning a Management VPN, as well as provisioning the Management CE (MCE) and Management PE (MPE). For related information, see the "Management VPN Technique" section on page 8-6 and the "Implementing the Management VPN Technique" section on page 8-12.

!! Topology:

```
!!
!! CE1---PE==MPE---MCE
!
!! --------------------------------
!! Provider Edge router: PE
!! CE1 is provisioned as a hub in the Blue VPN and as a spoke in the Management VPN.
!
! Hostname: PE
!
! Version 12.0
!
!! Provisioned routing forwarding instance for Blue VPN—vrf V6:blue.
!! The route-target 200:5 is for customer-hub connectivity.
!! The route-target 200:6 is for customer-spoke connectivity.
!! The route-target 200:1 is to import a route from the MCE into the VRF.
!! The export map exports only the PE-to-CE link subnet from the blue VRF.
!! The export map exports the management route-target 200:2 and exports the
!! Blue VPN target 200:5.
!! The CE attached to the Blue VPN is a spoke in the Management VPN.
!
ip vrf V6:blue
rd 200:6
route-target import 200:5
route-target import 200:6
route-target import 200:1
route-target export 200:5
export map grey_mgmt_vpn_VpnsRus_V6:blue
!
!! The subinterface on the PE faces CE1. The IP address is from the VPNSC Pool.
!
interface Serial2/1.555 point-to-point
description Serial2/1.555 fr dlci=555 : Provisioned by VPNSC: Service Request Id# = 16
ip vrf forwarding V6:blue
ip address 209.165.202.129 255.255.255.252
```

```
frame-relay interface-dlci 555
no shutdown
!
!! The routing protocol for the PE-to-CE link is RIP.
!! Definition for a RIP routing instance for VRF Blue.
!! Routes from IBGP core that are associated with route-targets 200:5, or 200:6,
!! or 200:1 are redistributed into RIP.
!
router rip
address-family ipv4 vrf V6:blue
redistribute bgp 200 metric transparent
network 209.165.202.0
exit-address-family
!
no auto-summary
version  2
!
!! Definition for the core-facing IBGP routing protocol routing instance for VRF Blue.
!! VRF Blue RIP routes are redistributed into the IBGP core.
!! Exported RIP routes are associated with route target 200:5 and 200:2.
!
router bgp 200
address-family ipv4 vrf V6:blue
redistribute rip
exit-address-family
!
!! The route map is used by the export map in the Blue VRF for filtering
!! routes to the Management VPN.
!! The match matches the PE-to-CE subnet with access-list VPNSC_GREY_MGMT_ACL.
!! Route-targets for Management 200:2 and Blue VPN route-target 200:5 are exported.
!
route-map grey_mgmt_vpn_VpnsRus_V6:blue permit 10
match ip address VPNSC_GREY_MGMT_ACL
set extcommunity rt 200:2 200:5
!
ip access-list extended VPNSC_GREY_MGMT_ACL
permit 209.165.202.128 0.0.0.3 255.255.255.255
!
!! Customer Edge router CE1 is provisioned as a hub in the Blue VPN
!! and as a spoke in the Management VPN.
!
! Hostname: CE1
!
! Version 12.0
!
interface Serial0.510 point-to-point
description Serial0.510 fr dlci=510 : Provisioned by VPNSC: Service Request Id# = 14
ip address 209.165.209.29 255.255.255.252
frame-relay interface-dlci 110
no shutdown
!
interface Serial0.555 point-to-point
description Serial0.555 fr dlci=555 : Provisioned By VPNSC: Service Request Id# = 16
```

```
ip address 209.165.202.130 255.255.255.224
frame-relay interface-dlci 555
no shutdown
!
!! The routing protocol for the PE-to-CE1 link is RIP.
!
router rip
network 209.165.202.0
no auto-summary
version  2
!
!! Management Provider Edge router: MPE
!! The attached Management CE (MCE) is a hub in the Management VPN.
!
! Hostname: MPE
! Version 12.0
!
!! The Management VPN uses route-target 200:1 as a hub and route-target 200:2 as a spoke.
!
ip vrf grey_mgmt_vpn_VpnsRus
rd 200:1
route-target import 200:1
route-target import 200:2
route-target export 200:1
!
!! The subinterface on the MPE faces the MCE.
!
interface Serial1/3
ip vrf forwarding grey_mgmt_vpn_VpnsRus
ip address 209.165.201.30 255.255.255.252
!
!! The routing protocol for the MPE-to-MCE link is RIP.
!! (Cisco recommends that you use a dynamic routing protocol.)
!! Definition for RIP routing instance for the VRF Grey Management VPN.
!! Routes from IBGP core that are associated with route-targets 200:1
!! and 200:2 are redistributed into RIP.
!! The subnet from the PE to CE1 link is imported with route-target 200:2.
!
router rip
address-family ipv4 vrf grey_mgmt_vpn_VpnsRus
redistribute static metric 1
redistribute bgp 200 metric transparent
network 209.165.201.0
exit-address-family
!
!! Routes are exported into the BGP core from RIP; connected and static routes
!! use route-target 200:1.
!
router bgp 200
address-family ipv4 vrf grey_mgmt_vpn_VpnsRus
redistribute rip
redistribute static
redistribute connected
```

```
exit-address-family
!
!! Customer Edge router MCE is provisioned as a hub in the Mgmt VPN.
!
! Hostname: MCE
!
! Version 12.0
!
router rip
network 209.165.201.0
!
rtr responder
```

A CE Configured as a Member of an Multiple VPNs

An *extranet* is a VPN with CEs that are members of multiple VPNs. Extranet provisioning provides a way to create multiple VPN connectivity to a single VRF. You can add multiple CERCs to your VPN in any topology to form extranets. You can join an extranet in such a way that a CE can be a spoke in one VPN and a hub in another VPN.

This configuration includes three CEs—two CEs in different VPNs and one CE that is a member of an extranet. A VRF name appended with ***-etc*** indicates that the VRF is a member of an extranet.

```
!! Topology:
!
!! CE1---PE---CE2
!!        |
!! CE3----
!
!! CE1 is a hub in the Blue VPN.
!! CE2 is a hub in the Red VPN.
!! CE3 is a hub in both the Blue and Red VPNs (Extranet).
!!
!
!! --------------------------------
!! Provider Edge router: PE
!
! Hostname: PE
!
! Version 12.0
!
!! Provisioned routing forwarding instance for blue VPN—vrf V6:blue
!! for CE1 hub connectivity.
!! Route target 200:5 is used for hub-to-hub routing connectivity.
!! Route-target 200:6 is used for spoke routing connectivity.
!
ip vrf V6:blue
rd 200:6
route-target import 200:5
route-target import 200:6
route-target export 200:5
!
```

```
!! Provisioned routing forwarding instance for Red VPN—vrf V10:red
!! for CE2 hub connectivity.
!! Route target 200:3 is used for hub-to-hub routing connectivity.
!! Route-target 200:4 is used for spoke routing connectivity.
!
ip vrf V10:red
rd 200:10
route-target import 200:3
route-target import 200:4
route-target export 200:3
!
!! Provisioned routing forwarding instance for blue VPN—vrf V6:blue-etc
!! for CE3 hub connectivity.
!! Route target 200:5 is used for hub-to-hub routing connectivity in the Blue VPN
!! Route-target 200:6 is used for spoke routing connectivity in the Blue VPN
!! Route target 200:3 is used for hub-to-hub routing connectivity in the Red VPN
!! Route-target 200:4 is used for spoke routing connectivity in the Red VPN
!! The VRF name with "-etc" indicates that the VRF is a member of an extranet.
!
ip vrf V11:blue-etc
rd 200:11
route-target import 200:3
route-target import 200:4
route-target import 200:5
route-target import 200:6
route-target export 200:3
route-target export 200:5
!
!! The subinterface on the PE is facing CE1; the IP address is from the VPNSC Pool.
!
interface Serial2/0.343 point-to-point
description Serial2/0.343 fr dlci=343 : Provisioned by VPNSC: Service Request Id# = 17
ip vrf forwarding V6:blue
ip address 209.165.200.229 255.255.255.255
frame-relay interface-dlci 343
no shutdown
!
!! The subinterface on the PE is facing CE2; the IP address is from the VPNSC Pool.
!
interface Serial2/3.888 point-to-point
description Serial2/3.888 fr dlci=888 : Provisioned by VPNSC: Service Request Id# = 18
ip vrf forwarding V10:red
ip address 209.165.200.233 255.255.255.252
frame-relay interface-dlci 888
no shutdown
!
!! The subinterface on the PE is facing CE3; the IP address is from the VPNSC Pool.
!
interface Serial2/5.777 point-to-point
description Serial2/5.777 fr dlci=777 : Provisioned by VPNSC: Service Request Id# = 19
ip vrf forwarding V11:blue-etc
ip address 209.165.200.237 255.255.255.252
frame-relay interface-dlci 777
```

```
no shutdown
!
!! The routing protocol is RIP on the PE-to-CE link.
!
router rip
!
!! Definition for the RIP routing instance for the VPN Blue.
!! Routes from the IBGP core that are associated with route-targets 200:5 or 200:6
!! are redistributed into RIP.
!! Hub VRF definition.
!
address-family ipv4 vrf V6:blue
redistribute bgp 200 metric transparent
network 209.165.200.0
exit-address-family
!
!! Definition for RIP routing instance for the VPN Red.
!! Routes from the IBGP core that are associated with route-targets 200:3 or 200:4
!! are redistributed into RIP.
!! Provides hub VRF definition.
!
address-family ipv4 vrf V10:red
redistribute bgp 200 metric transparent
network 209.165.200.0
exit-address-family
!
!! Definition for RIP routing instance for the VRF in both the Red and Blue VPNs.
!! Routes from the IBGP core that are associated with route-targets 200:5, 200:6, 200:3,
!! or 200:4 are redistributed into RIP.
!! Provides hub VRF definition.
!
address-family ipv4 vrf V11:blue-etc
redistribute bgp 200 metric transparent
network 209.165.200.0
exit-address-family
!
router bgp 200
!
!! Definition for the core-facing IBGP routing protocol routing instance for VRF Blue.
!! VRF Blue RIP routes are redistributed into the IBGP core.
!
address-family ipv4 vrf V6:blue
redistribute rip
exit-address-family
!
!! Definition of the core-facing IBGP routing protocol routing instance for the VRF Blue.
!! VRF Red RIP routes are redistributed into the IBGP core.
!
address-family ipv4 vrf V10:red
redistribute rip
exit-address-family
!
!! Core-facing IBGP routing protocol routing instance for the extranet VRF
```

```
!! VRF Red RIP routes are redistributed into the IBGP core
!
address-family ipv4 vrf V11:blue-etc
redistribute rip
exit-address-family
!
!! ------------------------------
!! Customer Edge router CE1 is provisioned as a hub in the Blue VPN.
! Hostname: CE1
!
! Version 12.0
!
interface Serial0
encapsulation frame-relay
!
interface Serial0.343 point-to-point
description Serial0.343 fr dlci=343 : Provisioned by VPNSC: Service Request Id# = 17
ip address 209.165.200.230 255.255.255.252
frame-relay interface-dlci 343
no shutdown
!
router rip
network 209.165.200.0
no auto-summary
version  2
!
!! ------------------------------
!! Customer Edge router CE2 is provisioned as a hub in the Red VPN.
!
! Hostname: CE2
!
! Version 12.0
!
interface Serial0.888 point-to-point
description Serial0.888 fr dlci=888 : Provisioned by VPNSC: Service Request Id# = 18
ip address 209.165.200.234 255.255.255.252
frame-relay interface-dlci 888
!
no shutdown
!
router rip
network 209.165.200.0
no auto-summary
version  2
!
!! ------------------------------
!! Customer Edge router CE3 is provisioned as a hub in the Red and Blue VPNs.
!
! Hostname: CE3
!
! Version 12.0
!
interface Serial0.777 point-to-point
```

```
description Serial0.777 fr dlci=777 : Provisioned by VPNSC: Service Request Id# = 19
ip address 209.165.200.238 255.255.255.252
frame-relay interface-dlci 777
no shutdown
!
router rip
network 209.165.200.0
no auto-summary
```

OSPF Routing for the PE-CE Link

This configuration file provides an example of using the Open Shortest Path First (OSPF) protocol on the PE-CE link, and using IP numbered provisioning from the PE to CE1. CE1 is a member of a VPN called Red. CE1 is provisioned as a hub in the Red VPN and as a spoke in the Management VPN.The export map exports only the PE-to-CE link subnet from the Red VRF. VRF Red OSPF routes are redistributed into the IBGP core. The route map is used by the export map in the Red VRF to filter routes to the Management VPN.

```
!!
!! Topology:
!!
!! CE1---PE
!! --------------------------------
!! Provider Edge router: PE
!! CE1 is provisioned as a hub in the Red VPN and as a spoke in the Management VPN.
!
! Hostname: PE
!
! Version 12.0
!! Provisioned routing forwarding instance for Red VPN—vrf V10:red.
!! The route-target 200:3 is for Red VPN hub connectivity.
!! The route-target 200:4 is for Red VPN spoke connectivity.
!! The route-target 200:1 is to import a route for management from the MCE into the VRF.
!! The export map exports only the PE-to-CE link subnet from the Red VRF.
!! The export map exports the management route-target 200:2.
!
ip vrf V10:red
rd 200:10
route-target import 200:3
route-target import 200:4
route-target import 200:1
route-target export 200:3
export map grey_mgmt_vpn_VpnsRus_V10:red
!

interface Serial2/3.323 point-to-point
description Serial2/3.323 fr dlci=323 : Provisioned by VPNSC: Service Request Id# = 21
ip vrf forwarding V10:red
ip address 209.165.200.225 255.255.255.252
frame-relay interface-dlci 323
no shutdown
```

```
!!
!! OSPF routing for vrf Red using Area 0.
!! IBGP routes that reference route-targets 200:3,200:4, or 200:1 are redistributed
!! into VRF Red.
!
router ospf 10 vrf V10:red
network 209.165.200.224 0.0.0.3 area 0
redistribute bgp 200 subnets
!
!
!! Definition for the core-facing IBGP routing protocol routing instance for VRF Red.
!! VRF Red OSPF routes are redistributed into the IBGP core.
!! Exported static routes are associated with route targets 200:3 and 200:2.
!
router bgp 200
address-family ipv4 vrf V10:red
!
redistribute ospf 10 match internal external 1 external 2
exit-address-family
!
!! The route map is used by the export map in the Red VRF to filter routes
!! to the Management VPN.
!! The match matches the PE-to-CE subnet with the extended access list.
!! Route-targets 200:2 and 200:3 are exported.
!
route-map grey_mgmt_vpn_VpnsRus_V10:red permit 10
match ip address VPNSC_GREY_MGMT_ACL
set extcommunity rt 200:2 200:3
!
ip access-list extended VPNSC_GREY_MGMT_ACL
permit 209.165.200.224 0.0.0.3 255.255.255.255
!! ------------------------------
!! Customer Edge router CE1 is provisioned as a hub in the Red VPN.
!
! Hostname: CE1
!
! Version 12.0
!
interface Serial0
!
encapsulation frame-relay
!
interface Serial0.323 point-to-point
description Serial0.323 fr dlci=323 : Provisioned by VPNSC: Service Request Id# = 21
ip address 209.165.200.226 255.255.255.252
frame-relay interface-dlci 323
!
no shutdown
!
router ospf 10
network 209.165.200.224 0.0.0.3 area 0
```

OSPF Routing Using IP Unnumbered Provisioning

This configuration file provides an example of using the Open Shortest Path First (OSPF) protocol on the PE-CE link, and using IP unnumbered provisioning from the PE to CE1. CE1 is a member of a VPN called Red. CE1 is provisioned as a hub in the Red VPN and as a spoke in the Management VPN. The export map exports only the PE-to-CE link subnet from the Red VRF. VRF Red OSPF routes are redistributed into the IBGP core.

The route map is used by the export map in the Red VRF to filter routes to the Management VPN. The Loopback interface is used for unnumbered connectivity to the PE. The static route points to the Loopback address used for the unnumbered interface on the PE.

Note Unlike standard interfaces, when loopback interfaces are provisioned in VPNSC, 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.

```
!! Area 1 used is for the PE-to-CE link without default information originate.
!!
!! Topology:
!!
!! CE1---PE
!
!! --------------------------------
!! Provider Edge router: PE
!! CE1 is provisioned as a hub in the Red VPN and as a spoke in the Management VPN.
!
! Hostname: PE
!
! Version 12.0
!
!! Provisioned routing forwarding instance for Red VPN—vrf V10:red.
!! The route-target 200:3 is for Red VPN hub connectivity.
!! The route-target 200:4 is for Red VPN spoke connectivity.
!! The route-target 200:1 is to import a route for management from the MCE into the VRF.
!! The export map exports only the PE-to-CE link subnet from the Red VRF.
!! The export map exports the management route-target 200:2.
!
ip vrf V10:red
rd 200:10
route-target import 200:3
route-target import 200:4
route-target import 200:1
route-target export 200:3
export map grey_mgmt_vpn_VpnsRus_V10:red
!
!! The Loopback interface is used for the unnumbered interface in the Red VRF
!! using the VPNSC IP address pool.
!
interface Loopback1
description Provisioned by VPN-SC
ip vrf forwarding V10:red
ip address 209.165.201.1 255.255.255.255
no shutdown
```

```
!
!! The subinterface on the PE faces CE1.
!
interface Serial2/1.343 point-to-point
description Serial2/1.343 fr dlci=343 : Provisioned by VPNSC: Service Request Id# = 22
ip vrf forwarding V10:red
ip unnumbered Loopback1
frame-relay interface-dlci 343
no shutdown
!
!! OSPF routing for VRF Red using Area 1.
!! IBGP routes that reference route-targets 200:3,200:4, or 200:1
!! are redistributed into VRF Red.
!
router ospf 13 vrf V10:red
network 209.165.201.0 0.0.0.0 area 1
redistribute bgp 200 subnets
!
!! Definition for the core-facing IBGP routing protocol routing instance for VRF Red.
!! VRF red OSPF routes are redistributed into the IBGP core.
!! Exported static routes are associated with route targets 200:3 and 200:2.
!
router bgp 200
address-family ipv4 vrf V10:red
redistribute ospf 13 match internal external 1 external 2
redistribute static
exit-address-family
!
!! The static route that points to the CE loopback address is redistributed
!! into the IBGP core.
!
ip route vrf V10:red 209.165.201.2 255.255.255.255 Serial2/1.343 1
!
!! The route map is used by the export map in the Red VRF to filter routes
!! to the Management VPN.
!! The match matches-the-PE to CE subnet with the extended access list.
!! Route-targets 200:2 and 200:3 are exported
!
route-map grey_mgmt_vpn_VpnsRus_V10:red permit 10
match ip address VPNSC_GREY_MGMT_ACL
set extcommunity rt 200:2 200:3
!
ip access-list extended VPNSC_GREY_MGMT_ACL
permit 209.165.201.0 0.0.0.3 255.255.255.255
!
!! --------------------------------
!! Customer Edge router CE1 is provisioned as a hub in the Red VPN.
!
! Hostname: CE1
!
! Version 12.0
!
```

```
!! The Loopback interface is used for unnumbered connectivity to the PE.
!
interface Loopback1
description Provisioned by VPN-SC
ip address 209.165.201.2 255.255.255.255
no shutdown
!
interface Serial0
encapsulation frame-relay
!
interface Serial0.343 point-to-point
description Serial0.343 fr dlci=343 : Provisioned by VPNSC: Service Request Id# = 22
ip unnumbered Loopback1
frame-relay interface-dlci 343
no shutdown
!
!! The OSPF routing protocol uses Area 1 for the PE-to-CE link.
!
router ospf 13
network 209.165.201.2 0.0.0.0 area 1
!
!! The static route points to the Loopback address used for the
!! unnumbered interface on the PE.
!
ip route 209.165.201.1 255.255.255.255 Serial0.343 1
```

Static Routing Example

This configuration file provides an example of static routing over the PE-CE link. This configuration file provisions a default static route to the PE. The static route to the PE-CE link is redistributed into the IBGP core. VPN Solutions Center supports default and specific static routes to other VPN sites. The CE uses default routing.

```
!
!! Topology:
!
!! CE1---PE
!
!! --------------------------------
!! Provider Edge router: PE
!! CE1 is provisioned as a hub in the Red VPN and as a spoke in the Management VPN.
!
! Hostname: PE
!
! Version 12.0
!
!! Provisioned routing forwarding instance for red VPN - vrf V10:red.
!! The route-target 200:3 is for Red VPN hub connectivity.
!! The route-target 200:4 is for Red VPN spoke connectivity.
!! The route-target 200:1 imports a route from the MCE into the VRF.
!! The export map exports only the PE-to-CE link subnet from the Red VRF.
!! The export map exports the management route-target 200:2.
!
```

```
ip vrf V10:red
rd 200:10
route-target import 200:3
route-target import 200:4
route-target import 200:1
route-target export 200:3
!
export map grey_mgmt_vpn_VpnsRus_V10:red
!
!! The subinterface on the PE faces CE1; the IP address is taken from the
!! VPNSC IP address pool.
!
interface Serial2/0.454 point-to-point
description Serial2/0.454 fr dlci=454 : Provisioned by VPNSC: Service Request Id# = 20
ip vrf forwarding V10:red
ip address 209.165.202.130 255.255.255.252
frame-relay interface-dlci 454
no shutdown
!
!! The static route to the PE-to-CE link is redistributed into the IBGP core.
!
ip route vrf V10:red 209.165.202.129 255.255.255.255 Serial2/4.454 1
!
!! Definition for the core-facing IBGP routing protocol routing instance for VRF Red.
!! VRF Red static routes are redistributed into the IBGP core.
!! Exported static routes are associated with route targets 200:3 and 200:2.
!
router bgp 200
address-family ipv4 vrf V10:red
redistribute static
exit-address-family
!
!! The route map is used by the export map in the Red VRF to filter routes
!! to the Management VPN.
!! The match matches-the-PE to CE subnet with the extended access list.
!! Route-targets 200:2 and 200:3 are exported.
!
route-map grey_mgmt_vpn_VpnsRus_V10:red permit 10
match ip address VPNSC_GREY_MGMT_ACL
set extcommunity rt 200:2 200:3
!
ip access-list extended VPNSC_GREY_MGMT_ACL
permit 209.165.202.128 0.0.0.3 255.255.255.255
!
!! --------------------------------
!! Customer Edge router CE1 is provisioned as a hub in the Red VPN.
!
! Hostname: CE1
!
! Version 12.0
!
interface Serial0
encapsulation frame-relay
```

```
!
interface Serial0.455 point-to-point
description Serial0.455 fr dlci=455 : Provisioned by VPNSC: Service Request Id# = 20
ip address 209.165.202.129 255.255.255.252
frame-relay interface-dlci 455
no shutdown
!
!! A default static route to the PE is provisioned.
!! VPNSC supports default and specific static routes to other VPN sites.
!
ip route 0.0.0.0 0.0.0.0  209.165.202.130 1
```

EBGP Routing from PE to CE

This configuration file shows an example of using External BGP connectivity from a PE to a CE. A route target is provisioned to import a route from the Management CE (MCE) into the Red VPN's VRF. The export map exports only the PE-to-CE subnet from the Red VRF for connectivity to the MCE.

```
!!
!! Topology:
!!
!! CE1---PE
!
!! --------------------------------
!! Provider Edge router: PE
!! CE1 is provisioned as a hub in the Red VPN and a spoke in the Management VPN.
!
! Hostname: PE
!
! Version 12.0
!!
!! Provisioned routing forwarding instance for Red VPN—vrf V10:red.
!! The route-target 200:3 is for Red VPN hub connectivity.
!! The route-target 200:4 is for Red VPN spoke connectivity.
!! The route-target 200:1 is to import a route from the MCE into the VRF.
!! The export map exports only the PE-to-CE link subnet from the Red VRF.
!! The export map exports the management route-target 200:2.
!
ip vrf V10:red
rd 200:10
route-target import 200:3
route-target import 200:4
route-target import 200:1
route-target export 200:3
export map grey_mgmt_vpn_VpnsRus_V10:red
!
interface Serial2/6
encapsulation frame-relay
!
!! The subinterface on the PE is facing CE1; the IP address is from the VPNSC Pool.
!
interface Serial2/6.555 point-to-point
description Serial2/6.555 fr dlci=555 : Provisioned by VPNSC: Service Request Id# = 23
```

```
ip vrf forwarding V10:red
ip address 209.165.200.225 255.255.255.252
frame-relay interface-dlci 555
no shutdown
!
!! Definition for core-facing IBGP routing protocol routing instance for VRF Red.
!! VRF Red EBGP neighbor for AS 10 on the CE.
!
router bgp 200
address-family ipv4 vrf V10:red
neighbor 209.165.200.226 remote-as 10
neighbor 209.165.200.226 activate
exit-address-family
!
!! Route map is used by the export map in Red VRF to filter routes to the Management VPN.
!! The match matches the PE-to-CE subnet with the extended access list.
!! Route-targets 200:2 and 200:3 are exported.
!
route-map grey_mgmt_vpn_VpnsRus_V10:red permit 10
match ip address VPNSC_GREY_MGMT_ACL
set extcommunity rt 200:2 200:3
!
ip access-list extended VPNSC_GREY_MGMT_ACL
permit 209.165.200.224 0.0.0.3 255.255.255.255
!
!! --------------------------------
!! Customer Edge router CE1 is provisioned as a hub in the Red VPN.
!
! Hostname: CE1
!
! Version 12.0
!
interface Serial0
encapsulation frame-relay
!
interface Serial0.555 point-to-point
description Serial0.555 fr dlci=555 : Provisioned By VPNSC: Service Request Id# = 23
ip address 209.165.200.226 255.255.255.252
frame-relay interface-dlci 555
no shutdown
!
!! EBGP neighbor to AS 200 on the PE.
!
router bgp 10
neighbor 209.165.200.225 remote-as 200
```

Provisioning EBGP Routing with IP Unnumbered Scheme

This configuration file provides an example of provisioning the PE-CE link using External BGP and an IP unnumbered addressing scheme. A route target is provisioned to import a route from the Management CE (MCE) into the VRF. The loopback interface on the CE is used for an unnumbered EBGP session to the PE.

```
!! EBGP routing PE-to-CE with unnumbered provisioning PE-to-CE1
!!
!! Topology:
!! CE1---PE
!
!! -------------------------------
!! Provider Edge router: PE member
!! CE1 is provisioned as a hub in the Red VPN and a spoke in the Management VPN
! Hostname: pe
!
! Version 12.0
!
!! Provisioned routing forwarding instance for Red VPN—vrf V10:red
!! The route-target 200:3 is for Red VPN hub connectivity.
!! The route-target 200:4 is for Red VPN spoke connectivity.
!! The route-target 200:1 is to import a route from the MCE into the VRF.
!! The export map exports only the PE-to-CE link subnet from the Red VRF.
!! The export map exports the management route-target 200:2.
!
ip vrf V10:red
rd 200:10
route-target import 200:3
route-target import 200:4
route-target import 200:1
route-target export 200:3
export map grey_mgmt_vpn_VpnsRus_V10:red
!
!! The Loopback interface is used for the unnumbered interface in the Red VRF;
!! the IP address is taken from the VPNSC IP address pool
!
interface Loopback1
description Provisioned by VPN-SC
ip vrf forwarding V10:red
ip address 209.165.200.228 255.255.255.255
no shutdown
!
!! The subinterface on the PE is facing CE1.
!
interface Serial2/4.766 point-to-point
description Serial2/4.766 fr dlci=766 : Provisioned By VPNSC: Service Request Id# = 24
ip vrf forwarding V10:red
ip unnumbered Loopback1
frame-relay interface-dlci 766
no shutdown
!
!! Definition for the core-facing IBGP routing protocol routing instance for VRF Red.
```

```
!! VRF Red EBGP neighbor is in AS 10.
!! EBGP multihop is used for neighbor connectivity to the CE loopback interface.
!
router bgp 200
address-family ipv4 vrf V10:red
neighbor 209.165.200.229 remote-as 10
neighbor 209.165.200.229 activate
!
neighbor 209.165.200.229 ebgp-multihop
neighbor 209.165.200.229 update-source Loopback1
redistribute static
exit-address-family
!
!! The static route to the CE loopback is redistributed into the IBGP core.
!
ip route vrf V10:red 209.165.200.229 255.255.255.255 Serial2/4.766 1
!
!! The static route to the CE loopback is in the global table used by a recursive lookup.
!
ip route 209.165.200.229 255.255.255.255 Serial2/4.766 1
!
!! The route map is used by the export map in the Red VRF for filtering routes
!! to the Management VPN.
!! The match matches the PE-to-CE subnet with the extended access list.
!! Route-targets 200:2 and 200:3 are exported.
!
route-map grey_mgmt_vpn_VpnsRus_V10:red permit 10
match ip address VPNSC_GREY_MGMT_ACL
set extcommunity rt 200:2 200:3
!
ip access-list extended VPNSC_GREY_MGMT_ACL
permit 209.165.200.229 0.0.0.0 255.255.255.255
!
!! --------------------------------
!! Customer Edge router CE1 is provisioned as a hub in the Red VPN.
!
! Hostname: CE1
!
! Version 12.0
!
interface Serial0
!
encapsulation frame-relay
!
!! The loopback interface on the CE is used for an unnumbered EBGP session to the PE.
!
interface Loopback1
description Provisioned by VPN-SC
ip address 209.165.200.229 255.255.255.255
!
no shutdown
!
interface Serial0.766 point-to-point
```

```
description Serial0.766 fr dlci=766 : Provisioned By VPNSC: Service Request Id# = 24
ip unnumbered Loopback1
frame-relay interface-dlci 766
no shutdown
!
!! EBGP neighbor to AS 200 on the PE
!
router bgp 10
neighbor 209.165.200.228 remote-as 200
!
neighbor 209.165.200.228 ebgp-multihop
neighbor 209.165.200.228 update-source Loopback1
!
no auto-summary
!
!! The static route points to the PE loopback interface
!
ip route 209.165.200.228 255.255.255.255 Serial0.766 1
```

Cable Network Example

This configuration file provides an example of a simple cable network configuration.

```
!hostname: widgets
!
! Version 12.0
!
ip vrf V5:WidgetVPN
!
rd 200:5
!
route-target import 301:1
!
route-target import 301:2
!
route-target import 200:1
!
route-target export 301:1
!
export map grey_mgmt_vpn_VpnsRus_V5:WidgetVPN
!
interface Cable1.1
description : Provisioned by VPNSC: Service Request Id# = 14
!
ip vrf forwarding V5:WidgetVPN
ip address 209.165.200.225 255.255.255.252
!
cable helper-address 3.4.5.6
!
no shutdown
!
```

```
router bgp 200
address-family ipv4 vrf V5:WidgetVPN
exit-address-family
!
route-map grey_mgmt_vpn_VpnsRus_V5:WidgetVPN permit 10
match ip address VPNSC_GREY_MGMT_ACL
set extcommunity rt 200:2 301:1
!
ip access-list extended VPNSC_GREY_MGMT_ACL
permit 209.165.200.224 0.0.0.3 255.255.255.255
!
```

Example of Migration Process for Numbered Access List Entries to Named Access List Entries

VPN Solutions Center 2.x generates *named* access list entries instead of *numbered* access list entries in the configuration file. To provide backward compatibility for Repositories that have service requests with numbered access lists, the following migration process occurs:

When you create and deploy a new service request, VPN Solutions Center 2.x generates only named access list entries in the configuration file.

When you modify or redeploy an existing service request—which has numbered access list entries—VPNSC 2.x recognizes numbered access lists but only provisions named access lists. As a result, when you modify or redeploy a service request, VPN Solutions Center creates a named access list and numbered access list entries are deleted. This migration process continues until all the service requests have only named access lists.

The following two configlets illustrate the difference in VPNSC 2.x and previous versions (1.x).

Configlet for a New Service Request Using VPNSC 1.x with Numbered Entries

The following is a configlet generated for a new service request using VPN Solutions Center 1.x, which has number access list entries:

```
!
! Version 12.1: Generated by VPNSC on Wed Apr 04 11:50:46 2001
!
ip vrf V2:fordextranet
!
rd 9996:101
!
route-target import 9996:102
!
route-target import 9996:103
!
route-target import 9996:104
!
route-target export 9996:102
!
```

```
export map grey_mgmt_vpn_widenet_V2:fordextranet
!
interface Loopback2
description Provisioned By VPN-SC
!
ip vrf forwarding V2:fordextranet
ip address 13.13.0.1 255.255.255.255
ip address 13.13.0.1 255.255.255.255
!
interface ATM4/0/0.3 point-to-point
description ATM4/0/0.3 atm pvc vpi=3 vci=3 : Provisioned By VPNSC: Service Request ID# = 3
!
ip vrf forwarding V2:fordextranet
ip unnumbered Loopback2
!
ip unnumbered Loopback2
!
pvc 3/3
!
encapsulation aal5snap
!
no shutdown
!
router rip
address-family ipv4 vrf V2:fordextranet
!
redistribute bgp 9996 metric transparent
!
network 13.0.0.0
exit-address-family
!
no auto-summary
!
version  2
!
router bgp 9996
address-family ipv4 vrf V2:fordextranet
!
redistribute static
!
redistribute rip
exit-address-family
!
ip route vrf V2:fordextranet 13.13.0.2 255.255.255.255 ATM4/0/0.3 1
!
route-map grey_mgmt_vpn_widenet_V2:fordextranet permit 10
match ip address 1 17

set extcommunity rt 9996:105 9996:102
!
```

```
access-list 1 permit 13.13.0.2 0.0.0.0
!
end
```

Configlet for a New Service Request Using VPNSC 2.x with Named Entries

The following is a configlet generated for a new service request using VPN Solutions Center 2.x, which generates named access list entries:

```
!
! Version 12.1: Generated by VPNSC on Wed Apr 04 12:08:28 2001
!
ip vrf V3:fordextranet
!
rd 9996:102
!
route-target import 9996:102
!
route-target import 9996:103
!
route-target import 9996:104
!
route-target export 9996:102
!
export map grey_mgmt_vpn_widenet_V3:fordextranet
!
interface Loopback2
description Provisioned By VPN-SC
!
ip vrf forwarding V3:fordextranet
ip address 13.13.0.5 255.255.255.255
ip address 13.13.0.5 255.255.255.255
!
interface ATM4/0/0.7 point-to-point
description ATM4/0/0.7 atm pvc vpi=7 vci=7 : Provisioned By VPNSC: Service Request Id# = 7
!
ip vrf forwarding V3:fordextranet
ip unnumbered Loopback2
!
ip unnumbered Loopback2
!
pvc 7/7
!
encapsulation aal5snap
!
no shutdown
!
router rip
address-family ipv4 vrf V3:fordextranet
!
redistribute bgp 9996 metric transparent
```

```
!
network 13.0.0.0
exit-address-family
!
no auto-summary
!
version  2
!
router bgp 9996
address-family ipv4 vrf V3:fordextranet
!
redistribute static
!
redistribute rip
exit-address-family
!
ip route vrf V3:fordextranet 13.13.0.4 255.255.255.255 ATM4/0/0.7 1
!
route-map grey_mgmt_vpn_widenet_V3:fordextranet permit 10
match ip address VPNSC_GREY_MGMT_ACL
set extcommunity rt 9996:105 9996:102
!
ip access-list extended VPNSC_GREY_MGMT_ACL
permit ip 13.13.0.4 0.0.0.0 0.0.0.0 255.255.255.255
!
end
```

Example of 1.x Configlet Redeployed in VPN Solutions Center 2.2

Assume that a VPN Solutions Center 1.x service request generated the following configlet:

```
route-map widenet_grey_mgmt_vpn_V2:fordextranet permit 10
match ip address 1
set extcommunity rt 9996:105 9996:102
!
access-list 1 permit 13.13.0.1 0.0.0.0
```

After this service request is redeployed in VPN Soultions Center 2.2, the following configlet is generated:

```
route-map grey_mgmt_vpn_widenet_V2:fordextranet permit 10
no match ip address 1
match ip address VPNSC_GREY_MGMT_ACL
!
ip access-list extended VPNSC_GREY_MGMT_ACL
permit ip 13.13.0.1 0.0.0.0 0.0.0.0 255.255.255.255
!
no access-list 1
```

Cisco VPNSC: MPLS Solution Command Reference

This appendix provides a command reference for the new or modified Cisco IOS commands used to configure MPLS VPNs. All other commands used with MPLS VPNs are documented in the *Cisco IOS Release 12.0 Command Reference*. The commands listed in this appendix are as follows:

address-family	route-target
clear ip route vrf	show ip bgp vpnv4
exit-address-family	show ip cef vrf
import map	show ip protocols vrf
ip route vrf	show ip route vrf
ip vrf	show ip vrf
ip vrf forwarding	show tag-switching forwarding vrf
neighbor activate	debug ip bgp
rd	

In Cisco IOS Release 12.0(1)T or later, you can search and filter the output for **show** and **more** commands. This functionality is useful when you need to sort through large amounts of output, or if you want to exclude output that you do not need to see.

To use this functionality, enter a **show** or **more** command followed by the pipe character (|), one of the keywords begin, include, or exclude, and an expression that you want to search or filter on:

command | {**begin** | **include** | **exclude**} *regular-expression*

Below is an example of the **show atm vc** command in which the command output begins with the first line where the expression "PeakRate" appears:

```
show atm vc | begin PeakRate
```

For more information on the search and filter functionality, refer to the Cisco IOS Release 12.0(1)T feature module titled *CLI String Search*.

address-family

To enter the address family submode for configuring routing protocols, such as BGP, RIP, and static routing, use the **address-family** global configuration command.To disable the address family submode for configuring routing protocols, use the **no** form of this command.

VPN-IPv4 unicast

```
address-family vpnv4 [unicast]
no address-family vpnv4 [unicast]
IPv4 unicast
address-family ipv4 [unicast]
no address-family ipv4 [unicast]
```

IPv4 unicast with CE router

```
address-family ipv4 [unicast] vrf vrf_name
no address-family ipv4 [unicast] vrf vrf_name
```

Syntax Description

ipv4	Configures sessions that carry standard IPv4 address prefixes.
vpnv4	Configures sessions that carry customer VPN-IPv4 prefixes, each of which has been made globally unique by adding an 8-byte route distinguisher.
unicast	(Optional) Specifies unicast prefixes.
vrf *vrf_name*	Specifies the name of a VPN routing/forwarding instance (VRF) to associate with submode commands.

Examples

Routing information for address family IPv4 is advertised by default when you configure a BGP session using the **neighbor...remote-as** command, unless you execute the **no bgp default ipv4-activate** command.

Usage Guidelines

Using the **address-family** command puts you in address family configuration submode. Its prompt is:

```
(config-router-af)#.
```

Within this submode, you can configure address-family specific parameters for routing protocols, such as BGP, that can accommodate multiple Layer 3 address families.

To leave address family configuration submode and return to router configuration mode, enter **exit-address-family**, or simply **exit**.

Examples

The **address-family** command in the following example puts the router into address family configuration submode for the VPNv4 address family. Within the submode, you can configure advertisement of NLRI for the VPNv4 address family using **neighbor activate** and other related commands:

```
(config)# router bgp 100
(config-router)# address-family vpnv4
(config-router-af)#
```

The command in the following example puts the router into address family configuration submode for the IPv4 address family. Use this form of the command, which specifies a VRF, only to configure routing exchanges between PE and CE devices. This **address-family** command causes subsequent commands entered in the submode to be executed in the context of VRF vrf2. Within the submode, you can use **neighbor activate** and other related commands to accomplish the following:

- Configure advertisement of IPv4 NLRI between the PE and CE routers.
- Configure translation of the IPv4 NLRI (that is, translate IPv4 into VPNv4 for NLRI received from the CE, and translate VPNv4 into IPv4 for NLRI to be sent from the PE to the CE).
- Enter the routing parameters that apply to this VRF.

Enter the address family submode as follows:

```
(config)# router bgp 100
(config-router)# address-family ipv4 unicast vrf v2:blue
(config-router-af)#
```

Related Commands

Command	Description
exit-address-family	Exits address family submode.
neighbor activate	Exchanges an address with a neighboring router.

clear ip route vrf

To remove routes from the VRF routing table, use the **clear ip route vrf** EXEC command.

clear ip route vrf *vrf_name* {***** | *network* [*mask*]}

Syntax Description

vrf_name	Name of the VPN routing/forwarding instance (VRF) for the static route.
*	Deletes all routes for a given VRF
network	Destination to be removed, in dotted-decimal format.
mask	(Optional) Mask for the specified network destination, in dotted-decimal format.

Usage Guidelines

Use this command to clear routes from the routing table. Use the asterisk (*) to delete all routes from the forwarding table for a specified VRF, or enter the address and mask of a particular network to delete the route to that network.

Examples

The following command removes the route to the network 10.13.0.0 in the v1 routing table:

```
Router#clear ip route vrf v1:red 10.13.0.0
```

Related Commands

Command	Description
show ip route vrf	Displays the IP routing table associated with a VRF

exit-address-family

To exit from the address family submode, use the **exit-address-family** address family submode command.

exit-address-family

Usage Guidelines

This command has no arguments or keywords. It has no default behavior or values.

You can abbreviate this command to **exit**.

Examples

The following example shows how to exit the address-family command mode:

```
(config-router-af)#exit-address-family
```

Related Commands

Command	Description
address-family	Enters the address family submode used to configure routing protocols.

import map

To configure an import route map for a VRF, use the **import** VRF submode command.

import map `route-map`

Syntax Description

route-map	Specifies the route map to be used as an import route map for the VRF.

Defaults

There is no default. A VRF has no import route map unless one is configured using the **import map** command.

Usage Guidelines

Use an import route map when an application requires finer control over the routes imported into a VRF than provided by the import and export extended communities configured for the importing and exporting VRF.

The **import-map** command associates a route map with the specified VRF. You can filter routes that are eligible for import into a VRF, based on the route target extended community attributes of the route, through the use of a route map.

The route map might deny access to selected routes from a community that is on the import list.

Examples

The following example shows how to configure an import route map for a VRF:

```
(config)#ip vrf v1:blue
(config-vrf)#import map blue_import_map
```

Related Commands

Command	Description
ip vrf	Enters VRF configuration mode.
route-target	Configures import and export extended community attributes for the VRF.
show ip vrf	Displays information about a VRF or all VRFs.

ip route vrf

To establish static routes for a VRF, use the **ip route vrf** global configuration command. To disable static routes, use the **no** form of this command.

ip route vrf *vrf_name prefix mask [next-hop-address] [interface {interface-number}]*

[**global**] [*distance*] [**permanent**] [**tag** *tag*]

no ip route vrf *vrf_name prefix mask* [*next-hop-address*] [*interface* {*interface-number*}]

[**global**] [*distance*] [**permanent**] [**tag** *tag*]

Syntax Description

vrf_name	Name of the VPN routing/forwarding instance (VRF) for the static route.
prefix	IP route prefix for the destination, in dotted-decimal format.
mask	Prefix mask for the destination, in dotted-decimal format.
next-hop-address	(Optional) IP address of the next hop (the forwarding router that can be used to reach that network).
interface	(Optional) Type of network interface to use: ATM, Ethernet, loopback, POS (packet over SONET), or null.
interface-number	Number identifying the network interface to use.
global	Specifies that the given next hop address is in the non-VRF routing table.
distance	(Optional) An administrative distance for this route.
permanent	(Optional) Specifies that this route will not be removed, even if the interface shuts down.
tag *tag*	(Optional) Label value that can be used for controlling redistribution of routes through route maps.

Usage Guidelines

Use a static route when the Cisco IOS software cannot dynamically build a route to the destination.

If you specify an administrative distance when you set up a route, you are flagging a static route that can be overridden by dynamic information. For example, IGRP-derived routes have a default administrative distance of 100. To set a static route to be overridden by an IGRP dynamic route, specify an administrative distance greater than 100. Static routes each have a default administrative distance of 1.

Static routes that point to an interface are advertised through RIP, IGRP, and other dynamic routing protocols, regardless of whether the routes are redistributed into those routing protocols. That is, static routes configured by specifying an interface lose their static nature when installed into the routing table.

However, if you define a static route to an interface not defined in a network command, no dynamic routing protocols advertise the route unless a redistribute static command is specified for these protocols.

Examples

The following command reroutes packets addressed to network 209.165.201.0 in VRF v3:blue to the router at IP address 209.165.200.250:

```
(config)#ip route vrf v3:b
```

Related Commands

Command	Description
show ip route vrf	Displays the IP routing table associated with a VRF.

ip vrf

To configure a VRF routing table, use the **ip vrf** global configuration command. To remove a VRF routing table, use the **no** form of this command.

ip vrf *vrf_name*

no ip vrf vrf_name

Syntax Description

vrf_name	Name assigned to a VRF.

Usage Guidelines

By default, no VRFs are defined. No import or export lists are associated with a VRF. No route maps are associated with a VRF.

The **ip vrf** *vrf_name* command creates a VRF routing table and a CEF (forwarding) table, both named *vrf_name*.

The default route distinguisher value route-distinguisher is also associated with these tables.

Examples

The following example imports a route map to a VRF:

```
(Router-config)#ip vrf v2:green
(config-vrf)#rd 100:2
 route-target both 100:2
 route-target import 100:1
```

Related Commands

Command	Description
ip vrf forwarding	Associates a VRF with an interface or a subinterface.

ip vrf forwarding

To associate a VRF with an interface or subinterface, use the **ip vrf forwarding** interface configuration command. To disassociate a VRF, use the **no** form of this command.

Executing this command on an interface removes the IP address. The IP address should be reconfigured.

ip vrf forwarding *vrf_name*

no ip vrf forwarding *vrf_name*

Syntax Description

vrf_name	Name assigned to a VRF.

Defaults

The default for an interface is the global routing table.

Examples

The following example shows how to link a VRF to ATM interface 0/0:

```
(config)#interface atm0/0
(config-if)#ip vrf forward
```

Related Commands

Command	Description
ip vrf	Defines a VRF.
ip route vrf	Establishes static routes for a VRF.

neighbor activate

To enable the exchange of information with a BGP neighboring router, use the **neighbor activate** router configuration command. To disable the exchange of an address with a neighboring router, use the **no** form of this command.

neighbor {*ip-address* | *peer-group-name*} **activate**

no neighbor {*ip-address* | *peer-group-name*} **activate**

Syntax Description

ip-address	IP address of the neighboring router.
peer-group-name	Name of BGP peer group.

Defaults

The exchange of IP addresses with neighbors is enabled by default for the VPN IPv4 address family. You can disable IPv4 address exchange using the general command **no default bgp ipv4 activate**, or you can disable it for a particular neighbor using the **no** form of this command.

For all other address families, address exchange is disabled by default. You can explicitly activate the default command using the appropriate address family submode.

Examples

In the following example, a BGP router activates the exchange of a customer's IP address 10.15.0.15 to a neighboring router.

```
router bgp 100
neighbor 10.15.0.15 remote-as 100
neighbor 10.15.0.15 update-source loopback0
address-family vpnv4 unicast
neighbor 10.15.0.15 activate
exit-address-family
```

Related Commands

Command	Description
address-family	Enters the address family submode.
exit-address-family	Exits the address family submode.

rd

For a VRF to be functional, a route-distinguisher must be configured. To create routing and forwarding tables for a VRF, use the **rd** VRF submode command.

rd *route-distinguisher*

Syntax Description

route-distinguisher	Adds an 8-byte value to an IPv4 prefix to create a VPN IPv4 prefix.

Defaults

There is no default.

Usage Guidelines

A route distinguisher (RD) creates routing and forwarding tables and specifies the default route-distinguisher for a VPN. The RD is added to the beginning of the customer's IPv4 prefixes to change them into globally unique VPN-IPv4 prefixes.

An RD is either ASN-relative, in which case it is composed of an autonomous system number and an arbitrary number, or it is IP-address-relative, in which case it is composed of an IP address and an arbitrary number.

You can enter an RD in either of these formats:

- 16-bit AS number: your 32-bit number

  For example, 101:3

- 32-bit IP address: your 16-bit number

  For example, 192.168.122.15:1

Examples

The following example configures a default RD for two VRFs. It illustrates the use of both AS-relative and IP address-relative RDs:

```
(config)#ip vrf v1:blue
(config-vrf)#rd 100:3
```

```
(config-vrf)#exit
(config)#ip vrf v2:red
(config-vrf)#rd 173.13.0.12:200
```

Related Commands

Command	Description
ip vrf	Enters VRF configuration mode.
show ip vrf	Displays information about a VRF.

route-target

To create a route-target extended community for a VRF, use the **route-target** VRF submode command. To disable the configuration of a route-target community option, use the **no** form of this command.

route-target {**import** | **export** | **both**} *route-target-ext-community*

no route-target {**import** | **export** | **both**} *route-target-ext-community*

Syntax Description

import	Imports routing information from the target VPN extended community.
export	Exports routing information to the target VPN extended community.
both	Imports both import and export routing information to the target VPN extended community.
route-target-ext-community	Adds the route-target extended community attributes to the VRF's list of import, export, or both (import and export) route-target extended communities.

Defaults

There are no defaults. A VRF has no route-target extended community attributes associated with it until specified by the **route-target** command.

Usage Guidelines

The **route-target** command creates lists of import and export route target extended communities for the specified VRF.

Execute the command one time for each target community. Learned routes that carry a specific route target extended community are imported into all VRFs configured with that extended community as an import route target. Routes learned from a VRF site (for example, by BGP, RIP, or static route configuration) contain export route targets for extended communities configured for the VRF added as route attributes to control the VRFs into which the route is imported.

The route-target specifies a target VPN extended community. Like a route-distinguisher, an extended community is composed of either an autonomous system number and an arbitrary number, or an IP address and an arbitrary number.

You can enter the numbers in either of these formats:

- 16-bit AS number: your 32-bit number

  For example, 101:3

- 32-bit IP address: your 16-bit number

For example, 192.168.122.15:1

Examples

The following example shows how to configure route-target extended community attributes for a VRF. The result of the command sequence is that VRF v1:blue has two export extended communities (1000:1 and 1000:2) and two import extended communities (1000:1 and 173.27.0.130:200).

```
(config)#ip vrf v1:blue
(config-vrf)#route-target both 1000:1
(config-vrf)#route-target export 1000:2
(config-vrf)#route-target import 173.27.0.130:200
```

Related Commands

Command	Description
ip vrf	Enters VRF configuration mode.
import	Configures an import route map for the VRF.

show ip bgp vpnv4

To display VPN address information from the BGP table, use the **show ip bgp vpnv4** EXEC command.

show ip bgp vpnv4 {**all** | **rd** *route-distinguisher* | **vrf vrf_name**}

[*ip-prefix/length* [**longer-prefixes**] [*output-modifiers*]]

[*network-address* [*mask*] [*longer-prefixes*] [*output-modifiers*]] [cidr-only] [community]

[community-list] [dampened-paths] [filter-list] [flap-statistics] [inconsistent-as]

[neighbors] [paths [*line*]] [peer-group] [quote-regexp] [regexp] [summary] [tags]

Syntax Description

all	Displays the complete VPNv4 database.
rd *route-distinguisher*	Displays NLRIs that have a matching route distinguisher.
vrf *vrf_name*	Displays NLRIs associated with the named VRF.
ip-prefix/length	(Optional) IP prefix address (in dotted decimal format) and length of mask (0 to 32).
longer-prefixes	(Optional) Displays the entry, if any, that exactly matches the specified prefix parameter, as well as all entries that match the prefix in a "longest-match" sense. That is, prefixes for which the specified prefix is an initial sub-string.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
network-address	(Optional) IP address of a network in the BGP routing table.
mask	(Optional) Mask of the network address, in dotted decimal format.
cidr-only	(Optional) Displays only routes that have nonnatural net masks.
community	(Optional) Displays routes matching this community.
community-list	(Optional) Displays routes matching this community list.
dampened-paths	(Optional) Displays paths suppressed due to dampening (BGP route from peer is up and down).

filter-list	(Optional) Displays routes conforming to the filter list.
flap-statistics	(Optional) Displays flap statistics of routes.
inconsistent-as	(Optional) Displays only routes that have inconsistent autonomous systems of origin.
neighbors	(Optional) Displays details about TCP and BGP neighbor connections.
paths	(Optional) Displays path information.
line	(Optional) A regular expression to match the BGP AS paths.
peer-group	(Optional) Displays information about peer groups.
quote-regexp	(Optional) Displays routes matching the AS path “regular expression.”
regexp	(Optional) Displays routes matching the AS path “regular expression.”
summary	(Optional) Displays BGP neighbor status.
tags	(Optional) Displays incoming and outgoing BGP labels for each NLRI.

Usage Guidelines

Use this command to display VPNv4 information from the BGP database. The command show ip bgp vpnv4 all displays all available VPNv4 information. The command **show ip bgp vpnv4 summary** displays BGP neighbor status.

Examples

The following example shows output for all available VPNv4 information in a BGP routing table:

```
Router#show ip bgp vpnv4 all
BGP table version is 18, local router ID is 14.14.14.14
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP,? - incomplete
Network            Next Hop       Metric     LocPrf    Weight  Path
Route Distinguisher: 100:1 (v1:blue)
*> 11.0.0.0       50.0.0.1         0         0          101 i
*>i12.0.0.0       13.13.13.13                0          100      0       102 i
*> 50.0.0.0       50.0.0.1         0         0          101 i
*>i51.0.0.0       13.13.13.13      0         100        0        102 i
```

Table B-1 Show IP BGP VPNv4 Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
Metric	Displays the BGP metric.
LocPrf	Displays the local preference.
Weight	Displays the BGP weight.
Path	Displays the BGP path per route.

The following example shows how to display a table of labels for NLRIs that have a route-distinguisher value of 100:1.

```
Router#show ip bgp vpnv4 rd 100:1 tags
Network          Next Hop        In tag/Out tag
Route Distinguisher: 100:1 (vrf1)
2.0.0.0          10.20.0.60        34/notag
10.0.0.0         10.20.0.60        35/notag
12.0.0.0         10.20.0.60        26/notag
                 10.20.0.60        26/notag
13.0.0.0         10.15.0.15        notag/26
```

Table B-2 Show IP BGP VPNv4 rd Tags Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Network	Displays the network address from the BGP table.
Next Hop	Specifies the BGP next hop address.
In Tag	Displays the label (if any) assigned by this router.
Out Tag	Displays the label assigned by the BGP next hop router.

The following example shows VPNv4 routing entries for the VRF called v1:red.

```
Router#show ip bgp vpnv4 vrf v1:red
BGP table version is 18, local router ID is 14.14.14.14
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP,? - incomplete
Network          Next Hop         Metric    LocPrf         WeightPath
Route Distinguisher: 100:1 (vrf1)
*> 11.0.0.0       50.0.0.1          0         0          101 i
*>i12.0.0.0       13.13.13.13               0          100         0     102 i
*> 50.0.0.0       50.0.0.1          0         0          101 i
*>i51.0.0.0       13.13.13.13       0         100        0           102 i
```

Related Command

Related Commands

Command	Description
show ip vrf	Displays the VRFs and their associated interfaces.

show ip cef vrf

To display the CEF forwarding table associated with a VRF, use the show ip cef vrf EXEC command.

show ip cef vrf *vrf_name* [*ip-prefix* [**mask** [*longer-prefixes*]] [**detail**] [**output-modifiers**]]

[interface interface-number] [**adjacency** [interface interface-number] [detail] [**discard**]

[**drop**] [**glean**] [**null**] [**punt**] [*output-modifiers*]] [**detail** [*output-modifiers*]]

[**non-recursive** [**detail**] [*output-modifiers*]] [**summary** [*output-modifiers*]]

[**traffic** [**prefix-length**] [*output-modifiers*]] [**unresolved** [**detail**] [*output-modifiers*]]

Syntax Description

vrf_name	Name assigned to the VRF.
ip-prefix	(Optional) IP prefix of entries to show, in dotted-decimal format (A.B.C.D).
mask	(Optional) Mask of the IP prefix, in dotted-decimal format.
longer-prefixes	(Optional) Displays table entries for all of the more specific routes.
detail	(Optional) Displays detailed information for each CEF table entry.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
interface	(Optional) Type of network interface to use: ATM, Ethernet, Loopback, POS (packet over SONET) or Null.
interface-number	Number identifying the network interface to use.
adjacency	(Optional) Displays all prefixes resolving through adjacency.
discard	Discards adjacency.
drop	Drops adjacency.
glean	Gleans adjacency.
null	Null adjacency.
punt	Punts adjacency.
non-recursive	(Optional) Displays only nonrecursive routes.
summary	(Optional) Displays a CEF table summary.
traffic	(Optional) Displays traffic statistics.
prefix-length	(Optional) Displays traffic statistics by prefix size.
unresolved	(Optional) Displays only unresolved routes.

Usage Guidelines

Used with only the *vrf_name* argument, the **show ip cef vrf** command shows a shortened display of the CEF table. Used with the *detail* argument, the **show ip cef vrf** command shows detailed information for all CEF table entries.

Examples

This example shows the forwarding table associated with the VRF called v3:green.

```
Router#show ip cef vrf v3:green
Prefix                  Next Hop        Interface
0.0.0.0/32              receive
11.0.0.0/8              50.0.0.1        Ethernet1/3
12.0.0.0/8              52.0.0.2        POS6/0
50.0.0.0/8              attached        Ethernet1/3
50.0.0.0/32             receive
50.0.0.1/32             50.0.0.1        Ethernet1/3
50.0.0.2/32             receive
50.255.255.255/32       receive
51.0.0.0/8              52.0.0.2        POS6/0
224.0.0.0/24            receive
255.255.255.255/32      receive
```

Table B-3 Show IP CEF VRF Field Descriptions

Field	Description
Prefix	Specifies the network prefix.
Next Hop	Specifies the BGP next hop address.
Interface	Specifies the VRF interface.

Related Commands

Command	Description
show ip vrf	Displays the VRFs and their associated interfaces.
show ip route vrf	Displays the IP routing table associated with a VRF.

show ip protocols vrf

To display the routing protocol information associated with a VRF, use the **show ip protocols vrf** EXEC command.

show ip protocols vrf *vrf_name*

Syntax Description

vrf_name	Name assigned to a VRF.

Examples

The following example shows information about a VRF called v2:red.

```
Router#show ip protocols vrf v2:red
Routing Protocol is "bgp 100"
Sending updates every 60 seconds, next due in 0 sec
Outgoing update filter list for all interfaces is
Incoming update filter list for all interfaces is
IGP synchronization is disabled
Automatic route summarization is disabled
Redistributing: connected, static
Routing for Networks:
Routing Information Sources:
     Gateway          Distance    Last Update
     13.13.13.13      200          03:26:15
     18.18.18.18      200          03:26:54
Distance: external 20 internal 200 local 200
```

Table B-4 Show IP Protocols vrf Field Descriptions

Field	Description
Gateway	Displays the IP address of the router identifier for all routers in the network
Distance	Displays the metric used to access the destination route.
Last update	Displays the last time the routing table was updated from the source.

Related Commands

Command	Description
show ip vrf	Displays the VRFs and their associated interfaces.

show ip route vrf

To display the IP routing table associated with a VRF (VPN routing/forwarding instance), use the **show ip route vrf** EXEC command.

show ip route vrf *vrf_name* [**connected**] [protocol [as-number] [tag] [output-modifiers]]

[**list** number [output-modifiers]] [**profile**] [**static** [output-modifiers]]

[**summary** [output-modifiers]] [**supernets-only** [output-modifiers]]

[**traffic-engineering** [output-modifiers]]

Syntax Description

vrf_name	Name assigned to the VPN routing/forwarding instance (VRF).
connected	Displays all the connected routes in a VRF.
protocol	To specify a routing protocol, use one of the following keywords: bgp, egp, eigrp, hello, igrp, isis, ospf, or rip.
as-number	Autonomous system number.
tag	IOS routing area label.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
list *number*	Specifies the IP access list to display.
profile	Displays the IP routing table profile.
static	Displays static routes.
summary	Displays a summary of routes.
supernets-only	Displays supernet entries only.
traffic-engineering	Displays only traffic-engineered routes.

Examples

This example shows the IP routing table associated with the VRF called v1:red.

```
Router#show ip route vrf v1:red
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
         U - per-user static route, o - ODR
       T - traffic engineered route
Gateway of last resort is not set
B    51.0.0.0/8 [200/0] via 13.13.13.13, 00:24:19
C    50.0.0.0/8 is directly connected, Ethernet1/3
B    11.0.0.0/8 [20/0] via 50.0.0.1, 02:10:22

B    12.0.0.0/8 [200/0] via 13.13.13.13, 00:24:20
```

This example shows BGP entries in the IP routing table associated with the VRF called v1:red.

```
Router#show ip route vrf v1:red bgp
B  51.0.0.0/8 [200/0] via 13.13.13.13, 03:44:14
B  11.0.0.0/8 [20/0] via 51.0.0.1, 03:44:12
B  12.0.0.0/8 [200/0] via 13.13.13.13, 03:43:14
```

Related Commands

Command	Description
show ip vrf	Displays VRFs and their associated interfaces.
show ip cef vrf	Displays the CEF forwarding table associated with a VRF.

show ip vrf

To display the set of defined VRFs (VPN routing/forwarding instances) and associated interfaces, use the **show ip vrf** EXEC command.

show ip vrf [{**brief** | **detail** | **interfaces**}] [*vrf_name*] [*output-modifiers*]

Syntax Description

brief	.(Optional) Displays concise information on the VRF(s) and associated interfaces.
detail	(Optional) Displays detailed information on the VRF(s) and associated interfaces.
interfaces	(Optional) Displays detailed information about all interfaces bound to a particular VRF, or any VRF.
vrf_name	Name assigned to the VPN routing/forwarding instance (VRF)
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.

Usage Guidelines

Use this command to display information about VRFs. Two levels of detail are available: use the **brief** keyword or **no** keyword to display concise information, or use the **detail** keyword to display all information. To display information about all interfaces bound to a particular VRF, or to any VRF, use the **interfaces** keyword.

When no optional parameters are specified, the command shows concise information about all configured VRFs.

Examples

This example shows brief information for the VRFs currently configured:

```
Router#show ip vrf
   Name                  Default RD          Interfaces
   vrf1:red              100:1               Ethernet1/3
   vrf2:blue             100:2               Ethernet0/3
```

Table B-5 Show IP vrf Field Descriptions

Field	Description
Name	Specifies the VRF name.
Default RD	Specifies the default route distinguisher.
Interfaces	Specifies the network interfaces.

This example shows detailed information for the VRF called v1:blue.

```
Router#show ip vrf detail v1:blue
VRF vrf1:blue; default RD 100:1
 Interfaces:
  Ethernet1/3
Export VPN route-target communities
  RT:100:1
 Import VPN route-target communities
  RT:100:1
 No import route-map
```

Table B-6 Show IP vrf Detail Field Descriptions

Field	Description
Interfaces	Specifies the network interfaces.
Export	Specifies VPN route-target export communities.
Import	.Specifies VPN route-target import communities.

This example shows the interfaces bound to a particular VRF:

```
router#show ip vrf interfaces
Interface           IP-Address      VRF                Protocol
Ethernet2           130.22.0.33     vrf3:blue          up
Ethernet4           130.77.0.33     hub                up
router#
```

Table B-7 Show IP VRF Interfaces Field Descriptions

Field	Description
Interface	Specifies the network interfaces for a VRF.
IP-Address	Specifies the IP address of a VRF interface.
VRF	Specifies the VRF name.
Protocol	Displays the state of the protocol (up/down) for each VRF interface.

Related Commands

Command	Description
ip vrf	Enters VRF configuration mode.
rd	Configures a default route distinguisher (RD) for a VRF.
route-target	Configures import and export extended community attributes for the VRF.
import	Configures an import route map for a VRF.
ip vrf forwarding	Associates a VRF with an interface or subinterface.

show tag-switching forwarding vrf

To display label forwarding entries associated with a particular VRF or IP prefix, use the **show tag-switching forwarding vrf** EXEC command. To disable the display of label forwarding information, use the **no** form of this command.

show tag-switching forwarding vrf *vrf_name* [**ip-prefix**/length [mask]] [detail] [*output-modifiers*]

no show tag-switching forwarding vrf *vrf_name* [**ip-prefix**/length [mask]] [detail] [*output-modifiers*]

Syntax Description

vrf_name	Displays NLRIs associated with the named VRF.
ip-prefix/length	(Optional) IP prefix address (in dotted-decimal format) and length of mask (0 to 32).
mask	(Optional) Destination network mask, in dotted-decimal format.
detail	(Optional) Displays detailed information on the VRF routes.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.

Examples

The following example shows label forwarding entries that correspond to the VRF called v2:green.

```
Router#show tag-switching forwarding vrf v2:green detail
```

Related Commands

Command	Description
show tag-switching forwarding	Displays label forwarding information.
show ip cef vrf	Displays the CEF forwarding table associated with a VRF.

debug ip bgp

To display information related to processing BGPs, use the **debug ip bgp** EXEC command. To disable the display of BGP information, use the **no** form of this command.

debug ip bgp [*A.B.C.D.* | **dampening** | **events** | **in** | **keepalives** | **out** | **updates** | **vpnv4**]

no debug ip bgp [*A.B.C.D.* | **dampening** | **events** | **in** | **keepalives** | **out** | **updates** | **vpnv4**]

Syntax Description

A.B.C.D.	(Optional) Displays the BGP neighbor IP address.
dampening	(Optional) Displays BGP dampening.
events	(Optional) Displays BGP events.
in	(Optional) BGP inbound information.
keepalives	(Optional) Displays BGP keepalives.
out	(Optional) Displays BGP outbound information.
updates	(Optional) Displays BGP updates.
vpnv4	(Optional) Displays VPNv4 NLRI information.

Examples

The following example displays the output from this command:

```
Router#debug ip bgp vpnv4
03:47:14:vpn:bgp_vpnv4_bnetinit:100:2:58.0.0.0/8
03:47:14:vpn:bnettable add:100:2:58.0.0.0 / 8
03:47:14:vpn:bestpath_hook route_tag_change for v2:58.0.0.0/255.0.0.0(ok)
03:47:14:vpn:bgp_vpnv4_bnetinit:100:2:57.0.0.0/8
```

A

API Application Programming Interface. APIs are supplied as CORBA IDL files with Cisco VPN Solutions Center products. After compiling these IDL files to produce language-specific implementation files for the *target language* of your choosing, you can use these APIs to incorporate MPLS-VPN features in third-party client-application source code.

area Segments and their attached devices. Areas are usually connected to other areas through routers, making up a single autonomous system. See also *AS*. See also *region*.

AS Autonomous System. A single network or group of networks that is controlled by a common system administration group and uses a single, clearly defined routing protocol. Autonomous systems are subdivided by *areas* or *regions*. An autonomous system must be assigned a unique 16-bit number by the *IANA*.

ASBR Autonomous System Boundary Router. An area border router located between an OSPF autonomous system and a non-OSPF network. ASBRs run both OSPF and another routing protocol, such as RIP. ASBRs must reside in a nonstub OSPF area.

ATM Asynchronous Transfer Mode.

ATM-LSR A label switch router with a number of LSC-ATM interfaces. The router forwards the cells among these interfaces using labels carried in the VPI/VCI field.

ATM edge LSR A router that is connected to the ATM-LSR cloud through LSC-ATM interfaces. The ATM edge LSR adds labels to unlabeled packets and strips labels from labeled packets.

autonomous system See *AS*.

B

baseline A set of data collected from targets. For example, the latest configuration files for a list of Cisco Routers, or the latest configuration files, IP unnumbered information, and PVC information for a list of Cisco Routers. MPLS VPN Solution software automatically maintains baselines that correspond to: 1) the latest PE configuration files in the Provider Administrative Domain (with one baseline per PAD); 2) the latest configuration files of the CEs and PEs in the VPNs that the customer has defined. VPN Solutions Center software uses these baselines to create audit and topology reports.

BGP Border Gateway Protocol. An interdomain routing protocol designed for the global Internet. Exterior border gateway protocols (EBGPs) communicate among different autonomous systems. Interior border gateway protocols (IBGPs) communicate among routers within a single autonomous system. It is defined in RFC 1163.

B

BGP confederation MPLS VPNs that divide a single autonomous system into multiple sub-autonomous systems, and classify them as a single, designated confederation. The network recognizes the confederation as a single autonomous system. The peers in the different autonomous systems communicate over EBGP sessions; however, they can exchange route information as if they were IBGP peers.

Border Gateway Protocol See *BGP*.

border router A router at the edge of a provider network that interfaces to another provider's border router using the EBGP protocol.

C

Cable-CE An object within VPN Solutions Center software only to allow for provisioning cable services. A cable-CE represents the cable modem and its associated hosts for a particular site.

CAR Committed Access Rate. CAR is Cisco's traffic policing tool for instituting a QoS policy at the edge of a network. CAR allows you to identify packets of interest for classification with or without rate limiting. CAR allows you to define a traffic contract in routed networks.

CE Customer Edge Router. A CE is part of a customer network and connects to a provider edge router (PE). A CE can join any set of virtual private networks (VPNs). Each CE connects a customer site to a PE, obtaining the VPN service for that customer site, and belongs to exactly one customer. CE routers are not aware of associated VPNs. Each CE may have many configlets and may be configured by multiple SRVC service requests.

CEF Cisco Express Forwarding. An advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with large and dynamic traffic patterns. VPN routing and forwarding tables (VRFs) use CEF technology, therefore MPLS VPNs must be CEF-enabled.

CERC CE routing community. A VPN can be organized into subsets called *CE routing communities*, or CERCs. A CERC describes how the CEs in a VPN communicate with each other. Thus, CERCs describe the logical topology of the VPN. MPLS VPN Solution can be employed to form a variety of VPN topologies between CEs by building hub and spoke or full mesh CE routing communities. CERCs are building blocks that allow you to form complex VPN topologies and CE connectivity.

Class of Service See *CoS*.

configlet Router configuration commands generated by MPLS VPN Solution that are added to the CE and PE router configuration files to enable a VPN between customer sites. A configuration fragment that can be downloaded to a CE or PE to modify its current IOS command-set configuration.

CORBA Common Object Request Broker Architecture.

C

CoS CoS refers to the methods that provide *differentiated service*, in which the network delivers a particular kind of service based on the class of service specified for each packet. CoS provides specific categories of service such as Gold, Silver, and Best-Effort service classes.

CoS is a set of concrete device features in which a single network router treats traffic in different classes differently. CoS techniques provide a means of specifying policies to control network resource allocation in support of customer and applications requirements. The implementation of CoS techniques delivers measurable Quality of Service (QoS).

CoS profile Represents a set of CoS configurations offered by a provider to its customer. Each CoS profile consists of a set of CoS classes that record configuration information on how traffic is shaped and policed across the PE-CE link.

CSM Cisco Service Management System. The name of Cisco's large-picture project for service management. Many interdependent products fall within this project.

confederation A confederation divides an autonomous system into subautonomous systems and assigns a confederation identifier to the autonomous systems. In a confederation, each subautonomous system is fully meshed with other subautonomous systems. The subautonomous systems communicate using an IGP, such as Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS).

customer Requests VPN service from a *provider*. Each customer may own many customer sites and may have many service request objects.

customer edge router See *CE*.

customer network A network under the control of an end customer. The VPN connects the single customer network by connecting the isolated sites.

customer site A set of IP systems with mutual IP connectivity between them without the use of a VPN. Each customer site belongs to exactly one customer. A customer site can contain any number of CEs.

D

DES Data Encryption Standard (DES) encrypts packet data. Cisco IOS implements the mandatory 56-bit DES-CBC with the explicit initialization vector (IV). Cipher Block Chaining (CBC) requires an initialization vector to start encryption. The initialization vector is given in the IPsec packet. Triple DES (3DES) adds security by performing the operation three times with different subkeys.

DHCP Dynamic Host Configuration Protocol. A protocol that provides a mechanism for allocating IP addresses dynamically so that addresses can be reused when hosts no longer need them.

D

DLCI Data-Link Connection Identifier. A value that specifies a private virtual circuit (PVC) or a switched virtual circuit (SVC) in a Frame Relay network.

DSCP The Differentiated Service Code Point (DSCP) octet. In the IP header, DSCP classifies the packet service level. The DSCP maps to a particular observable forwarding behavior called a Per Hop Behavior (PHB). The DSCP replaces the ToS octet in the IPv4 header, and the Class octet in the IPv6 header. Currently, only the first six bits are used, allowing up to 64 different classifications for service levels. The DSCP is unstructured, but it does reserve some values to maintain limited backward compatibility with the precedence bits in the ToS octet.

E

EBGP Exterior Border Gateway Protocol. EBGPs (see *BGP*) communicate among different network domains or autonomous systems. The primary function of EBGP is to exchange network reachability information between autonomous systems, including information about the list of autonomous system routes. The autonomous systems use EGBP border edge routers to distribute the routes, which include label switching information. Each border edge router rewrites the next-hop and MPLS labels.

G

GRE Generic routing encapsulation. A tunneling protocol developed by Cisco that can encapsulate a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link to Cisco routers at remote points over an IP internetwork. By connecting multiprotocol subnetworks in a single-protocol backbone environment, IP tunneling that uses GRE allows network expansion across a single-protocol backbone environment.

I

IANA Internet Assigned Numbers Authority. Organization operated under the auspices of the ISOC as a part of the IAB. IANA delegates authority for IP address-space allocation and domain-name assignment to the InterNIC and other organizations. IANA also maintains a database of assigned protocol identifiers used in the TCP/IP stack, including BGP autonomous system numbers.

IBGP Interior Border Gateway Protocol. IBGPs (see *BGP*) communicate among routers within a single network domain or autonomous system.

IDL Interface Definition Language. Generic language for describing *APIs* for *API* servers. IDL API files must be compiled using an IDL compiler from an approved CORBA vendor to produce language-specific API files in a CORBA-supported *target language*. Using the generated target-language files you can add API-supported features to third-party client-application source code.

IGP Interior Gateway Protocol. An Internet protocol used to exchange routing information within an autonomous system. Examples of common IBGPs include IGRP, OSPF, and RIP.

IPv4 Internet Protocol, version 4. A version of IP that support a 32-bit address space.

I

IPv6 A new version of IP that will replace IPv4. The key difference between IPv4 and IPv6 is that IPv6 supports a 128-bit address space to allow many more devices to be uniquely addressed as the Internet continues its exponential growth and expands into new types of devices such as telephones, automobiles, and so on.

IS-IS Intermediate system-to-intermediate system. IS-IS is an OSI link-state hierarchical routing protocol in which ISs (routers) exchange routing information based on a single metric to determine network topology.

ISP Internet Service Provider. Provider of internet access and services through single BGP autonomous system.

L

L2TP Layer 2 tunneling protocol. Protocol used for implementing VPDNs and VPNs by tunneling PPP with multivendor interoperability and acceptance. This protocol was proposed as an alternative to IPsec, but is often used with IPsec for authentication. This protocol merges Microsoft's PPTP and Cisco's Layer 2 Forwarding (L2F) technologies.

Label-switched path (LSP) A sequence of hops (R0...Rn) in which a packet travels from R0 to Rn through label-switching mechanisms. A label-switched path can be established dynamically, based on normal routing mechanisms, or it can be established through configuration.

Label-switched path (LSP) tunnel A configured connection between two routers in which MPLS is used to carry the packet.

LFIB Label Forwarding Information Base. The LFIB manages the labels and routes that the PE routers and EBGP border edge routers receive during the exchange of VPN information.

loopback address A logical interface on a Cisco router that is always "up" and does not connect to anything.

LSA Link-state advertisement. A broadcast packet used by link-state protocols. The LSA contains information about neighbors and path costs and is used by the receiving router to maintain a routing table.

M

maintenance helper address The IP address of the DHCP server in the Multiple Service Operator (MSO) network.

MCE Management customer edge router. The network management subnet is connected to the Management CE (MCE). The MCE *emulates* the role of a customer edge router (CE), but the MCE is in provider space and serves as a network operations center gateway router. The MCE is part of a management site as defined in the MPLS VPN Solution software.

M

MD5 The MD5 algorithm takes as input a message of arbitrary length and produces as output a 128-bit "fingerprint" or "message digest" of the input. It is computationally not possible to produce two messages having the same message digest, or to produce any message having a given prespecified target message digest. The MD5 algorithm is intended for digital signature applications, where a large file must be compressed in a secure manner before being encrypted with a private (secret) key under a public-key cryptosystem such as RSA.

MIB Management Information Base.

MP-BGP Multiprotocol Border Gateway Protocol. An extension to the BGP protocol that allows VPN information to remain unique within the MPLS VPN backbone. MP-BGP also allows BGP speakers to identify routing updates that do not carry standard IPv4 prefix information. Multiprotocol (MP-BGP) and VPN-IPv4 routing information provide these extensions.

iBGP (interior BGP) refers to BGP running within a single autonomous system. *eBGP* (exterior BGP) refers to BFP running between autonomous systems.

MPE Management provider edge router. The Management PE (MPE) *emulates* the role of a PE in the provider core network. The MPE connects the MCE to the provider core network. An MPE can have a dual role as both a PE and the MPE.

MPLS Multi protocol Label Switching. An emerging standard based on a Cisco Tag Switching technology. The operating concept is to "route at the edge and switch in the core." That is, routers are used at the ingress and egress edges of the network, and switches are used in the core network.

MPLS VPN Multiprotocol Label Switching virtual private network. For MPLS VPN Solution, it is a set of *PEs* that are connected via a common "backbone" network to supply private IP interconnectivity between two or more *customer sites* for a given *customer*. Each VPN has a set of provisioning templates and policies and can span multiple *provider administrative domains* (PADs). CE Routing Communities (CERCs) in a VPN break down complex topologies into manageable subgroups.

MSO Multiple Service Operator. A company that operates more than one cable TV system.

N

network In MPLS VPN Solution, a collection of targets (routers and NetFlow Collector devices) with unique names. A target can be a member of only one network. An MPLS VPN network allows a provider to partition the working space into manageable segments that are unique and do not overlap other networks.

network management subnet Consists of the MPLS VPN Solution and Cisco IP Manager workstations on a single LAN. The MPLS VPN Solution network management subnet is required when the provider's service offering entails the management of customer edge routers (CEs). Once a CE is in a VPN, it is no longer accessible by means of conventional IPv4 routing unless one of the techniques described in this chapter is employed. The network management subnet connects directly to an *MCE*.

NLRI Network layer reachability information. BGP sends routing update messages containing NLRI to describe a route and how to get there. In this context, an NLRI is a prefix. A BGP update message carries one or more NLRI prefixes and the attributes of a route for the NLRI prefixes; the route attributes include a BGP next hop gateway address, community values, and other information.

P

Provider Administrative Domain (PAD)

A Provider Administrative Domain (PAD) is an administrative domain defined by an Internet Service Provider. Set of all PE devices in one BGP autonomous system (AS). The network owned by the PAD is called a *backbone network*. Each PAD includes a route distinguisher and route target and IP address pools. Each Provider Administrative Domain can have many *Regions* within it. If an ISP requires two AS numbers, it must consist of two provider administrative domains. Each provider administrative domain has Regions that have a route distinguisher (*RD*), a route target (*RT*), and an IP address pool from which to automatically generate IP address values during provisioning.

Provider Edge Router (PE)

A router at the edge of a provider network that interfaces to a customer's CE routers. All VPN processing occurs in the PE router. Each PE belongs to exactly one *region* of a *provider administrative domain* and connects to one or more *customer sites*. Each PE can have many *VRF* definitions and configlets, and each can be configured by many service requests.

permanent virtual circuit (PVC)

PVCs save bandwidth associated with circuit establishment and circuit tear-down in situations where virtual circuits must exist all the time. This is applicable to Frame Relay and Asynchronous Transfer Mode (ATM). In ATM terminology, a PVC is called a *permanent virtual connection*.

provider

A party supplying internet service for its *customer*. See also *ISP*.

Q

QoS

Quality of Service. The mechanisms that give network managers the ability to control the mix of bandwidth, delay, jitter, and packet loss in the network. QoS is not a device feature, it is an end-to-end system architecture. See also *CoS*.

R

RBE

See *Routed Bridged Encapsulation*.

RD

Route Distinguisher. A route distinguisher (RD) creates routing and forwarding tables and specifies the default route-distinguisher for a VPN. The RD is 8-byte value added to the beginning of the customer's IPv4 prefixes to change them into globally unique VPN-IPv4 prefixes.
An RD is either ASN-relative, in which case it is composed of an autonomous system number and an arbitrary number, or it is IP-address-relative, in which case it is composed of an IP address and an arbitrary number.

Each VPN route forwarding table (VRF) has an RD. Prefixes should use the same RD if they are associated with the same set of route targets (RTs). The community of interest association is based on the route target (RT) extended community attributes distributed with the Network Layer Reachability Information (NLRI).The RD value must be a globally unique value to avoid conflict with other prefixes.

redistribution

Redistribution allows routing information discovered through another routing protocol to be distributed in the update messages of the current routing protocol. For example, when a RIP router receives routing information from another protocol (say OSPF), it updates all of its RIP neighbors with the routing information already discovered by the OSPF protocol.

R

region A group of provider edge routers (PEs) within a single BGP autonomous system. Provider Administrative Domains are divided into regions just as customers are divided into sites. Each region belongs to exactly one provider administrative domain and can have many PEs. Regions allow a provider to employ unique IP address pools in large geographical regions. Each region is represented in the VPN Inventory Repository by a Region object.

RIP Routing Information Protocol. The simplest Interior Gateway Protocol (IGP) in the Internet. This protocol is used to exchange routing information within an autonomous system. RIP uses hop count as its primary routing metric.

route distinguisher See *RD*.

route target See *RT*.

Routed Bridged Encapsulation RBE, also known as ATM half-bridging, is the process of routing traffic from a bridged LAN without using integrated routing and bridging (IRB). RBE was developed to address the known RFC1483 bridging issues, including broadcast storms and security. Except for the fact that it operates exclusively over ATM, RBE and half-bridging functionality are identical.

Routing Information Protocol See *RIP*.

RT Route Target. A 64-bit value by which the IOS discriminates routes for route updates in VRFs.

RTR Response Time Reporter. Renamed to Service Assurance Agent (SA Agent).

RTT Round-trip time. The total time required for a packet to traverse a network to its destination and back again.

S

Service Assurance Agent Service Assurance Agent (SA Agent) provides round-trip times for various protocol: DNS, Echo, HTTP, Jitter, TCP Connect, and UDP Echo.

service level agreement See *SLA*.

service provider network A backbone network under the control of a service provider that provides transport services between customer sites.

service request VPN configuration See *SRVC*.

SHA Secure Hash Algorithm. Computes a condensed representation of a message or a data file. When a message of any length is input, the SHA-1 produces a 160-bit output called a *message digest*. The message digest can then be input to the Digital Signature Algorithm (DSA), which generates or verifies the signature for the message. The creator of the digital signature and the verifier of the digital signature must use the same hash algorithm.

S

shadow CE A simulated CE used to measure data travel time between two devices. A shadow CE is connected directly to a PE via Ethernet. SLAs configured on a shadow CE monitor the VPN routes.

site A component of a VPN customer. A collection of one or more customer edge (CE) routers. Two CEs must be in the same site if they are connected outside the VPN. A site is defined by MPLS VPN Solution software as an attribute of a VPN customer.

SLA Service Level Agreement. Service-Level Agreements (SLAs) are negotiated contracts between VPN providers and their subscribers. An SLA defines the criteria for the specific services that the subscriber expects the provider to deliver. The SLA is the only binding mechanism at the subscriber's disposal to ensure that the VPN provider delivers the services as agreed.

SRVC Service Request VPN Configuration. Represents a PE-CE link provisioning request. A complete SRVC has a unique *VRF* definition and does not belong to a PRG. It can contain no more than two *configlets*—one for the PE and one for the CE. Each SRVC can configure a PE-CE pair and is initiated by one customer. Each SRVC can attach one PE interface to one VRF table.

SNMP Simple Network Management Protocol.

SP Service Provider.

static route Route that is explicitly configured and entered into the routing table. Static routes take precedence over routes chosen by dynamic routing protocols.

T

target Single device from which information may be collected. A target may be a router or NetFlow Collector, and so on. Any device (customer edge router, provider edge router, or NetFlow Collector) from which the MPLS VPN Solution software can collect information.

target language *CORBA*-supported programming language to be generated by the IDL compiler based on the IDL *API* files. The generated target-language files can then be used to incorporate API-supported features in third-party client-application source code. For a complete list of CORBA-supported target languages, see the Object Modeling Group web site.

TCP Transmission Control Protocol.

traffic engineering The techniques and processes used to cause routed traffic to travel through the network on a path other than the one that would have been chosen if standard routing methods had been used.

traffic engineering tunnel A label-switched path tunnel that is used for engineering traffic. It is set up through means other than normal Layer 3 routing and is used to direct traffic over a path different from the one that Layer 3 routing would cause it to take.

tunneling Architecture providing the services necessary to implement any standard point-to-point data encapsulation scheme.

U

UDP User Datagram Protocol.

V

virtual private network See *VPN*.

VPIM VPN Provisioning and Inventory Manager.

VPN Virtual Private Network. At its simplest, a virtual private network (VPN) is a collection of sites that share the same routing table. A VPN is also aframework that provides private IP networking over a public infrastructure such as the Internet. In MPLS VPN Solution, a VPN is a set of customer sites that are configured to communicate through a VPN service. 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.

VPN-IPv4 addresses A VPN-IPv4 address (also referred to as a *VPNv4* address) is the combination of the IPv4 address and the 8-byte route distinguisher (RD). Combining the RD and the IPv4 address makes the IPv4 route globally unique across the MPLS VPN network. BGP considers an IPv4 address as different from another IPv4 address that has the same network and subnet mask when the route distinguishers are different.

VRF VPN Routing and Forwarding instance. The VRF is a key element in the MPLS VPN technology. VRFs exist on PEs only. A VRF is populated with VPN routes and allows multiple routing tables in a PE. One VRF is required per VPN on each PE in the VPN. The configuration information for a VPN routing and forwarding table for *PEs* that share a common route-target (*RT*) signature.

A VRF consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table. In general, a VRF includes the routing information that defines a customer VPN site that is attached to a PE router.

In the VPN inventory repository, a VRF definition is a template by which to define a VRF table in a *PE*. A VRF definition is created automatically when a service request is created. Each VRF definition belongs to exactly one *Provider Administrative Domain* and has a specific set of *CERC memberships* (the *RT* signature). All services with the same connectivity in a VPN are added to a common VRF definition.

INDEX

Numerics

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

X