AMVS

Advanced MPLS VPN Solutions

Volume 2

Version 1.0

Student Guide

Text Part Number: 97-0625-01

The products and specifications, configurations, and other technical information regarding the products in this manual are subject to change without notice. All statements, technical information, and recommendations in this manual are believed to be accurate but are presented without warranty of any kind, express or implied. You must take full responsibility for their application of any products specified in this manual.

LICENSE

PLEASE READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THE MANUAL, DOCUMENTATION, AND/OR SOFTWARE ("MATERIALS"). BY USING THE MATERIALS YOU AGREE TO BE BOUND BY THE TERMS AND CONDITIONS OF THIS LICENSE. IF YOU DO NOT AGREE WITH THE TERMS OF THIS LICENSE, PROMPTLY RETURN THE UNUSED MATERIALS (WITH PROOF OF PAYMENT) TO THE PLACE OF PURCHASE FOR A FULL REFUND.

Cisco Systems, Inc. ("Cisco") and its suppliers grant to you ("You") a nonexclusive and nontransferable license to use the Cisco Materials solely for Your own personal use. If the Materials include Cisco software ("Software"), Cisco grants to You a nonexclusive and nontransferable license to use the Software in object code form solely on a single central processing unit owned or leased by You or otherwise embedded in equipment provided by Cisco. You may make one (1) archival copy of the Software provided You affix to such copy all copyright, confidentiality, and proprietary notices that appear on the original. EXCEPT AS EXPRESSLY AUTHORIZED ABOVE, YOU SHALL NOT: COPY, IN WHOLE OR IN PART, MATERIALS; MODIFY THE SOFTWARE; REVERSE COMPILE OR REVERSE ASSEMBLE ALL OR ANY PORTION OF THE SOFTWARE; OR RENT, LEASE, DISTRIBUTE, SELL, OR CREATE DERIVATIVE WORKS OF THE MATERIALS.

You agree that aspects of the licensed Materials, including the specific design and structure of individual programs, constitute trade secrets and/or copyrighted material of Cisco. You agree not to disclose, provide, or otherwise make available such trade secrets or copyrighted material in any form to any third party without the prior written consent of Cisco. You agree to implement reasonable security measures to protect such trade secrets and copyrighted Material. Title to the Materials shall remain solely with Cisco.

This License is effective until terminated. You may terminate this License at any time by destroying all copies of the Materials. This License will terminate immediately without notice from Cisco if You fail to comply with any provision of this License. Upon termination, You must destroy all copies of the Materials.

Software, including technical data, is subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. You agree to comply strictly with all such regulations and acknowledge that it has the responsibility to obtain licenses to export, re-export, or import Software.

This License shall be governed by and construed in accordance with the laws of the State of California, United States of America, as if performed wholly within the state and without giving effect to the principles of conflict of law. If any portion hereof is found to be void or unenforceable, the remaining provisions of this License shall remain in full force and effect. This License constitutes the entire License between the parties with respect to the use of the Materials

Restricted Rights - Cisco's software is provided to non-DOD agencies with RESTRICTED RIGHTS and its supporting documentation is provided with LIMITED RIGHTS. Use, duplication, or disclosure by the U.S. Government is subject to the restrictions as set forth in subparagraph "C" of the Commercial Computer Software - Restricted Rights clause at FAR 52.227-19. In the event the sale is to a DOD agency, the U.S. Government's rights in software, supporting documentation, and technical data are governed by the restrictions in the Technical Data Commercial Items clause at DFARS 252.227-7015 and DFARS 227.7202.

DISCLAIMER OF WARRANTY. ALL MATERIALS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. In no event shall Cisco's or its suppliers' liability to You, whether in contract, tort (including negligence), or otherwise, exceed the price paid by You. The foregoing limitations shall apply even if the above-stated warranty fails of its essential purpose.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with Cisco's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of

the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

You can determine whether your equipment is causing interference by turning it off. If the interference stops, it was probably caused by the Cisco equipment or one of its peripheral devices. If the equipment causes interference to radio or television reception, try to correct the interference by using one or more of the following measures:

• Turn the television or radio antenna until the interference stops.

• Move the equipment to one side or the other of the television or radio.

• Move the equipment farther away from the television or radio.

• Plug the equipment into an outlet that is on a different circuit from the television or radio. (That is, make certain the equipment and the television or radio are on circuits controlled by different circuit breakers or fuses.)

Modifications to this product not authorized by Cisco Systems, Inc. could void the FCC approval and negate your authority to operate the product.

The following third-party software may be included with your product and will be subject to the software license agreement:

CiscoWorks software and documentation are based in part on HP OpenView under license from the Hewlett-Packard Company. HP OpenView is a trademark of the Hewlett-Packard Company. Copyright © 1992, 1993 Hewlett-Packard Company.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Network Time Protocol (NTP). Copyright © 1992, David L. Mills. The University of Delaware makes no representations about the suitability of this software for any purpose.

Point-to-Point Protocol. Copyright © 1989, Carnegie-Mellon University. All rights reserved. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

The Cisco implementation of TN3270 is an adaptation of the TN3270, curses, and termcap programs developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981-1988, Regents of the University of California.

Cisco incorporates Fastmac and TrueView software and the RingRunner chip in some Token Ring products. Fastmac software is licensed to Cisco by Madge Networks Limited, and the RingRunner chip is licensed to Cisco by Madge NV. Fastmac, RingRunner, and TrueView are trademarks and in some jurisdictions registered trademarks of Madge Networks Limited. Copyright © 1995, Madge Networks Limited. All rights reserved.

XRemote is a trademark of Network Computing Devices, Inc. Copyright © 1989, Network Computing Devices, Inc., Mountain View, California. NCD makes no representations about the suitability of this software for any purpose.

The X Window System is a trademark of the X Consortium, Cambridge, Massachusetts. All rights reserved.

Advanced MPLS VPN Solutions, Revision 1.0: Student Guide

Table of Contents

Volume 1

Volume 2

5

MPLS VPN Topologies

Overview

This chapter describes the most commonly used MPLS VPN topologies and the design and implementation issues associated with them.

It includes the following topics:

- Simple VPN with optimal Intra-VPN routing
- Using BGP as the PE-CE routing protocol
- Overlapping Virtual Private Networks
- Central Services VPN solutions
- Hub-and-Spoke VPN solutions
- Managed CE Router Service

Objectives

Upon completion of this chapter, you will be able to perform the following tasks:

- Design and implement simple VPN solutions with optimal intra-VPN routing
- Design and implement various routing protocols within VPNs
- Design and implement central services VPN topologies
- Design and implement hub-and-spoke VPN topologies
- Design and implement VPN topology required for managed router services

Simple VPN with Optimal Intra-VPN Routing

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe the requirements of simple VPN solutions
- Describe the routing model of these solutions
- Describe the optimal intra-VPN routing data flow
- Select the optimal PE-CE routing protocol based on user requirements
- Integrate the selected PE-CE routing protocol with the MPLS VPN backbone MP-BGP routing

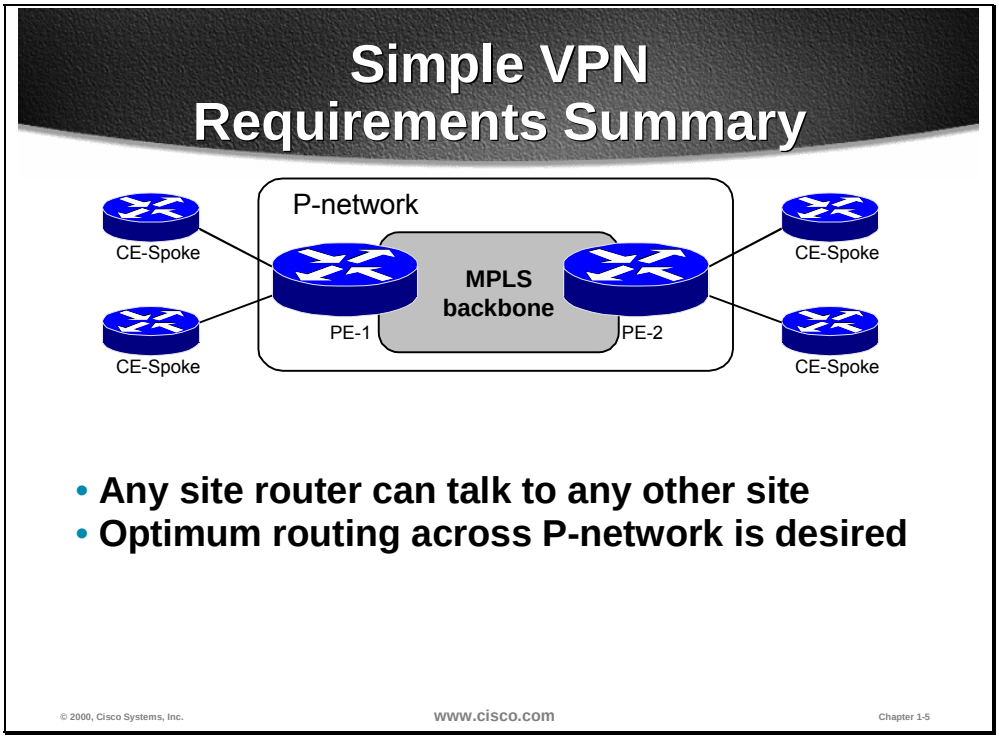

In contrast with other VPN technologies, MPLS VPN supports optimum any-to-any connectivity between customer sites (equivalent to the full mesh of overlay VPN networks) without the end customer having to manually configure anything. The provider only needs to configure the VPN in the Provider Edge (PE) routers. The so-called “hub-and-spoke” topology, which was primarily used to reduce the cost of the network, is no longer needed. The interconnection of CE sites is done automatically by using BGP and an IGP to find the shortest path.

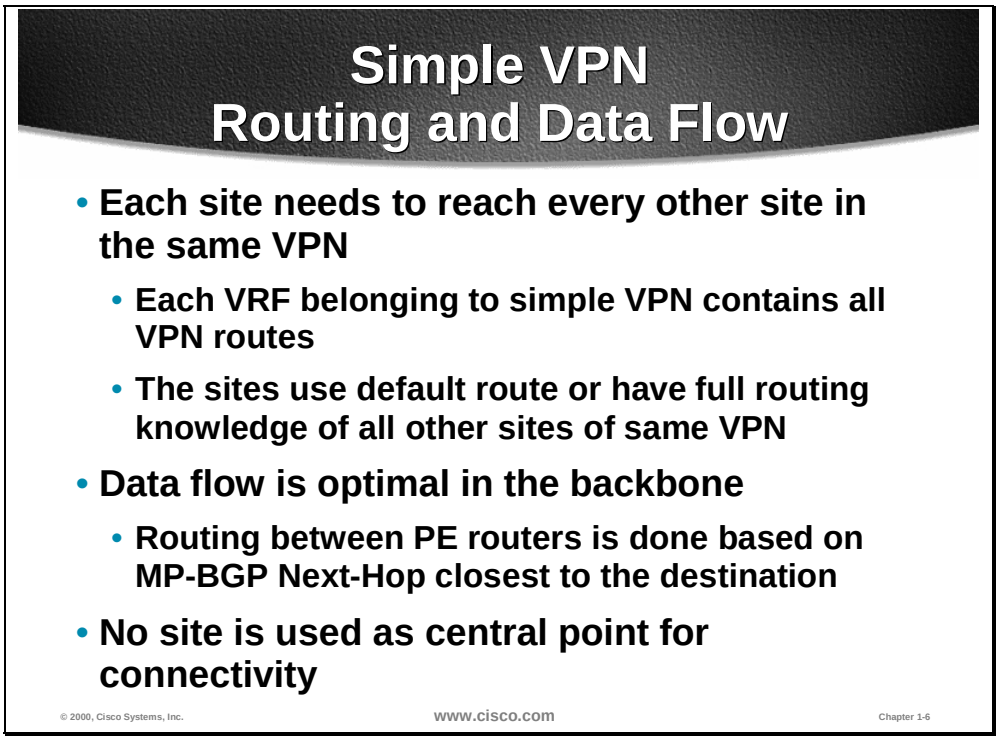

MPLS VPN architecture by default provides optimal routing between CE sites. A CE site can have full internal routing for its VPN or just a default route pointing to the PE router. The PE routers, however, need to have full routing information for the MPLS VPN network in order to provide connectivity and optimal routing. A MP-BGP next-hop address is used to find a label for a VPN destination network and the backbone IGP provides the optimal routing towards the next-hop address.

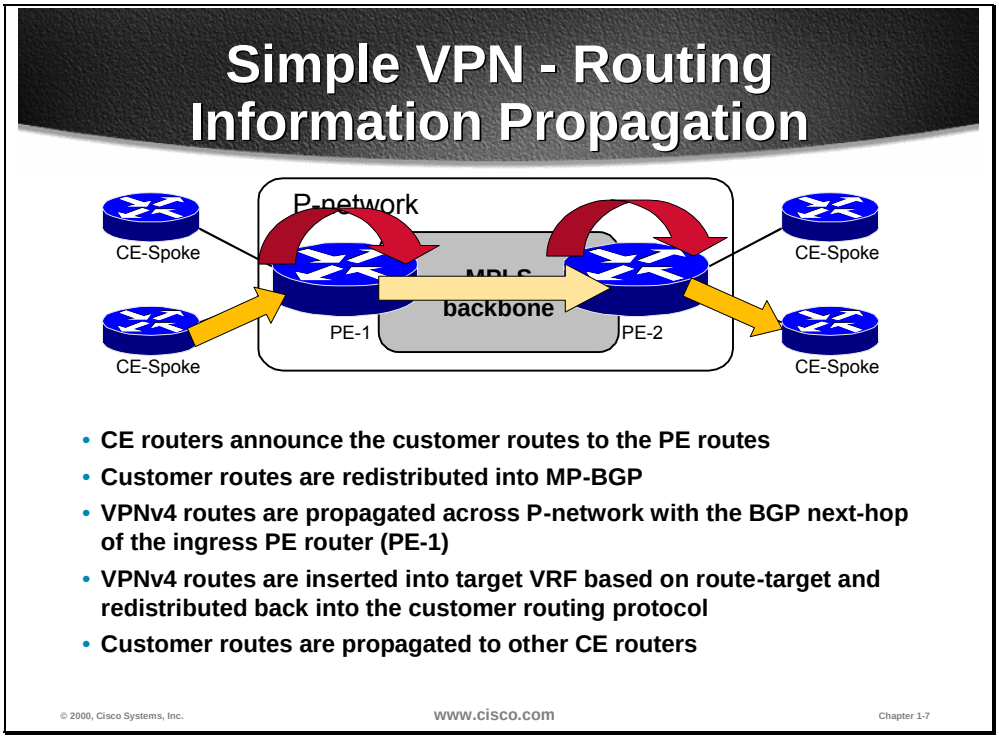

When a Customer Edge (CE) router announces a network through an IGP, the PE router will redistribute and export it into Multiprotocol BGP, converting an IPv4 address into a VPNv4 address. The following list contains the most significant changes that happen with redistribution and export:

- IPv4 Network Layer Reachability Information (NLRI) is converted into VPNv4 NLRI by pre-pending a route distinguisher (for example, a route distinguisher 12:13 could be prepended to an IPv4 prefix 10.0.0.0/8 resulting in a VPNv4 prefix 12:13:10:10.0.0.0/8)

Note NLRI is a BGP term for a prefix (address and subnet mask)

- VPNv4 NLRI also contains a label that will be used to identify the outgoing interface or the VRF where a routing lookup should be performed
- A *route target* extended community is added based on the VRF configuration

The PE router will forward VPN_IPv4 networks to all other PE routers that will use the route target community to identify the VRFs where this information has to be imported. The received VPN label will be used as the second label and the BGP next-hop label (learned via LDP) will be used as the top label for packets going to CE routers connected to distant PE routers.

The PE router will then redistribute the VPN_IPv4 network into the IGP used between the PE and the CE and send it to the CE router.

The MPLS VPN core network is not visible to the CE routers. The BGP part of the routing information propagation is only seen as slower convergence.

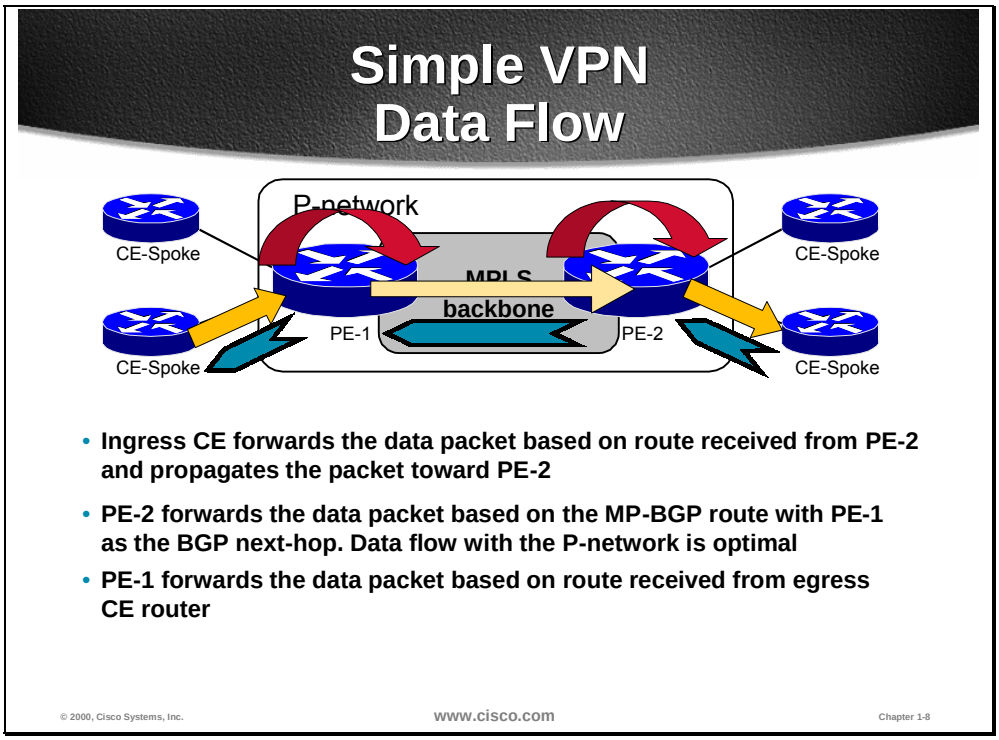

In the slide above, the CE router finds the destination in its IP routing table (learned through IGP or based on a static default route). PE-2 has learned about the destination through MP-BGP and labels each packet from the CE router with the VPN label (second label) and the next-hop label (top label).

The core routers are doing label switching based on the top label. The last core router before PE-1 will pop the top label (penultimate hop popping). PE-1 will identify the outgoing interface or the VRF by looking at the second label, which at this time is the top and only label. The packet sent to the CE is no longer labeled.

Note Please refer to **MPLS VPN Technology** lesson for more information on MPLS VPN packet forwarding.

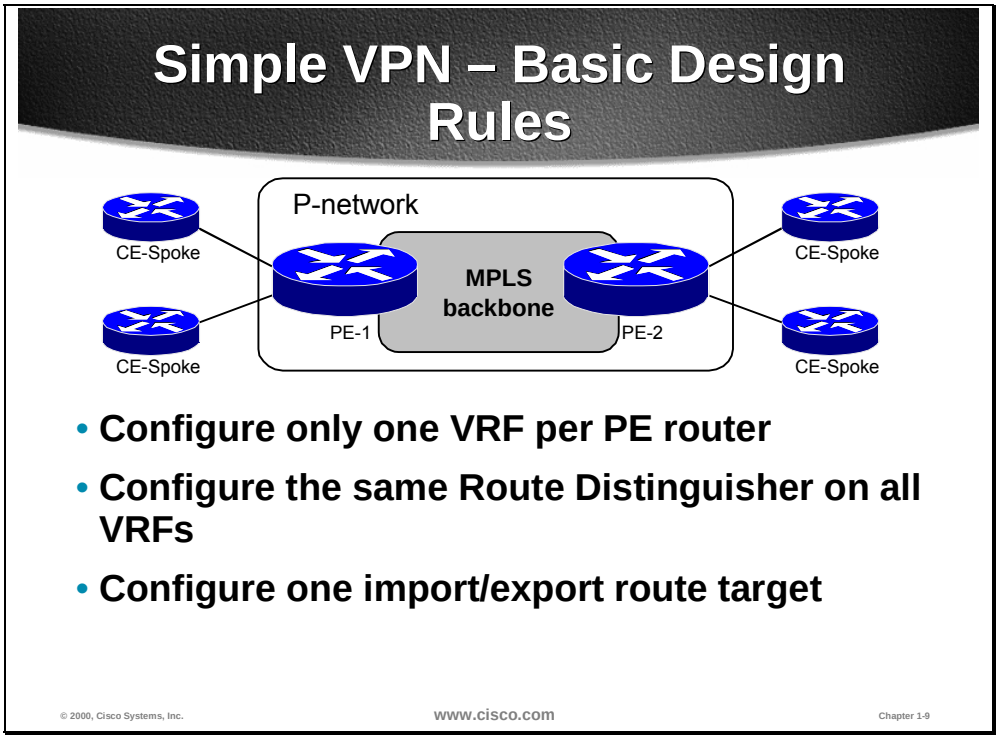

To optimize performance, reduce configuration efforts and conserve memory on the PE router on which you should minimize the number of VRFs per router.

Using one VRF per VPN per PE router will reduce memory requirements and CPU load. This is possible because the routing requirements for all CE routers in the same VPN are the same. Using one VRF per VPN can also improve convergence between CE routers connected to the same PE router.

Using the same route distinguisher for VRFs that are used for the same VPN will also conserve memory.

Only one route target is needed for a simple VPN. Any additional route targets are unnecessary and will consume at least 64 bits per routing update.

Using the same route distinguisher and route target for a simple VPN helps to ease the management, monitoring, and troubleshooting of the MPLS VPN network.

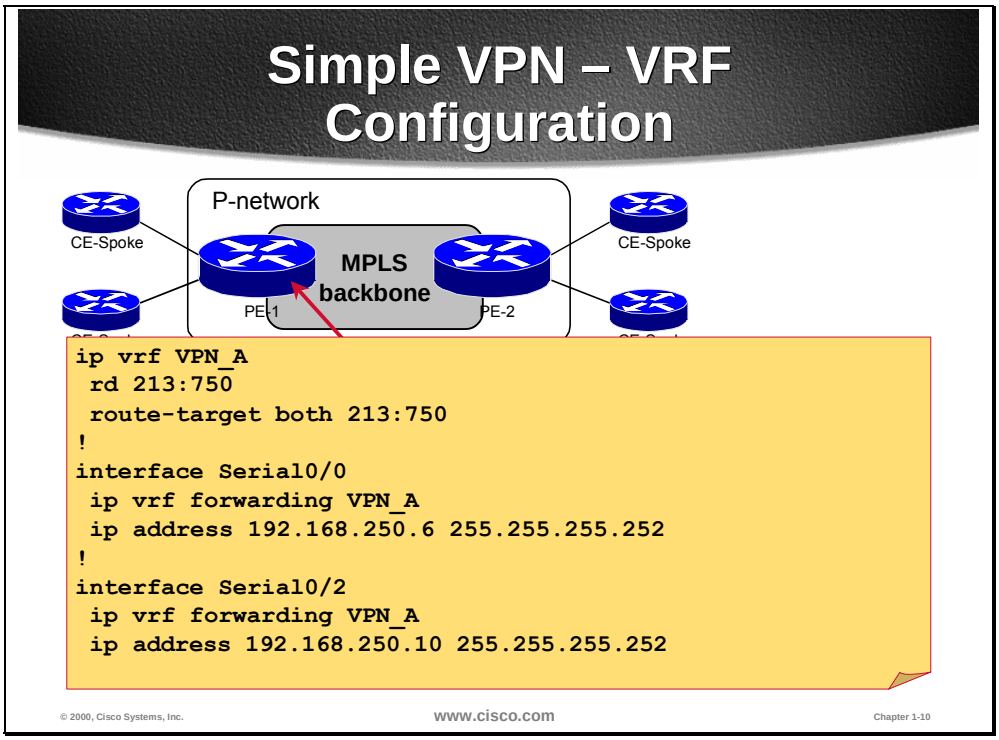

In the example above, we have two interfaces in the same VRF. We are using the same numbering scheme for route distinguishers and route targets.

Note There is no routing configuration in this example. This example only shows how to create a virtual router (VRF – virtual routing and forwarding instance) and to assign interfaces to it.

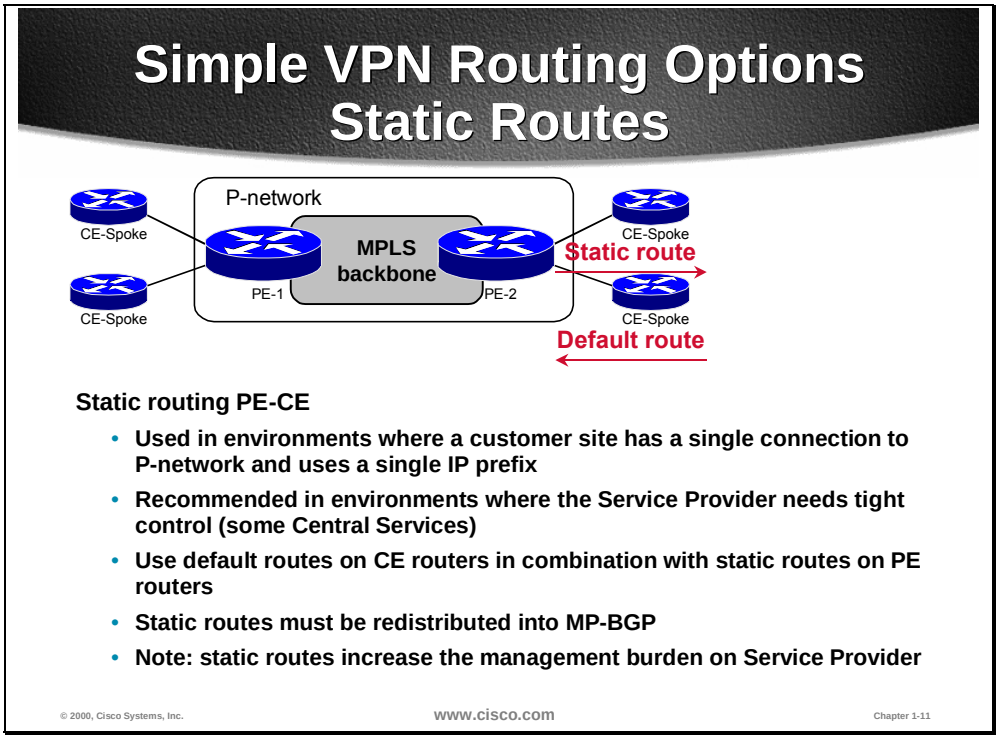

One of the routing options in a simple VPN is to use a static route on the PE and a static default route on the CE. This is an optimal solution for simple spoke VPN sites (sites with only one link into the P-network) that have only one IP subnet per site.

Using static routes also prevents the customer or the service provider from intentionally or accidentally flooding the other with a false and possibly overwhelming amount of routing information and thus strengthens the Service Provider's control over customer routing.

You must redistribute the static routes into MP-BGP to inform other PE routers of remote networks belonging to the customer VPN.

Note The static routes increase the management burden on the Service Provider as every change inside the customer's network must be coordinated with the Service Provider.

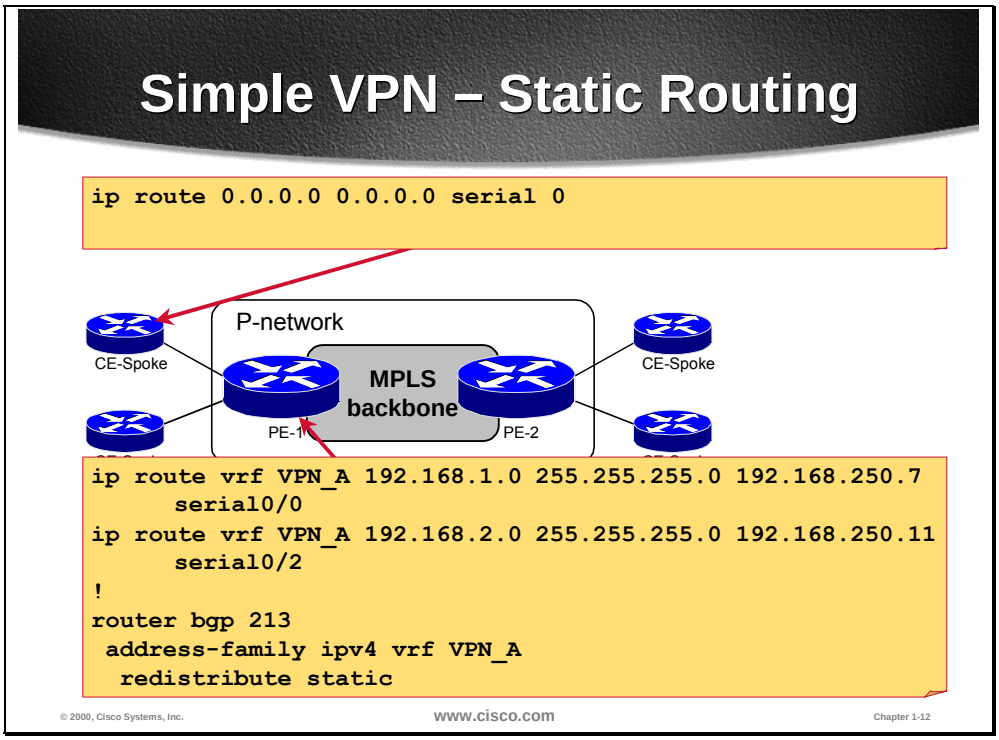

This example shows how to create a static route in a VRF routing table. The redistribution of static route into BGP should be configured in the address family of the VRF where the static route has been inserted.

Note You have to configure at least one export route target in the VRF to start advertising this network via MP-BGP.

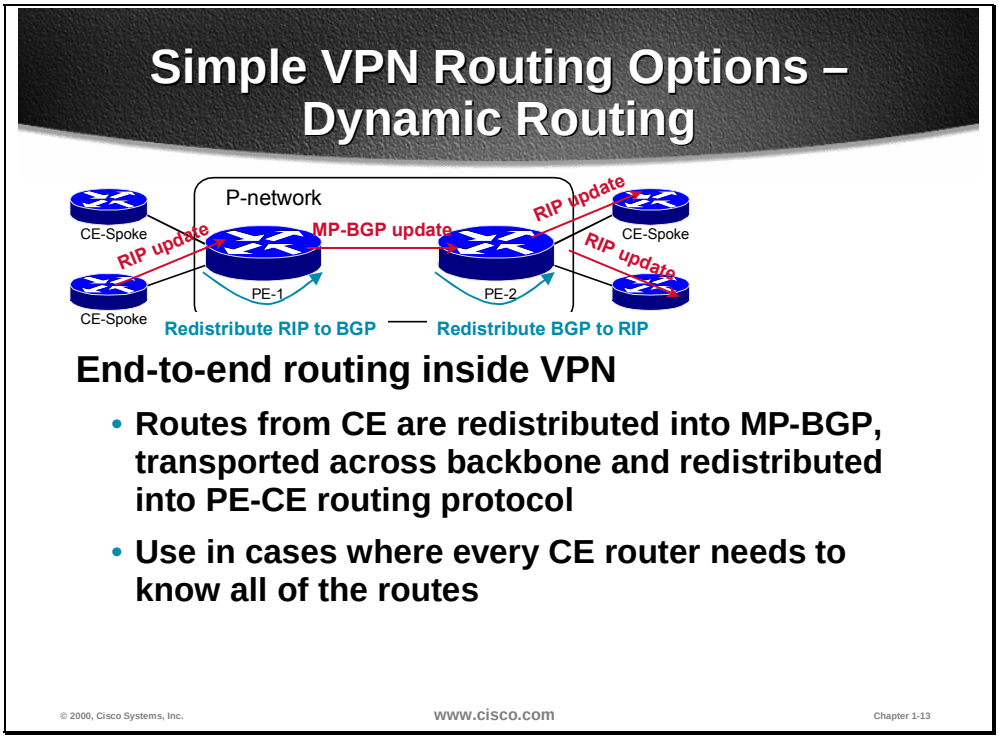

Instead of using static routing you can use an IGP, such as RIP version 2 or OSPF, to advertise customer networks between the PE-routers and the CE-routers. This option is normally used when the customer manages the CE routers, when there is more than one IP prefix per customer site, or when the site is multi-homed (has more than one link into the P-network or a separate Internet connection).

The IGP metric can be preserved by copying it into the BGP MED attribute (default action) and copying it back from the MED attribute into the IGP metric (configured with **metric transparent** option of the **redistribute** command).

Note Using transparent redistribution can be dangerous if you use different CE-PE routing protocols. For example: a redistributed OSPF update can create a BGP update where the MED attribute holds the OSPF cost taken from the routing table and this value can be large. When such update is redistributed into RIP, the hop count would have a large value, which is interpreted as an unreachable destination. In networks where the CE-routers use different routing protocols, the IGP metric cannot be deduced from BGP MED attribute and has to be specified manually.

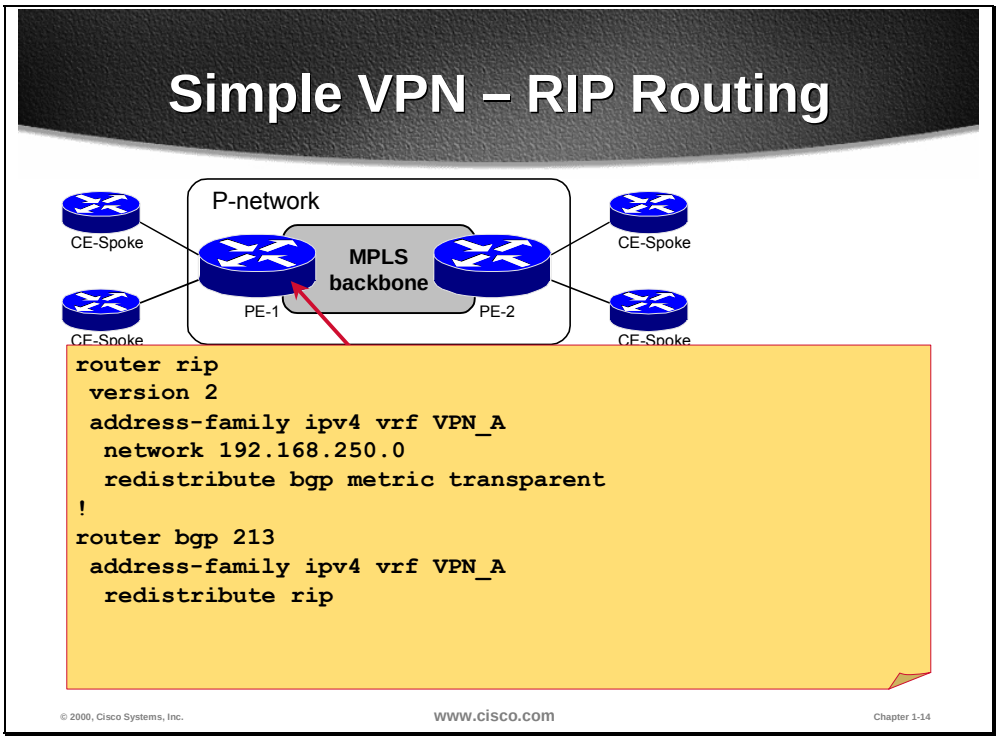

This example shows the configuration of RIP and BGP with RIP hop count propagation where RIP hop count is preserved while the route is transported across MPLS VPN backbone via MP-IBGP by being stored in the BGP MED attribute.

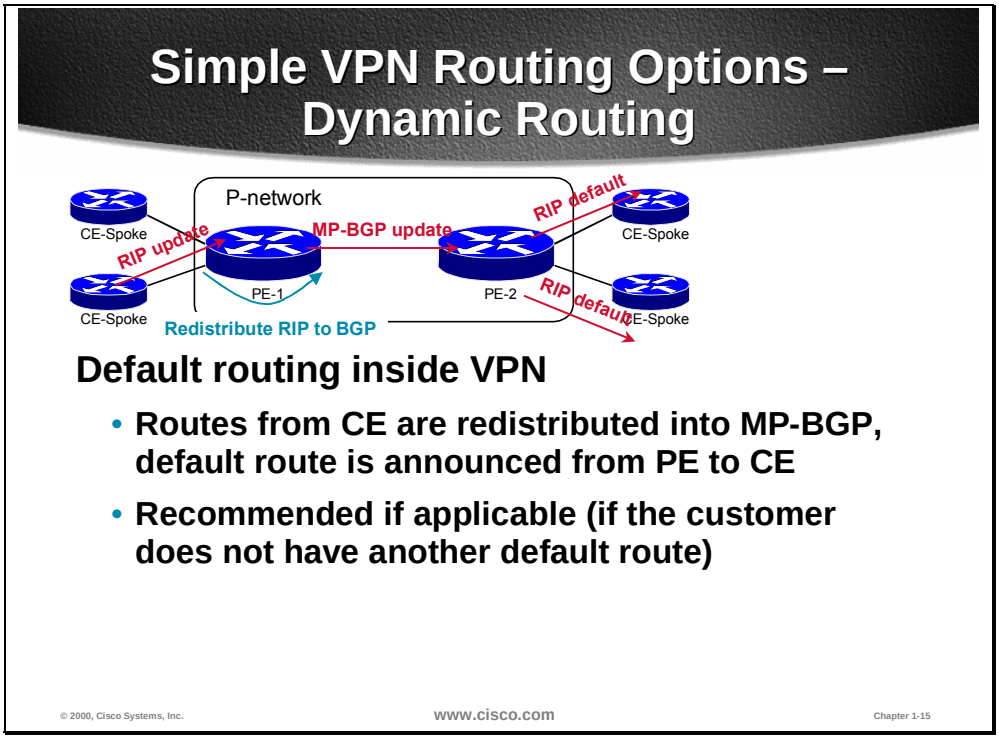

Instead of sending all the networks to the customer, we can send only a default route toward the CE routers. The PE router will accept IGP updates from the CE routers and send them to other PE routers via MP-iBGP, but it will only send a default route to the CE routers.

This approach can be used when customer sites have more than one IP prefix per site, which forces us to use a routing protocol instead of static routes. The CE routers, however, have one single connection to the MPLS VPN backbone (stub sites).

Note Default routing from the PE-router toward central VPN sites may not work well if these sites already have a different default route, for example, toward the Internet firewall. A similar situation might apply in situations where the customer is using a large number of Internet exit points throughout the VPN.

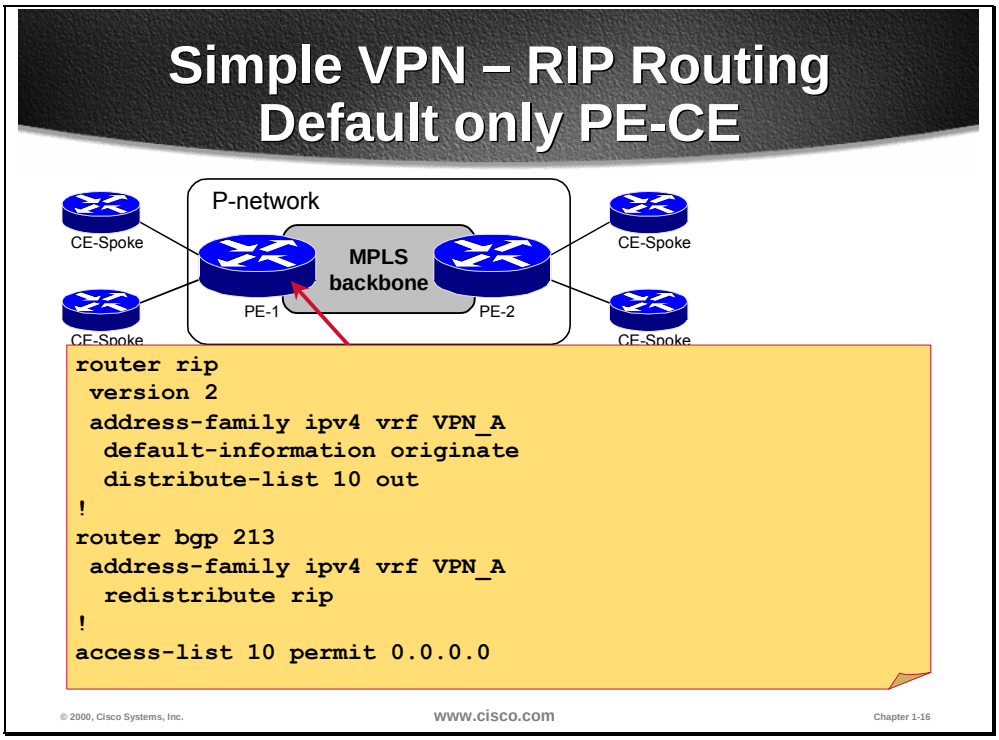

The example above shows the configuration steps needed to generate a default route in the RIP updates and a filter that denies everything but the default route. RIP neighbors will only receive a default route while other PE routers will receive all customer subnets via MP-iBGP. Redistribution from BGP to RIP is no longer necessary.

Note Classless routing has to be configured on the CE routers with the **ip classless** configuration command in order for this setup to work in all circumstances.

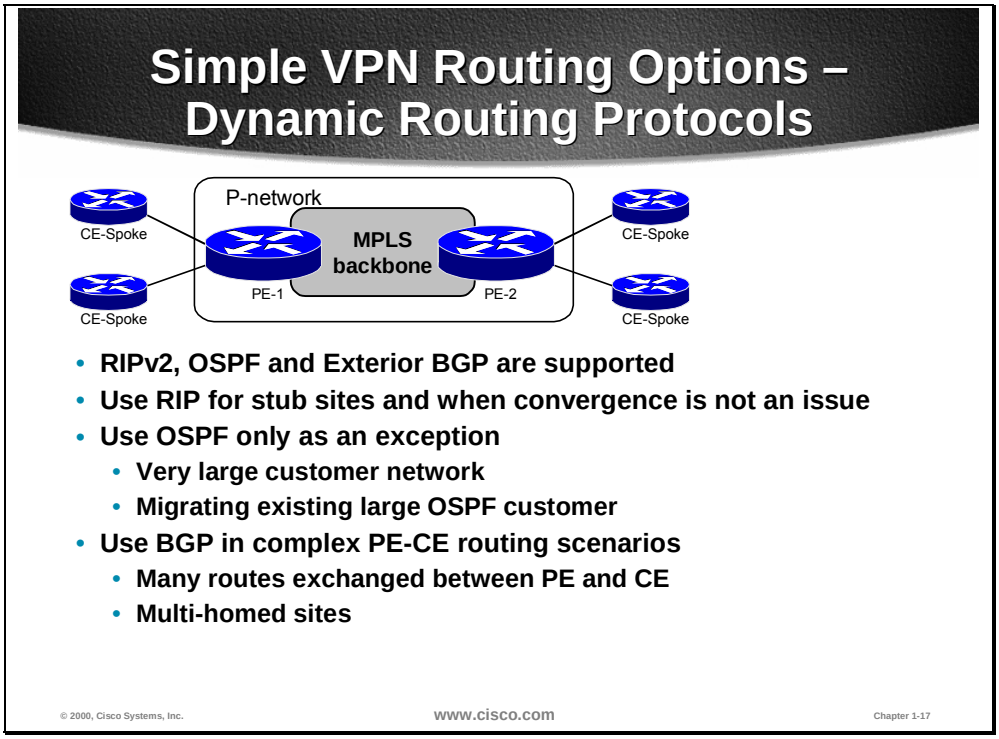

The following dynamic routing protocols are supported for PE-CE routing information exchange:

- RIP for stub sites where there are more subnets per site or where the service provider does not manage the CE; only the default route should be sent to the CE
- BGP for multi-homed sites – highly recommended to prevent suboptimal routing
- OSPF – should only be used for extremely large VPN customers where the customer insists on using OSPF for migration or intra-site routing purposes

Note OSPF is not recommended as the default IGP between the PE-routers and the CE-routers, as the number of VRFs that can support OSPF on a single PE-router is limited. Please refer to **MPLS VPN Implementation** lesson for more details.

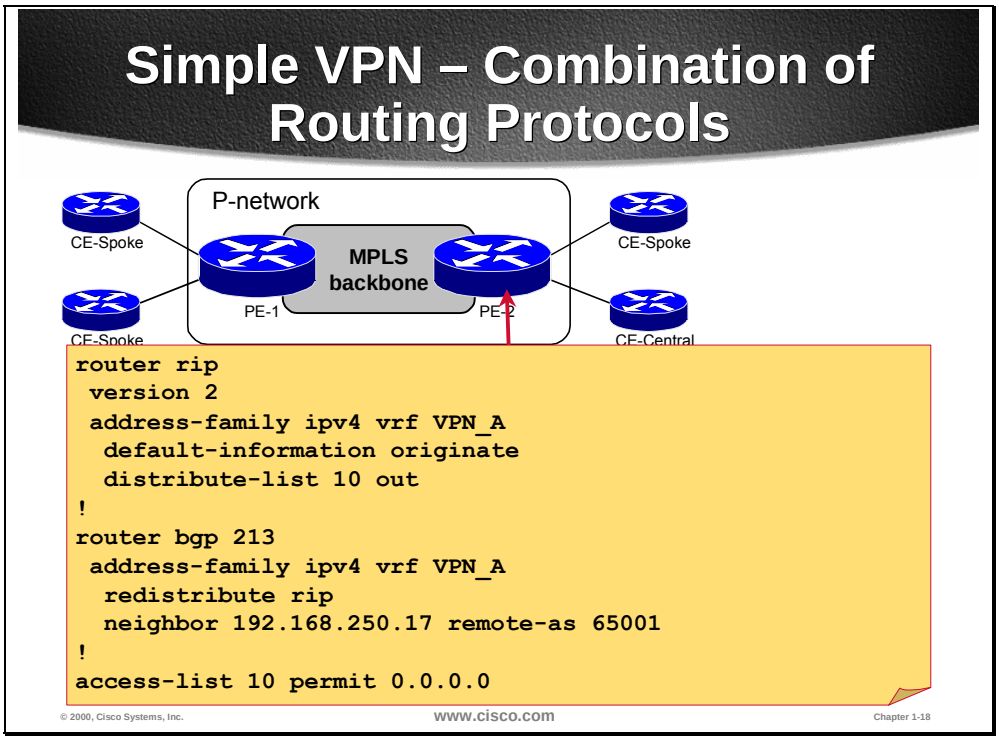

The example above shows a sample customer configuration where two different routing protocols are used between the PE-routers and the CE-routers in the same VPN. RIP is used with a spoke customer site and BGP is used to propagate the full VPN_A routing information to the central VPN_A site. In this example we only need to do redistribution from RIP to BGP because there is no need to send the full VPN routing information to other CE routers. Instead we are just sending the default route and filtering out everything else.

Summary

A MPLS VPN solution requires MPLS to be enabled on all core routers, MP-BGP to propagate the information about customer networks and an IGP within the core to find the shortest path to the loopback s of PE routers.

To learn about the customer networks we can use static routes for simple stub sites, RIPv2 for larger stub sites or sites that that are not managed by the service provider, BGP for multi-homed sites and OSPF only if really necessary.

When an update is received from a CE router, a PE router has to redistribute and export it into MP-BGP with at least one Route Target extended community. The Route Target is the used to identify the appropriate VRF on other PE routers where the update is imported and redistributed back into the routing protocol used within the VPN.

Review Questions

Answer the following questions

- What are the basic requirements for simple VPN service?
- What are the routing requirements for simple VPN service?
- What should the CE-PE-PE-CE data flow be for simple VPN service?
- Which PE-CE routing protocol would you use for simple VPN service?
- How many VRFs per PE-router do you need to implement simple VPN service?
- How do you integrate RIP running between PE and CE with MP-BGP running in the MPLS VPN backbone?
- When would you use static routing between PE and CE routers?
- When would you be able to use default routing from PE toward CE?
- When would you use OSPF between PE and CE routers?
- What are the drawbacks of offering OSPF as the PE-CE routing protocol to your customers?

Using BGP as the PE-CE Routing Protocol

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe the situations that warrant using BGP as the PE-CE routing protocol
- Describe the different design models that can be used when running BGP between PE and CE routers
- Explain the implications of using the same AS number on multiple customer sites

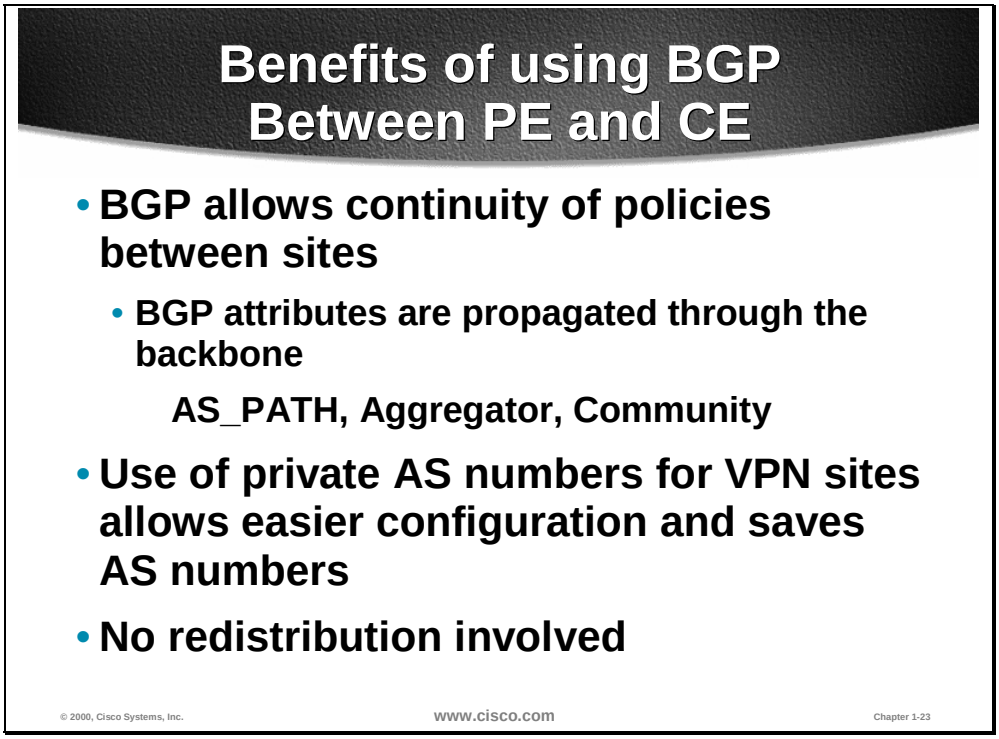

BGP is considered to be a complex routing protocol by most customers and is therefore avoided by some of the MPLS VPN customers. While BGP is best avoided in simple scenarios where the customers only have single-homed spoke sites, its complexity is more than compensated in scenarios where a complex routing policy is needed between the Service Provider and the customer network.

Deploying BGP as the routing protocol between the PE-routers and the CE-routers enable establishment of a consistent end-to-end routing policy as the BGP attributes set by one customer site are transparently propagated to other customer sites. There is also no need for route redistribution, since the same routing protocol is used across the whole network.

When using the BGP as the routing protocol between the PE and the CE router, the BGP session established between these two routers is a standard BGPv4 session. The updates received from the neighboring CE routers end up in the appropriate address family of the BGP table and no redistribution is required. Exporting from VRF into the multi-protocol BGP is still required to prepend a route distinguisher to the IPv4 prefix and to attach the route target(s) to the resulting VPNv4 route.

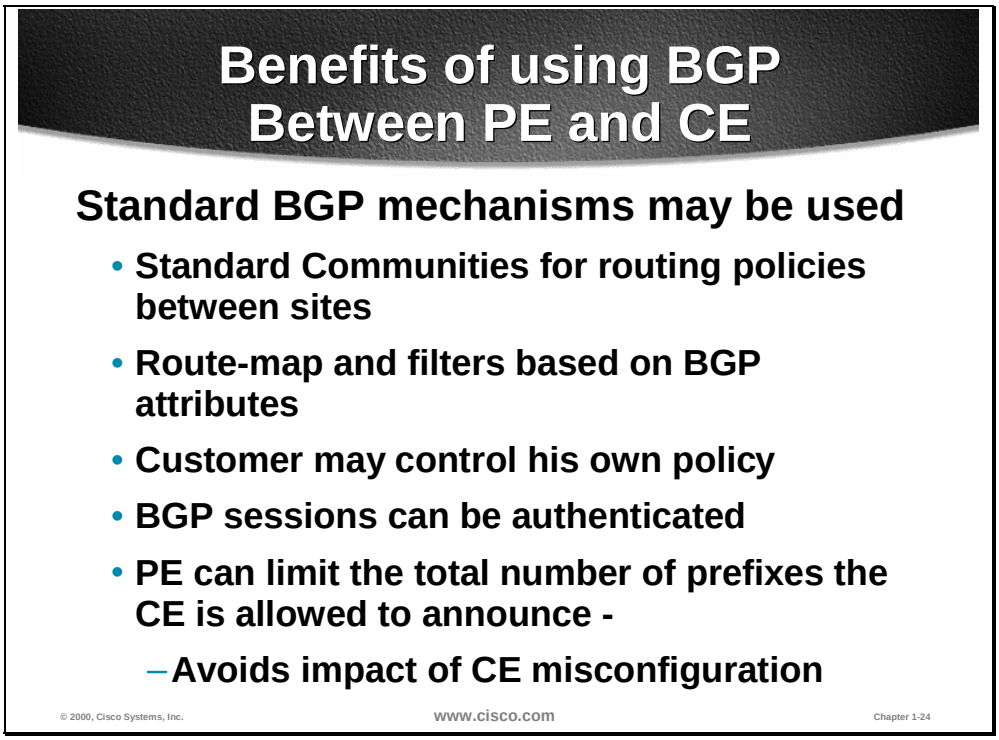

BGP has a wide range of filtering and other options that either the service provider (PE) or the customer (CE) can deploy to implement desired routing policies:

- Distribute lists to filter based on networks and/or subnet masks
- Prefix lists to filter based on networks and/or subnet masks
- Filter lists to filter based on the AS path
- Route maps to filter on subnets, subnet masks, AS path, communities, next-hop addresses
- Route maps to change BGP parameters (weight, local preference, MED, BGP communities or prepend local AS number to the AS path)
- Setting per-neighbor weight
- Setting the maximum number of updates accepted from a neighbor

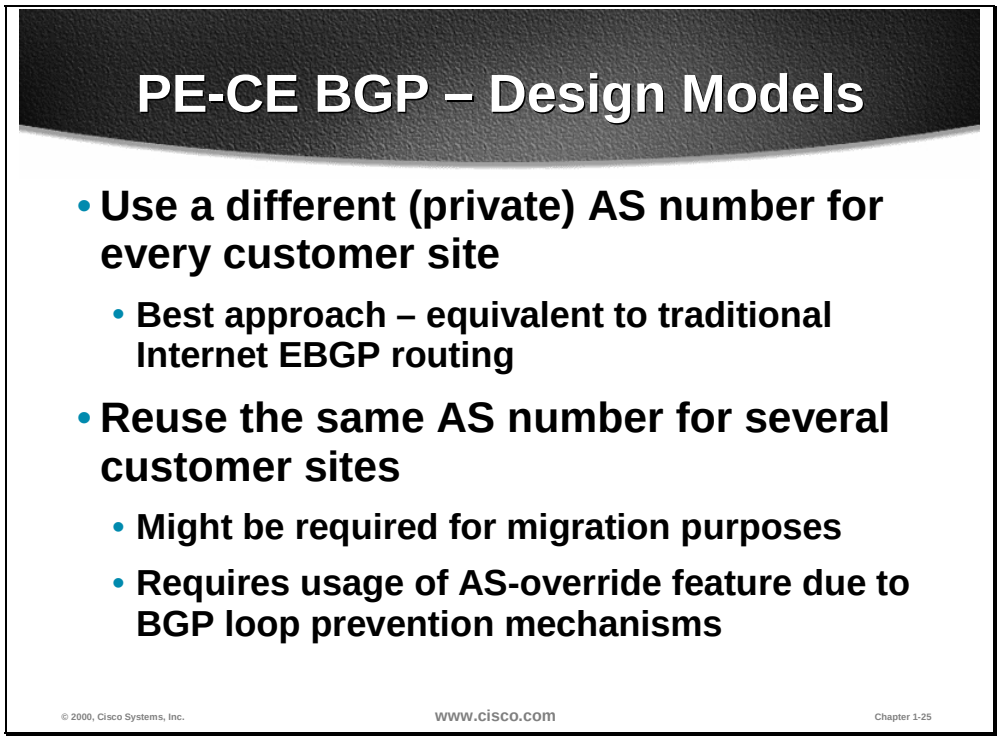

There are a number of different options for choosing the AS numbers for customer sites:

- Each site has a different private AS number (easy to configure; consumes a large number of private AS numbers)
- Each VPN has a different private AS number that is used for all the sites (AS-override feature is needed)
- Some VPNs use registered AS numbers (if the customer is also a service provider)
- All VPNs use the same private AS number (only one private AS number needed, but you need as-override feature)

Using a different AS number for every site simplifies the configuration, but consumes a large number of private AS numbers (from 64512 to 65535). For VPNs with less than 1024 sites that don't overlap, this limitation is not an issue.

Note The private AS numbers used by one VPN can be reused by another VPN as long as these VPNs do not overlap. If they do overlap, the AS-override feature is needed, similarly to the next case where you reuse the same AS number at multiple sites.

Reusing the same AS number for all (or multiple) sites belonging to the same VPN requires the usage of the AS-override feature, which is explained on the next page.

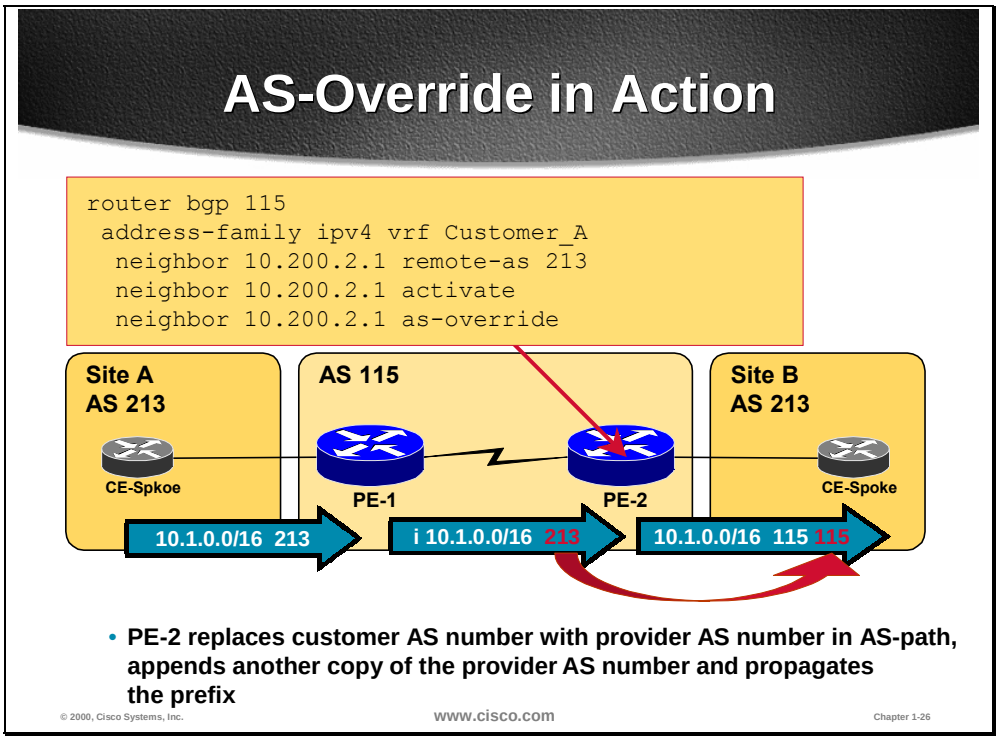

If site A and site B use the same AS number, then an update originating in either site will not be accepted by the other site because the receiving CE router finds its own AS number in the AS path and assumes that it's faced with a BGP routing information loop.

Because this is not a routing loop, we can overwrite the original AS-number (in this example 213) with the service provider's AS-number (115). PE2 will automatically prepend the service provider's AS number once more as part of normal EBGP update processing. Now site B will accept the update because it does not contain its own number in the AS path.

Summary

BGP is primarily used with those CE sites that have multiple connections to the MPLS VPN core. Using any other routing protocol can cause some traffic to be sub-optimally routed through the multi-homed site. BGP will normally prevent this from happening without any special configuration.

When designing BGP one can use private AS numbers for customer sites. AS numbers can also be reused which requires the AS-override feature to be used on the PE routers to allow updates from one site to be accepted on another site with the same AS number.

Review Questions

Answer the following questions

- When would you use BGP as the PE-CE routing protocol?
- When would you use the same AS number for several sites?
- When would you use a different AS number for every site?
- Which BGP features would you use to support the customers that use the same AS number at multiple sites?

Overlapping Virtual Private Networks

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe the requirements and typical usages of overlapping VPN solutions
- Describe the routing model and data flow of these solutions
- Design and configure overlapping VPNs in an MPLS VPN backbone

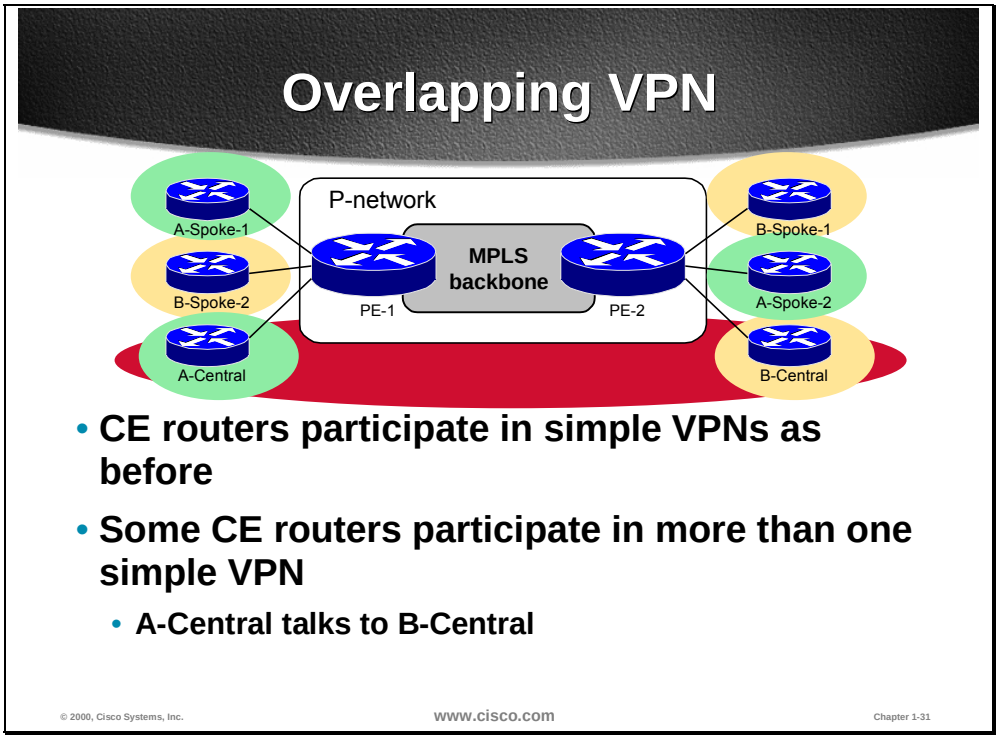

In the case where two VPN customers want to share some information, they may decide to interconnect their central sites. To achieve this, we can create a third VPN that partially overlaps with the customer VPNs and connects only the central sites of the customer VPNs. The result is that the central sites can talk to each other, but not to other sites belonging to the other customer's VPN.

The addresses used in the central sites, however, have to be unique in both VPNs. The other option is to use dual NAT with registered address to be imported and exported between the two central sites.

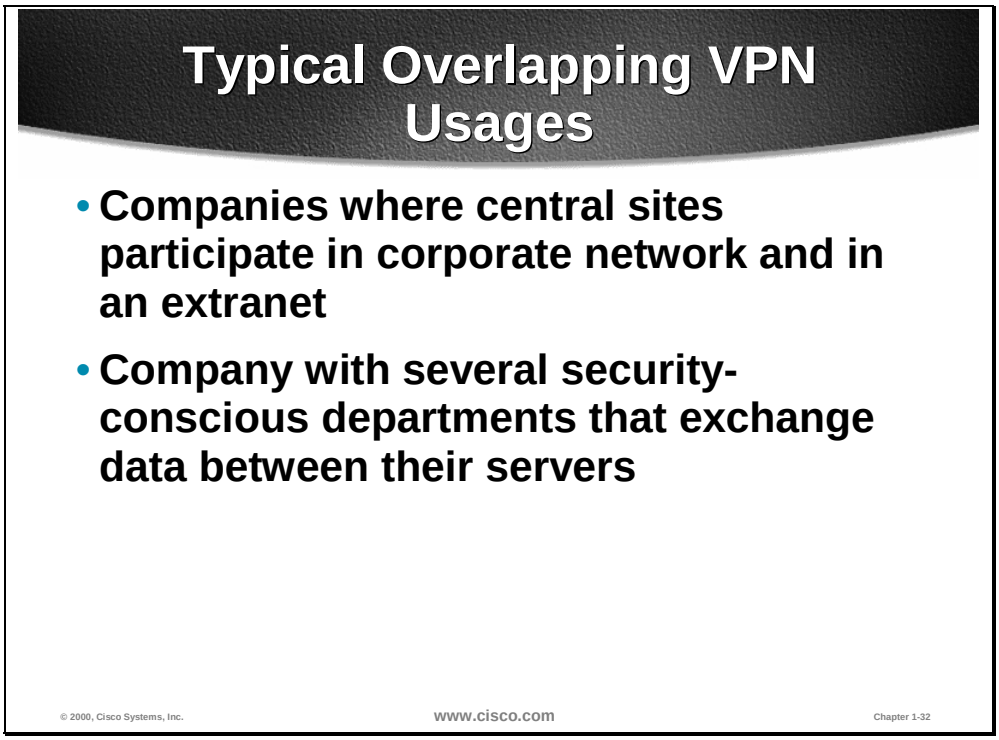

There are two typical usages for overlapping VPNs:

- Companies that use MPLS VPN to implement both intranet and extranet services. In this scenario each company participating in the extranet VPN would probably deploy a security mechanism on its CE routers to prevent other companies participating in the VPN from gaining access to other sites in the customer VPN.
- Some security-conscious companies might decide to deploy limited visibility between different departments in the same organization because of security reasons. Overlapping VPNs might be used as a solution in this case.

Note Security issues might force an enterprise network to be migrated to MPLS VPN even if it's not using MPLS VPN services from a service provider.

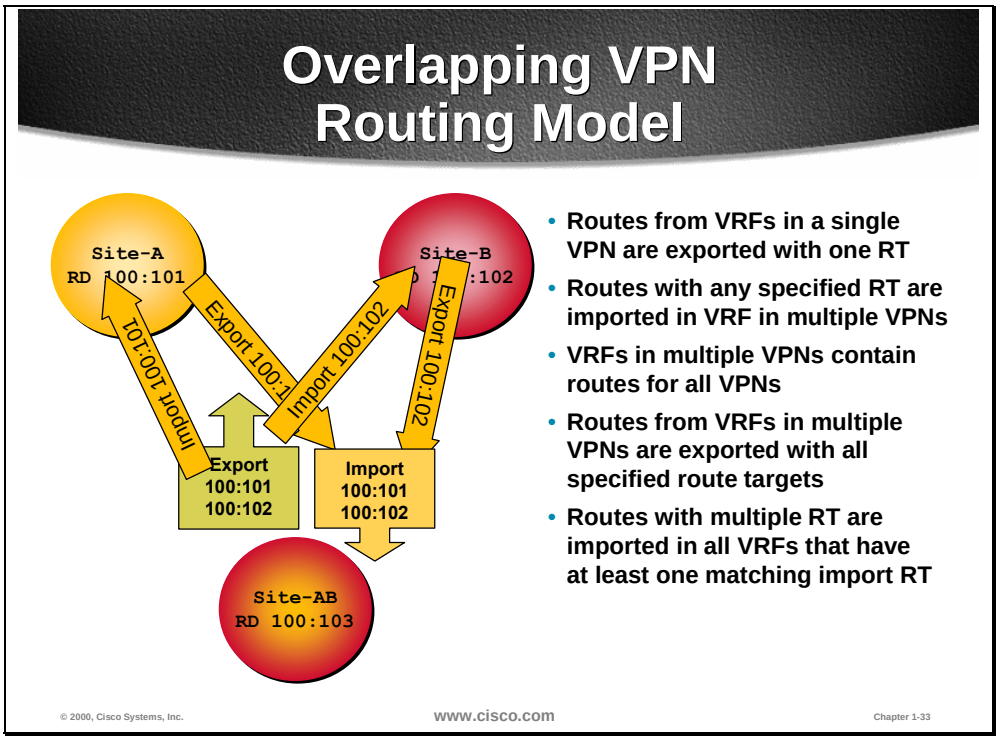

The slide above shows how to implement overlapping VPNs:

- Each VPNs has its own route target (100:101, 100:102) that the sites participating in the VPN import and export
- The sites that participate in more than one VPN import routes with route targets from any VPN in which they participate and export routes with route targets for all the VPNs in which they participate

Site A (participating only in VPN-A):

- Exports all networks with route target 100:101
- Imports all networks that carry route target 100:101 (VPN A)

Site B (participating only in VPN-B):

- Exports all networks with route target 100:102
- Imports all networks that carry route target 100:102 (VPN B)

Site AB (which participates in VPN-A and VPN-B):

- Exports all networks with route targets 100:101 **and** 100:102
- Imports all networks that carry route target 100:101 (VPN A) **or** 100:102 (VPN B)

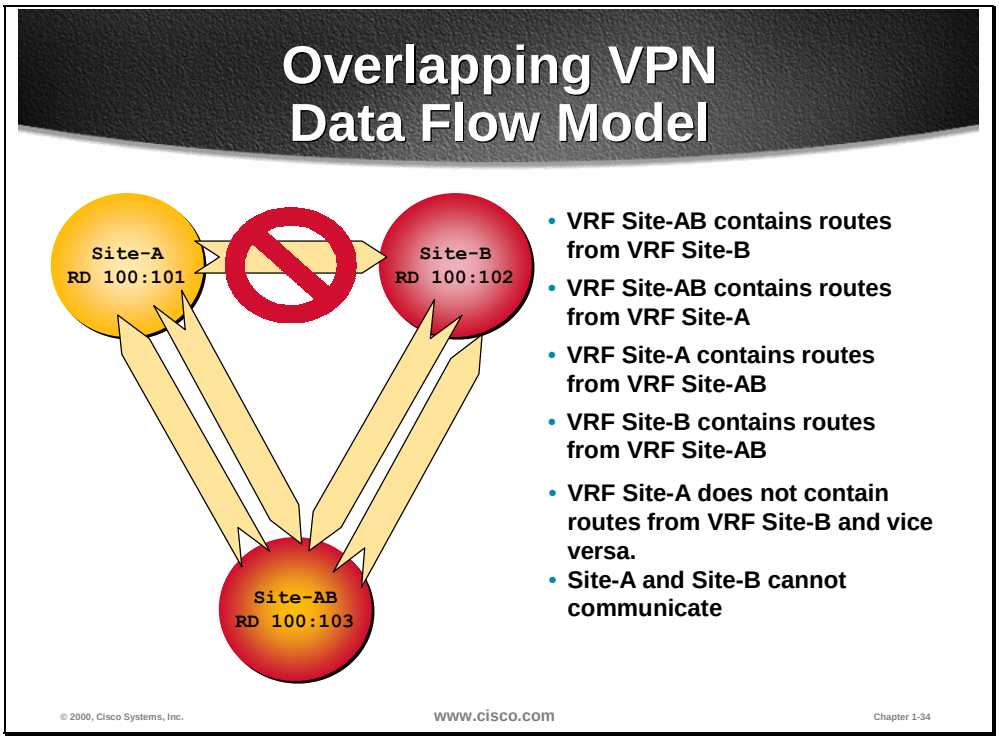

Because sites belonging to different VPNs don't share any routing information, they can't talk to each other.

Note If one of the sites participating in more than one VPN is propagating a default route to other sites, it can attract traffic from those sites and start acting like a transit site between VPNs, enabling sites that were not supposed to communicate to establish two-way communication.

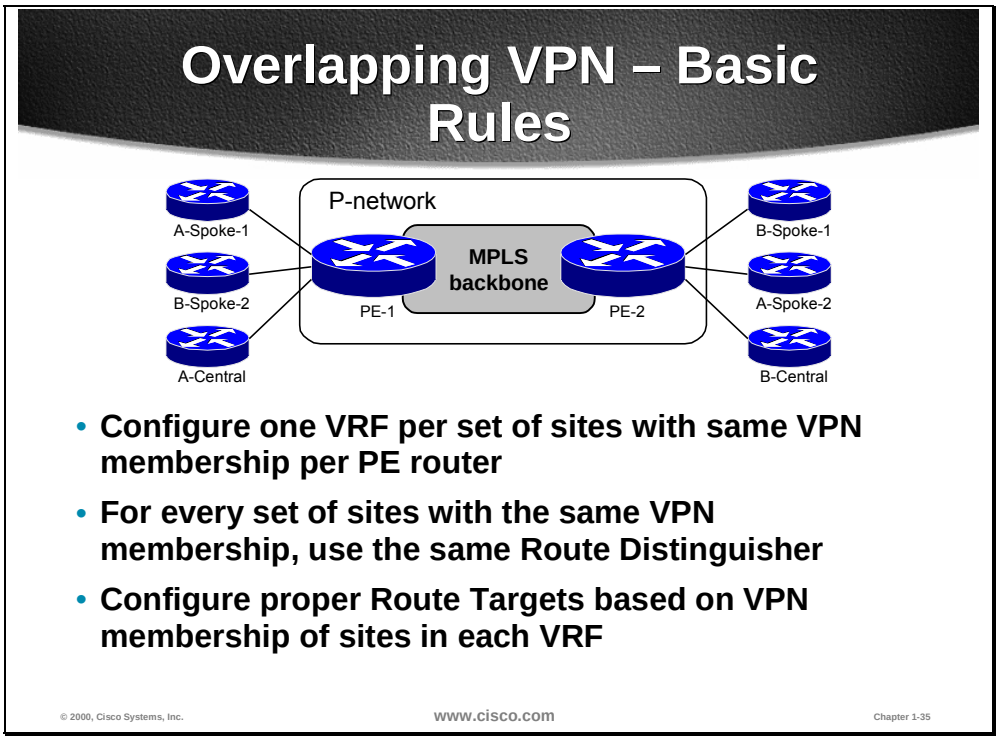

In this example we have four types of sites with different VPN memberships. This means we have to have at least four VRFs:

- A-Spoke-1 and A-Spoke-2 are members of VPN A only (we need two VRFs because they are not connected to the same PE router; we can, however, use the same route distinguisher)
- B-Spoke-1 and B-Spoke-2 are members of VPN B only (we need two VRFs because they are not connected to the same PE router; we can, however, use the same route distinguisher)
- A-Central is a member of VPN-A and VPN-AB (we need an additional route distinguisher)
- B-Central is a member of VPN-B and VPN-AB (we cannot use the same route distinguisher as for A-Central because B-Central has different routing requirements than A-Central)

The following table shows a route target and route distinguisher numbering scheme for PE-1:

VRF	Route Distinguisher	Import Route Target	Export Route Target
VPN_A	123:750	123:750	123:750
VPN_B	123:760	123:760	123:760
VPN_A_Central	123:751	123:750 123:1001	123:750 123:1001

The following table shows a route target and route distinguisher numbering scheme for PE-2:

VRF	Route Distinguisher	Import Route Target	Export Route Target
VPN_A	123:750	123:750	123:750
VPN_B	123:760	123:760	123:760
VPN_B_Central	123:761	123:760 123:1001	123:760 123:1001

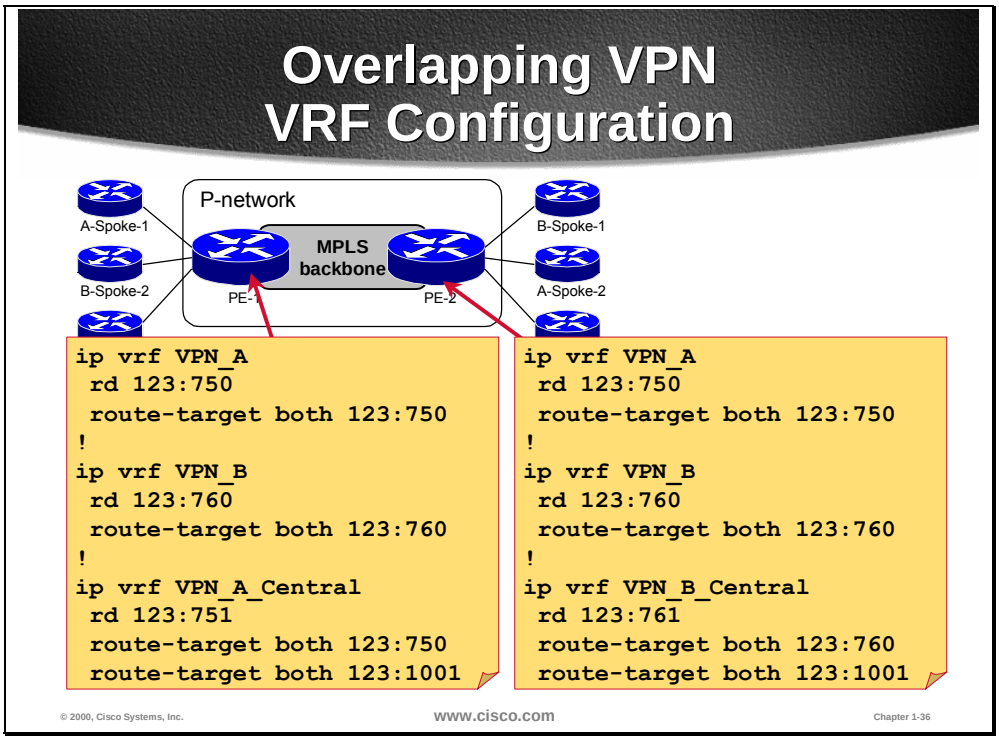

The IOS configuration for PE-1 and PE-2 reflects the route target and route distinguisher numbering scheme from the previous page. The example shows only VRF configuration and does not show VPN routing or MP-BGP routing between the PE-routers.

Overlapping VPN – Routing

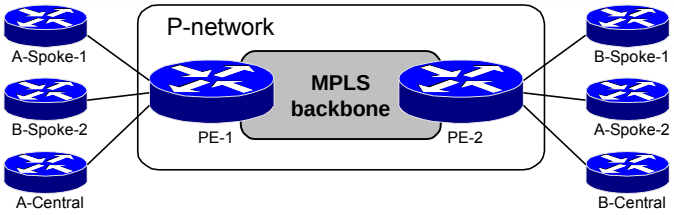

- **Use the same rules for routing protocol selection and routing design as with simple VPN**

Routing within individual VPNs does not change. You can still use any supported PE-CE routing protocol based on the design criteria already covered in the “Simple VPN with Optimal Intra-VPN Routing” section.

Summary

Overlapping VPNs are usually used when two separate VPNs want to interconnect parts of their networks. A third VPN is created within the MPLS VPN network that contains sites from both VPNs. A new Route Target extended community is used for networks originating in the sites that are also in the new VPN. This action may also require a new VRF resulting in a new Route Distinguisher.

Networks originating in these sites are exported with two Route Target extended communities – one for its VPN and one for the overlapping VPN.

Review Questions

Answer the following questions

- What are the typical usages for overlapping Virtual Private Networks?
- What are the connectivity requirements for overlapping VPNs?
- What is the expected data flow within overlapping VPNs?
- How many VRFs do you need at most to implement three partially overlapping VPNs? How many route distinguishers? How many route targets?
- How would you select a routing protocol to use in an overlapping VPN solution?

Central Services VPN Solutions

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe the situations when the central services VPN topology is appropriate
- Describe the routing model of the central services VPN topology
- Describe the data flow of the central services VPN topology
- Design and configure Central Services VPN
- Explain the implications of combining Central Services VPN with simple customer VPN

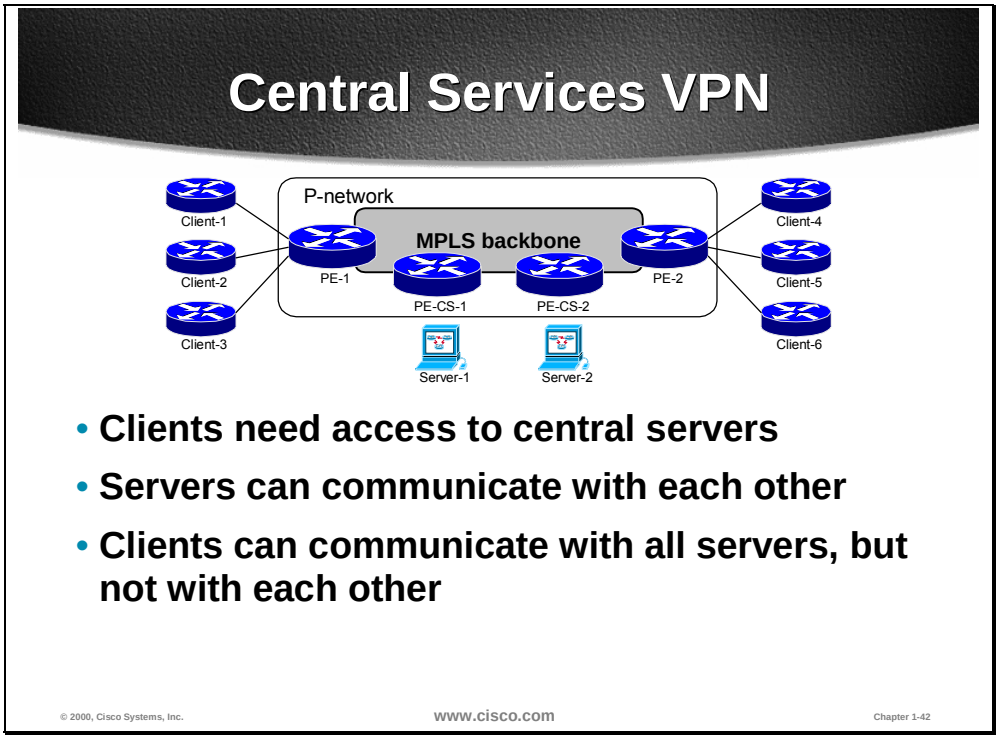

Central Services VPN is a topology where:

- Some sites (**server** sites) can communicate with all other sites
- All the other sites (**client** sites) can communicate only with the server sites

This topology can be used in the following situations:

- The service provider offers services to all his customers by allowing them access to a common VPN
- Two (or more) companies want to exchange information by sharing a common set of servers
- A security conscious company separates its departments and only allows them access to common servers

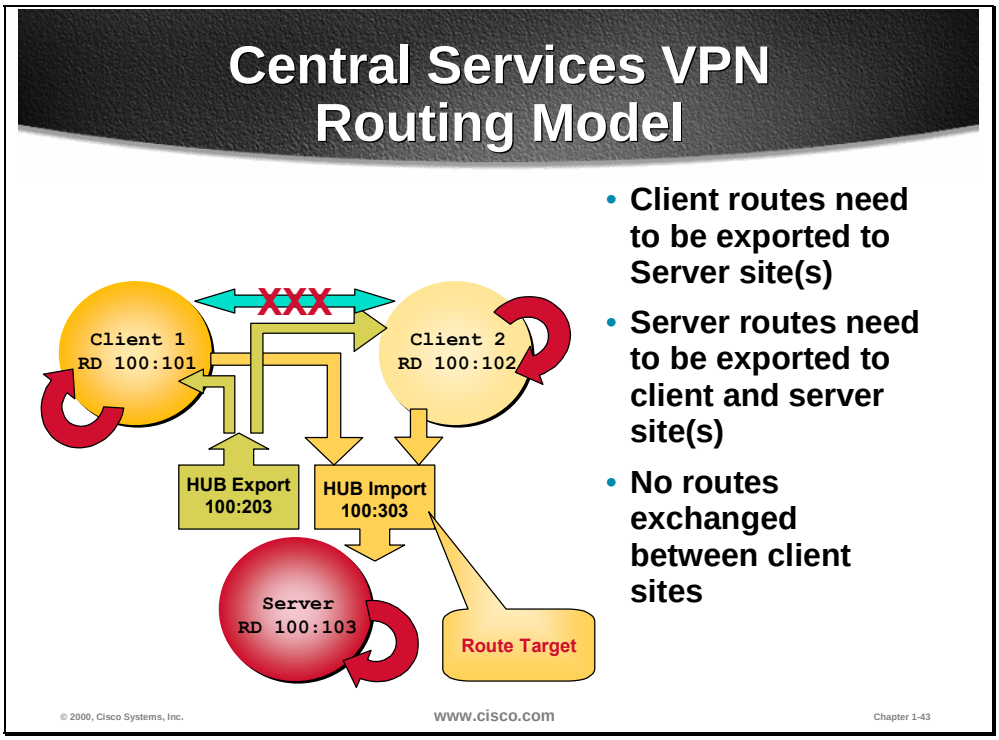

The example above describes the MPLS VPN routing model used to implement central services VPN:

- Client 1 and Client 2 have their own route target (100:101, 100:102) that they import and export; they also export networks with route target 100:303 and import networks with route target 100:203

Note The client-specific route targets are introduced to comply with the implementation requirements of IOS release 12.0T, where each VRF has to have at least one of its export route targets configured as its import route target.

- The central site imports and exports networks with the route target of its VPN, but it also imports networks with route target 100:303 and exports networks with route target 100:203

Client 1:

- Exports all networks with route target 100:101 **and** 100:303
- Imports all networks that carry route target 100:101 **or** 100:203

Client 2:

- Exports all networks with route target 100:102 **and** 100:303
- Imports all networks that carry route target 100:102 **or** 100:203

Central site:

- Exports all networks with route targets 100:203
- Imports all networks that carry route target 100:303

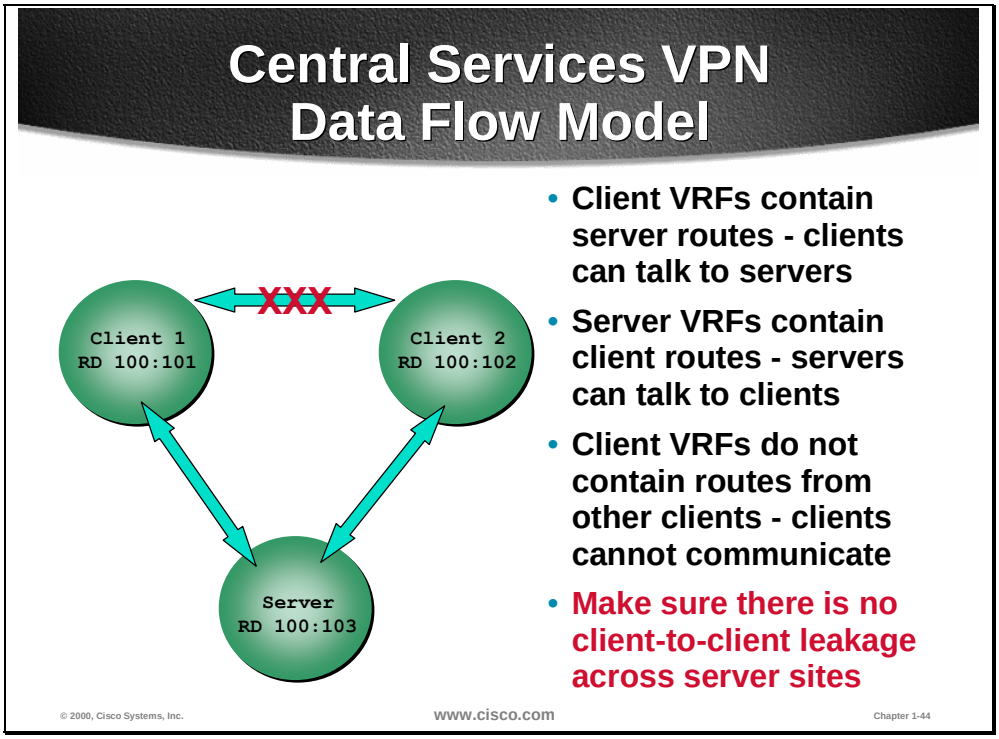

In the central services VPN topology, the client VRF contains only routes from the client site and routes from the server sites – the client sites thus cannot communicate with other client sites.

A server VRF in this topology contains routes from the site(s) attached to the VRF as well as routes from all other client and server sites. Hosts in server sites can therefore communicate with hosts in all other sites.

Note If the central site is propagating a default route to other sites, it can result in client sites seeing each other through the CE in the central site.

Central Services VPN Basic Rules

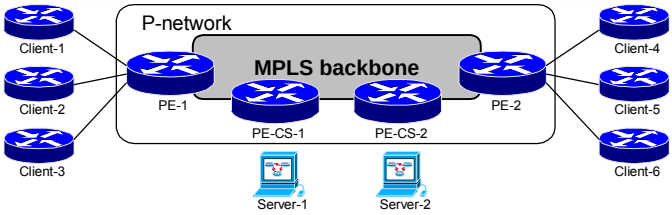

- **Configure a separate VRF per client site**
- **Configure one VRF per server site(s) per PE router**
- **Configure a unique route distinguisher on each client site**
- **Configure an import/export route target with same value as RD for each client site**

In this example we have six client sites and two server sites:

- We need a separate VRF for each client
- We need one VRF per PE router connecting a server site

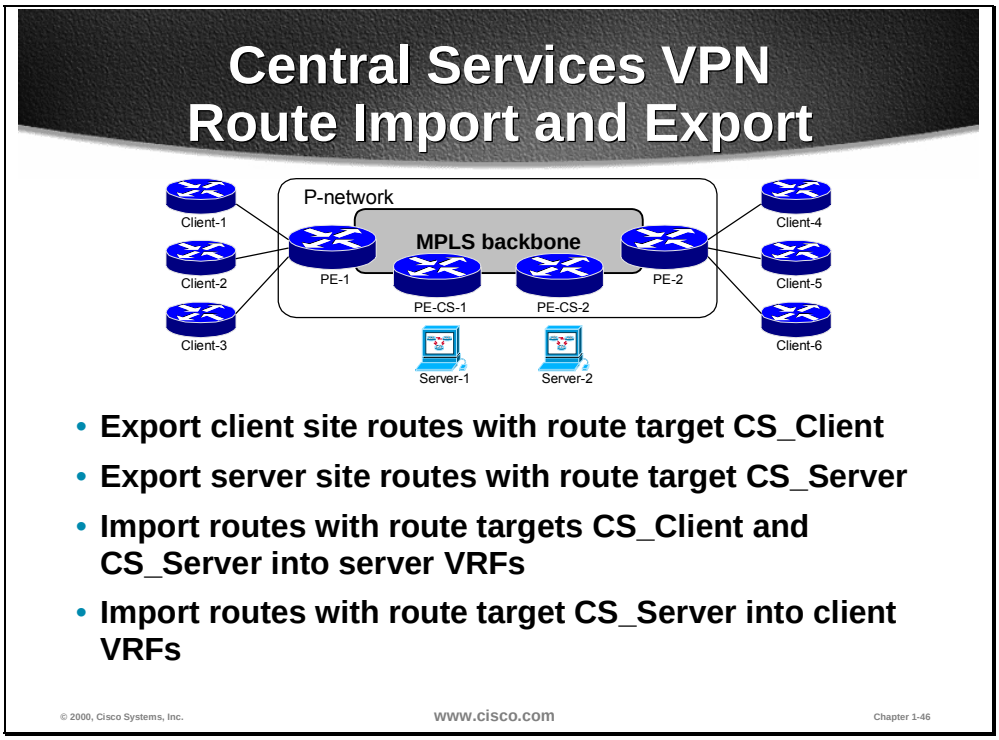

The following table shows a route target and route distinguisher numbering scheme for PE-1:

VRF	Route Distinguisher	Import Route Target	Export Route Target
Client_1	123:101	123:101 123:203	123:101 123:303
Client_2	123:102	123:102 123:203	123:102 123:303

The following table shows a route target and route distinguisher numbering scheme for PE-2:

VRF	Route Distinguisher	Import Route Target	Export Route Target
Client_4	123:111	123:111 123:203	123:111 123:303
Client_5	123:112	123:112 123:203	123:112 123:303

The following table shows a route target and route distinguisher numbering scheme for PE-CS-1:

VRF	Route Distinguisher	Import Route Target	Export Route Target
Server	123:103	123:103 123:303	123:103 123:203

The following table shows a route target and route distinguisher numbering scheme for PE-CS-2:

VRF	Route Distinguisher	Import Route Target	Export Route Target
Server	123:103	123:103 123:303	123:103 123:203

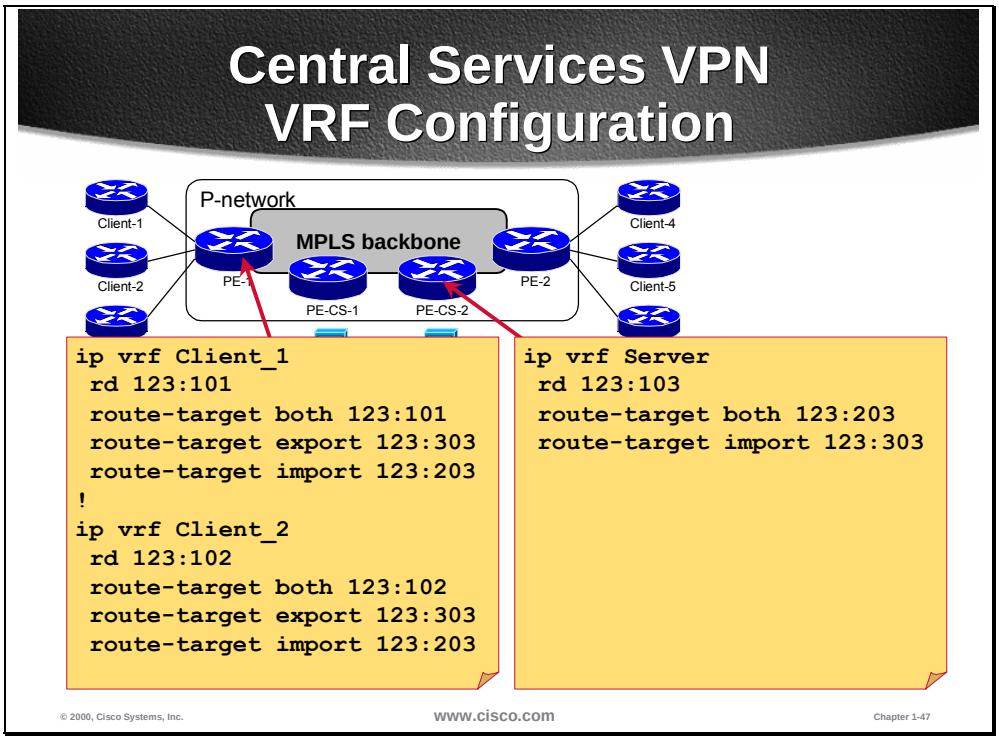

The example above shows a fraction of the configuration according to the route distinguisher and route-target numbering scheme shown on the previous page.

Central Services VPN + Simple VPN

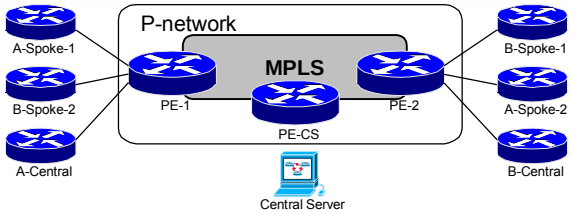

- **Customers run simple VPN (all A-sites in A-VPN, all B-sites in B-VPN)**
- **Only some sites of the customer VPN (A-Central and B-Central) need access to central servers**
- **Combination of rules from overlapping VPN and Central Services VPN**

In this design some of the customer sites need access to the central server and all other sites just need optimal intra-VPN access (as described in the “Simple VPN with Optimal Intra-VPN Routing” section). The design is consequently a mixture of Simple VPN topology and Central Services VPN topology.

Central Services + Simple VPN VRF Creation

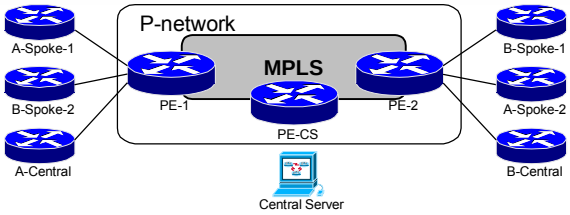

- **For all sites participating in a simple VPN, configure a separate VRF per set of sites participating in the same VPNs per PE router**
- **For sites that are only clients of central servers, create a VRF per site**
- **Create one VRF for central servers per PE router**

We need one VRF per VPN for sites that have access to other sites in the customer VPN, but no access to the Central Services VPN, one VRF per VPN for sites that have access to Central Services VPN, and one VRF for the Central Services VPN (on another PE router in our example).

Central Services + Simple VPN Route Distinguishers

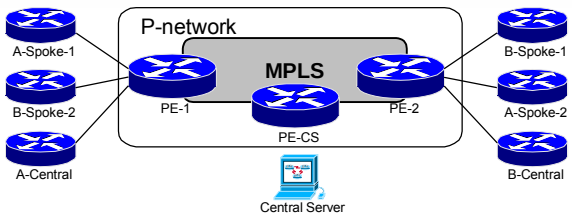

- **Configure a unique RD for every set of VRFs with unique membership requirements**
  - **A-Spoke-1 and A-Spoke-2 can share the same RD**
  - **A-Central needs a unique RD**
- **Configure a unique RD for each site that is only a client of central servers**
- **Configure one RD for all central server VRFs**

For this design we need two route distinguishers per VPN:

- One route distinguisher for simple VPN sites; same value should also be used for import and export route target
- One route distinguisher for VPN sites that also have access to the Central Services VPN
- One route distinguisher for the Central Services VPN

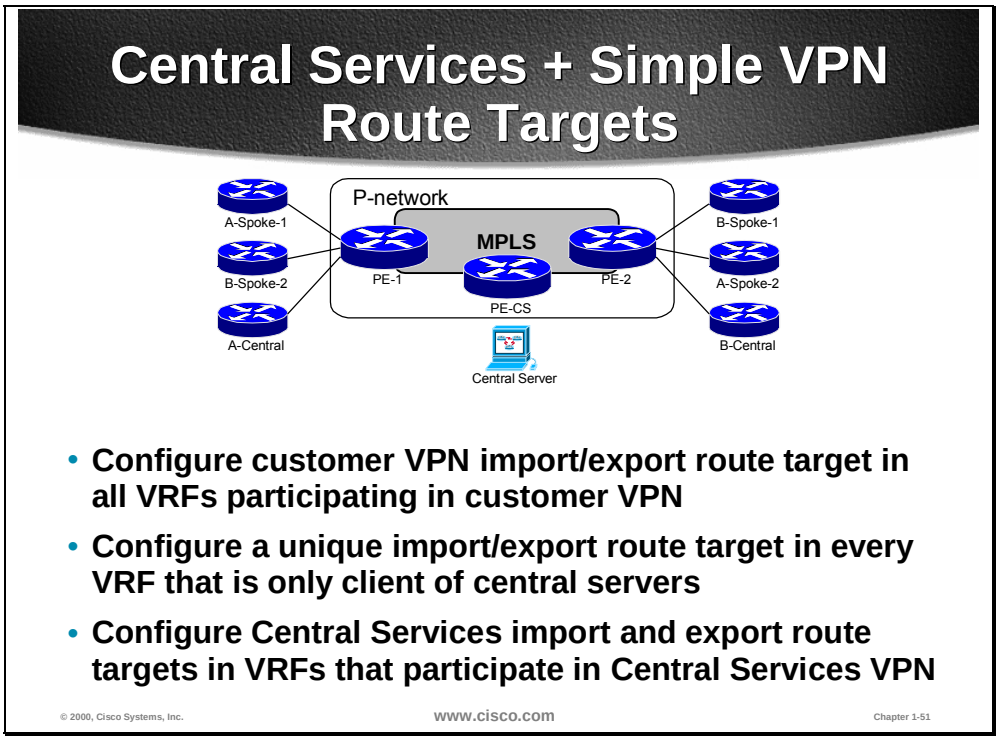

The following table shows a route target and route distinguisher numbering scheme for PE-1:

VRF	Route Distinguisher	Import Route Target	Export Route Target
VPN_A	123:750	123:750	123:750
VPN_B	123:760	123:760	123:760
VPN_A_Central	123:751	123:750 123:101	123:750 123:100

The following table shows a route target and route distinguisher numbering scheme for PE-2:

VRF	Route Distinguisher	Import Route Target	Export Route Target
VPN_A	123:750	123:750	123:750
VPN_B	123:760	123:760	123:760
VPN_B_Central	123:751	123:760 123:101	123:760 123:100

The following table shows a route target and route distinguisher numbering scheme for PE-CS:

VRF	Route Distinguisher	Import Route Target	Export Route Target
Server	123:101	123:101 123:100	123:101

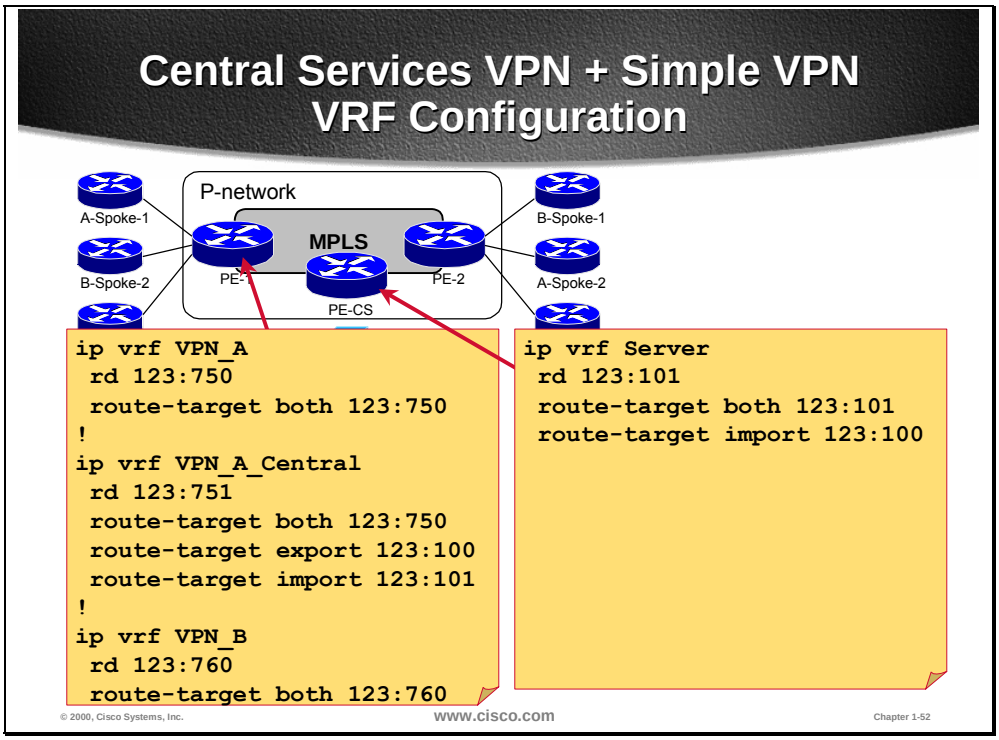

The example above shows a fraction of the configuration according to the route distinguisher and route-target numbering scheme shown on the previous page.

Summary

Central services VPN is used when more VPNs need to share a common set of servers. These servers reside in the Central Services VPN and all other VPNs have access to this VPN. Those VPNs, however, are not able to see one another.

The Central Services VPN is implemented using two Route Target extended communities where one is used to import networks into the VPN and the other to export networks. The client sites do the opposite. Two Route Target extended communities are needed to prevent client sites from exchanging routing information.

Review Questions

Answer the following questions

- What are the typical usages for central services VPN topology?
- What is the connectivity model for central services VPN topology?
- How do you implement central services VPN topology?
- How many route targets do you need for a central services VPN solution with two server sites and 50 client sites? How many route distinguishers?
- How do you combine central services VPN topology with simple VPN topology?

Hub-and-Spoke VPN Solutions

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe the situations when the hub-and-spoke VPN topology is appropriate
- Describe the routing model of the hub-and-spoke VPN topology
- Describe the data flow of the hub-and-spoke VPN topology
- Select the optimal PE-CE routing protocol in hub-and-spoke topology
- Explain the implications of using BGP as the PE-CE routing protocol at the hub site

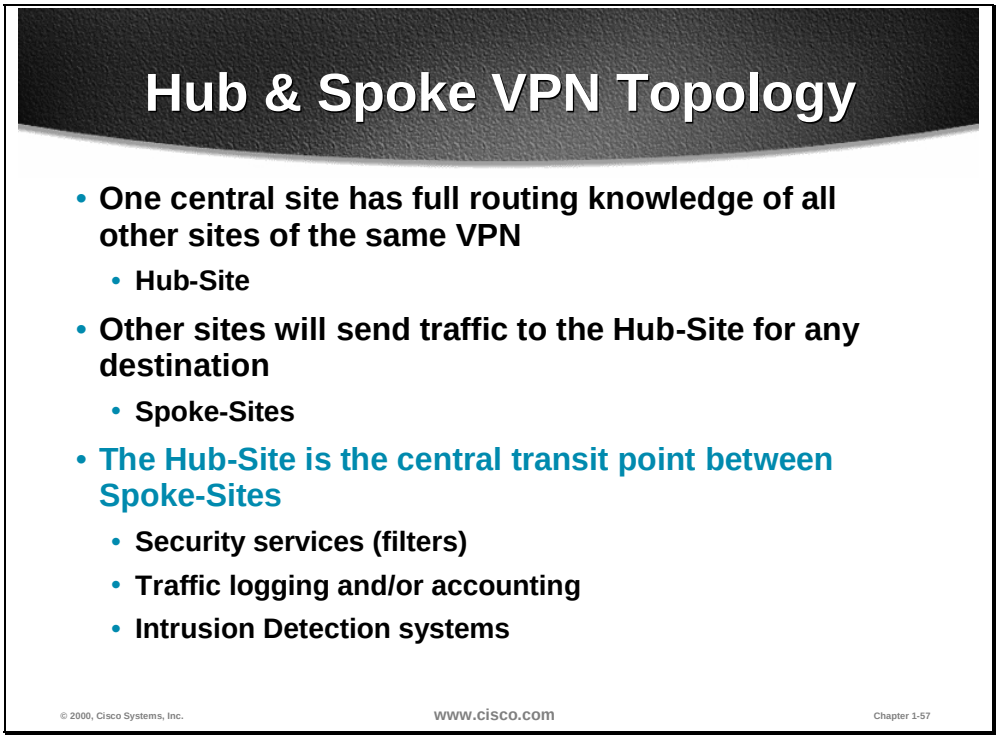

MPLS VPN core networks seamlessly create a full mesh between sites belonging to the same VPN with optimal routing within the service provider core network. The traffic exchanged between individual customer sites never flows through other customer routers.

Some customers would still like to retain the centralized control that was inherent with overlay VPN hub-and-spoke topology where all the traffic was exchanged through the central site (or sites). For customers that need a hub & spoke topology implemented over a MPLS VPN backbone, a special design is needed to force the VPN core to forward all packets to the central site. To achieve that we have to prevent spoke sites from exchanging routing information. They can only exchange routing information with the hub site.

Note Hub-and-spoke VPN topology defeats the scalability of MPLS VPN and the optimum inter-site routing provided by MPLS VPN.

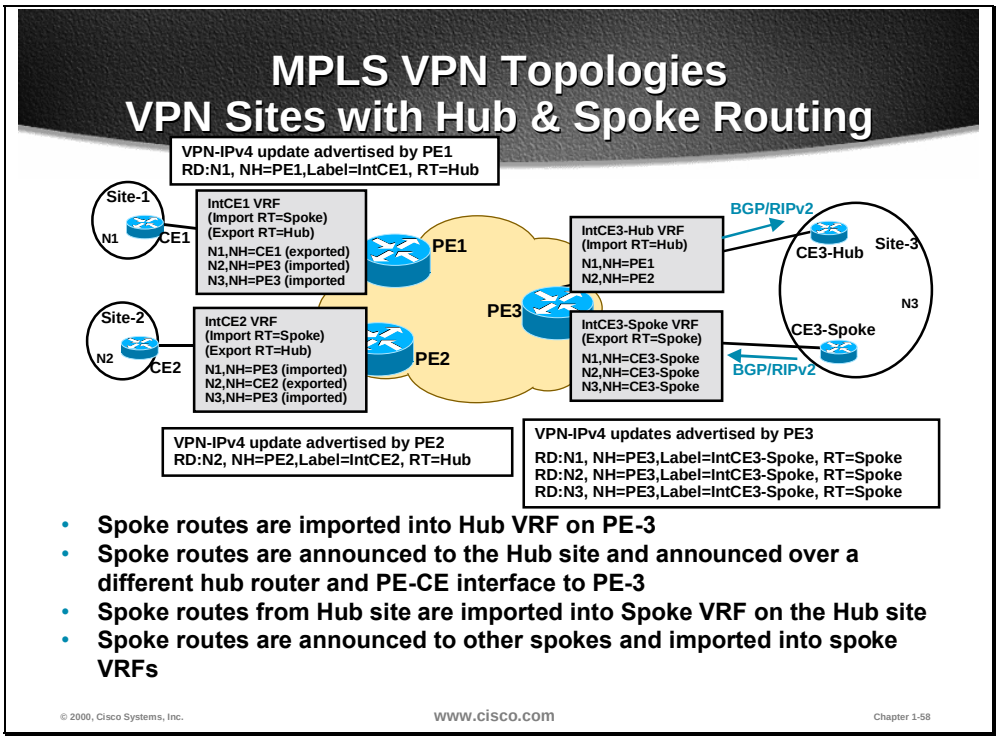

To make sure that the packets are forwarded to the hub site we need two interfaces in two separate VRFs for the hub site. One is used to receive packets from the hub site and propagate them to spoke sites; the other is used to collect packets from spoke sites and send them to the hub site.

In the picture above the spoke sites are propagating routing information to the PE, which is marking them with the route target “Hub” that can only be imported into the VRF Hub and sent to the central site.

The central site then propagates the same information out through the second interface, and the PE to which the central site is attached is marking this information with the route-target “spoke” that can be imported in all other spoke VRFs.

Note We need a separate VRF for every spoke site even if they are connected to the same PE to prevent spoke sites from exchanging routing information directly.

Note Using OSPF between PE and CE routers is not recommended because you have one VRF per site, which requires a separate OSPF process per interface. This would increase the CPU load and memory requirements on the PE router as well as very possibly exceed the limit of 32 routing protocols per router.

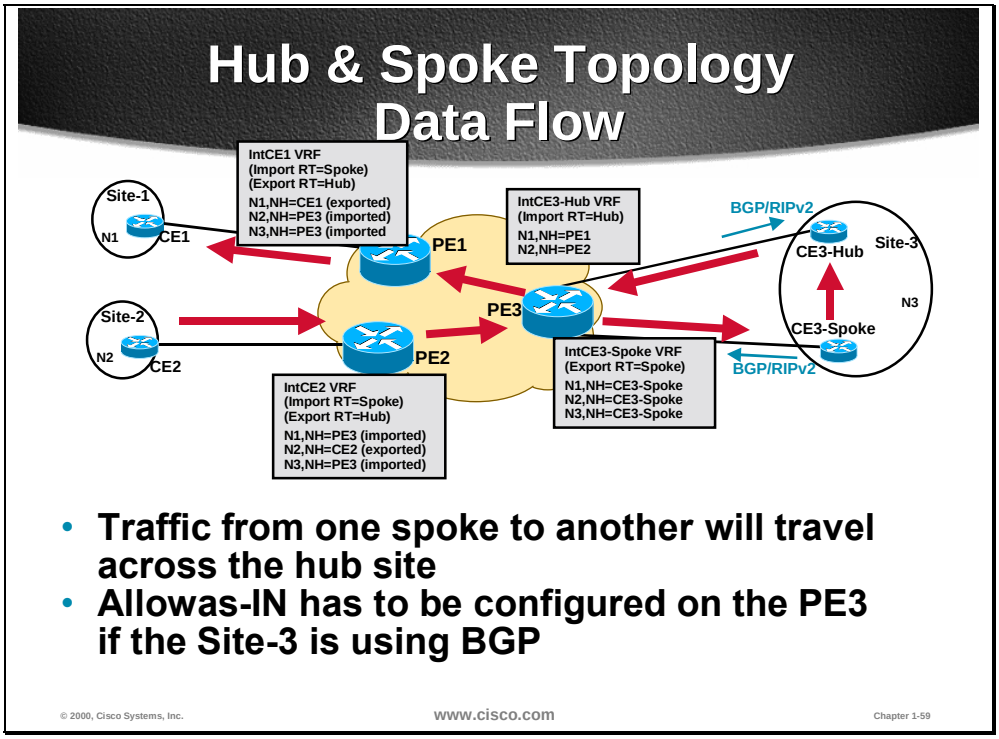

The routing table for all spoke VRFs is pointing to the central site for all destinations in the VPN, which results in packets flowing through the hub site.

Note This design causes asymmetric routing. For example: a packet going from Spoke-1 to Spoke-2 will enter the hub site through interface Hub-1 and exit through interface Hub-2; the returning packet will also enter the hub site through interface Hub-1 and exit through interface Hub-2. This side effect may prevent the customer from deploying stateful filters or similar mechanisms that also check the direction of packets.

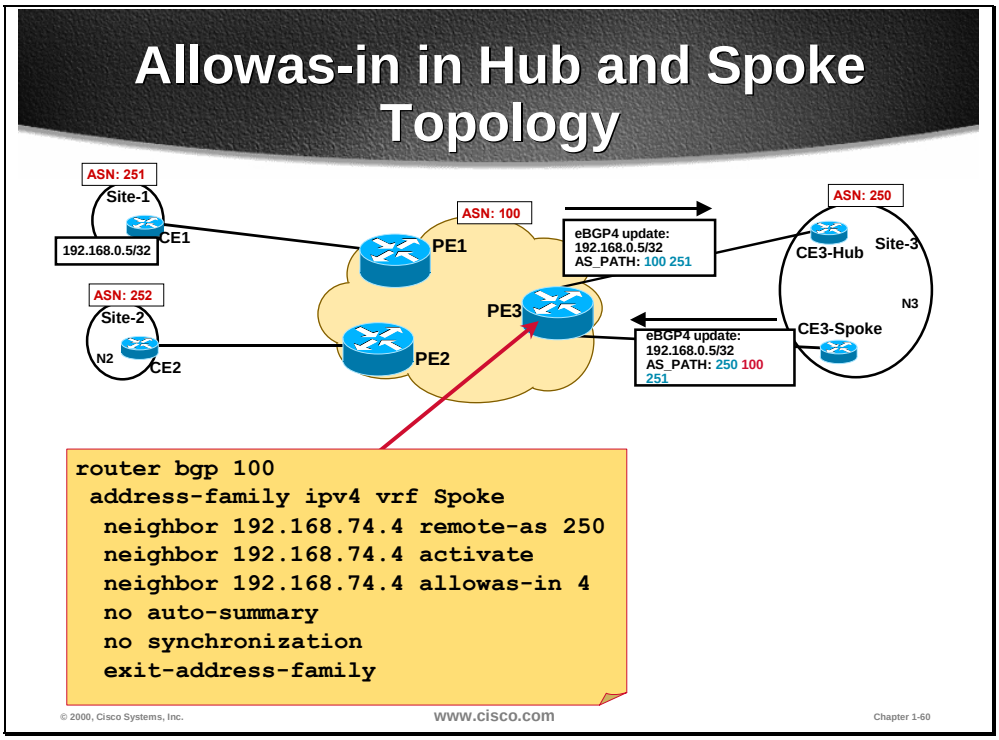

In the case where the customer is using BGP, the service provider does not accept updates coming from the hub site if they have previously been sent to the hub site through the other BGP session. This is because it regards it as a routing loop (it finds its own AS number in the AS path).

To overcome this problem there is an option where we specify the maximum number of occurrences of our own AS numbers in the AS path. In the example above the service provider will accept all updates as long as they don't contain its AS number in the AS path more than four times.

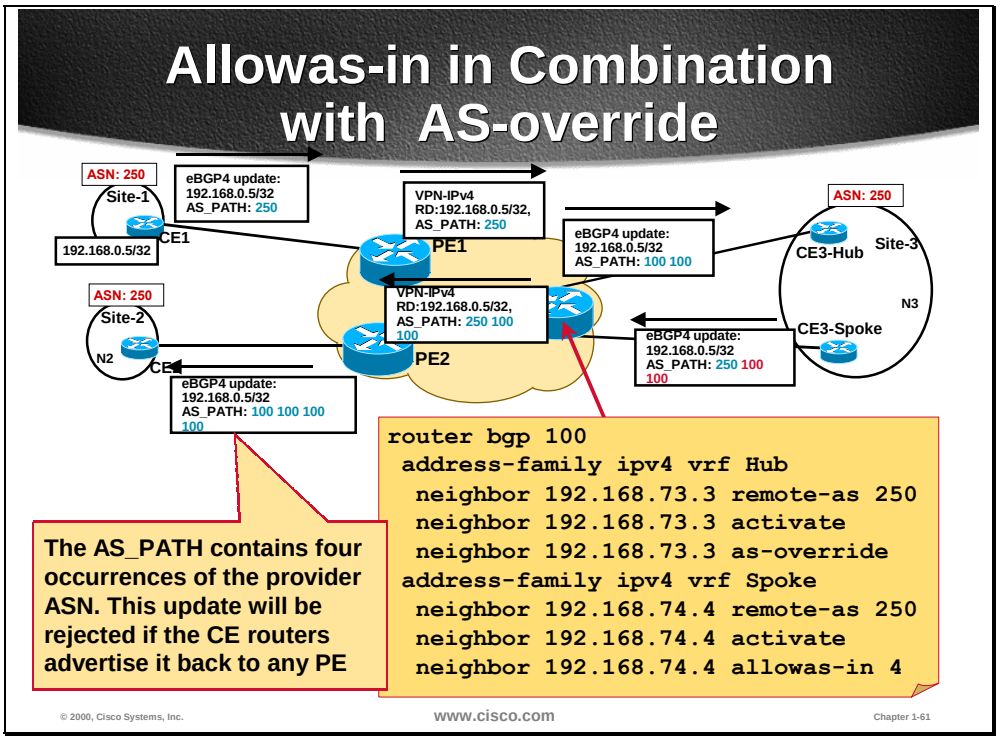

If the customer is using one single AS number for all the sites, a similar problem will occur in the hub site: a spoke site originating a route will prepend AS 250 and send it to the service provider; the service provider will prepend its own AS number and send the update to the hub site; the hub site will then ignore the update because it contains its AS number. To overcome this, the service provider can overwrite the customers AS number with its own. To accomplish this we use the **as-override** feature on all BGP sessions between PE and CE routers (for hub and spoke sites).

Summary

One of the major benefits of a MPLS VPN solution is that it provides a full mesh between the CE sites with optimal routing in the core. There is no longer any need for a central site. If, however, there is a need for all the packets to go through a central site, a special design is needed for the VPN.

To force the packets to go through the central site, we need two links for the central site (also called hub site) – one is importing other CE's routes, the other is exporting them. To prevent spoke CE sites from exchanging routing information the PE routers to which the spoke sites are attached have to export routing updates from CE sites with a different Route Target extended community from the one they import. This also requires each CE site to have its own VRF.

Routing is no longer optimal because all packets between spoke CE sites have to traverse the core network twice.

If BGP is used between the PE router and the hub site's CE router we have to use Allow-AS feature to prevent returning routing updates from being ignored on the PE router.

Review Questions

Answer the following questions

- When would you deploy hub-and-spoke VPN topology?
- What is the main difference between central services VPN topology and hub-and-spoke VPN topology?
- What is the main difference between simple VPN topology and hub-and-spoke VPN topology?
- Describe the routing information flow in hub-and-spoke topology.
- Describe the packet forwarding in hub-and-spoke topology.
- How many PE-CE links do you need at the spoke sites?
- How many PE-CE links do you need at the hub sites?
- Do you need two CE routers at the hub site?
- Do you need two PE routers to connect the hub site?
- Which routing protocol would you use between the P-network and the hub site?
- Which BGP features are necessary to support BGP as the routing protocol at the hub site?
- Which BGP features are necessary to support BGP as the routing protocol at the spoke site if all sites use the same AS number?

Managed CE-Router Service

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe the requirements of managed CE-router service
- Design a VPN topology based on central-services VPN topology, which solves the managed CE-router service requirements
- Implement the Managed CE-Router Service VPN solution

Managed CE Routers

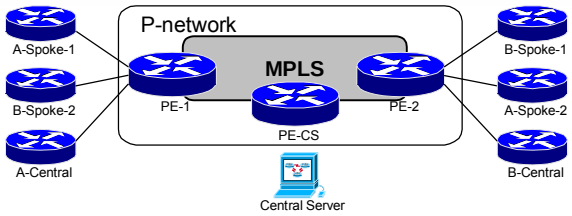

- **Central Server (NMS) needs access to loopback addresses of all CE routers**
- **Very similar to Central Services + Simple VPN**
  - **All CE routers participate in Central Services VPN**
  - **Only loopback addresses of CE routers need to be exported into Central Services VPN**

If the service provider is managing the customer routers, it is convenient to have a central point that has access to all CE routers, but not to the other destinations at customer sites. This requirement is usually implemented by deploying a separate VPN for management purposes. This VPN has to see all the loopback interfaces of all the CE routers. All CE routers have to see the Management VPN. The design is very similar to the Central Services VPN with the only difference being that we only mark loopback addresses to be imported into the Management VPN.

Note The topology described in this section is sometimes also referred to as **gray** management VPN implementation as all CE routers are accessed through a single link between the NMS CE router and the network core. An alternate solution (**rainbow** management VPN), where the NMS CE router has separate connections to each managed CE router is usually used in combination with overlay VPNs (for example, Frame Relay networks).

Managed CE Routers VRF Creation and RD

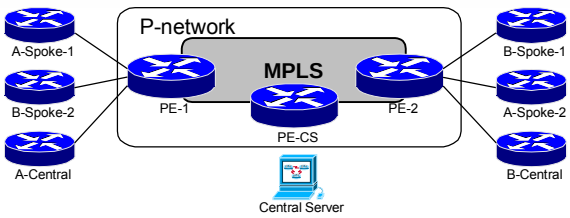

- **Create one VRF per customer VPN per PE router**
- **Assign the same RD to each customer VRF**
- **Create NMS VRF on the PE-CS router**
- **Assign a unique RD to NMS VRF**

The VRF and route distinguisher design is the same as with Central Services VPNs. The only difference between this topology and Central Services VPN topology combined with Simple VPN topology is the route-target marking process during route export.

Managed CE Routers Route Targets

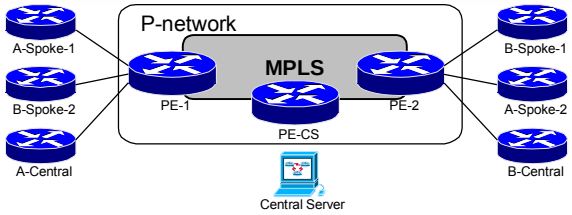

- **Configure per-customer import/export route target in all customer VRFs**
- **Configure NMS import/export route target in NMS VRF**
- **Import routes with NMS RT into customer VRF**
- **Export loopback addresses from customer VRF with RT NMS_Client**
- **Import routes with RT NMS_Client into NMS VRF**

The following table shows a route target and route distinguisher numbering scheme for PE-1:

VRF	Route Distinguisher	Import Route Target	Export Route Target
VPN_A	123:750	123:750 123:101	123:750 123:100 (lo0)
VPN_B	123:760	123:760 123:101	123:760 123:100 (lo0)

The following table shows a route target and route distinguisher numbering scheme for PE-CS:

VRF	Route Distinguisher	Import Route Target	Export Route Target
Server	123:101	123:101 123:100	123:101

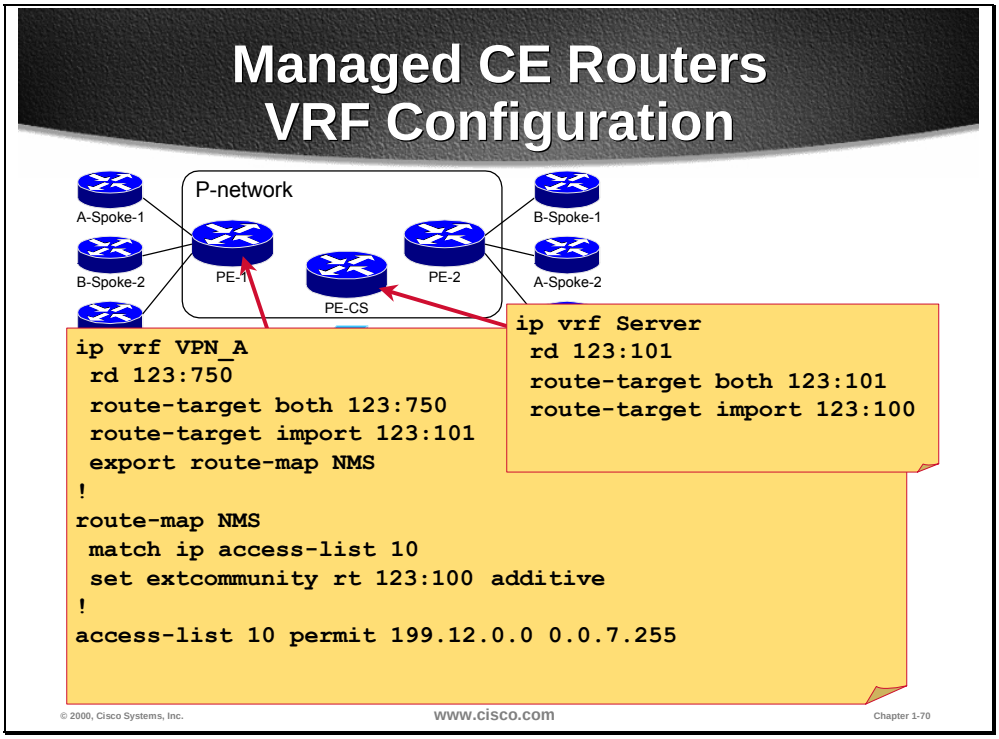

The example above shows a sample configuration for a customer VRF with differentiated route target export for loopback addresses according to the numbering scheme shown on the previous page. Export route-map is used to match on part of the IP address space and attach an additional route-target to the routes within this address space (CE router loopback addresses).

Note The routing protocol between PE and CE routers has to be secured (with distribute-lists or prefix-lists) to prevent customers from announcing routes in the address space dedicated to the network management; otherwise, the customers can gain two-way connectivity to the network management station.

The CE router loopback addresses are then imported into the Server VPN based on the additional route-target attached to them during the export process.

Note This design allows client sites to send packets to the Management VPN regardless of the source address. Special precautions should be taken to protect the Management VPN from potential threats and denial-of-service attacks coming from customer sites.

Summary

A separate Management VPN can be used by the service provider to manage the CE routers of all the VPNs.

A pair of Route Target extended communities is used to accomplish this. One is used to export CE routers' loopback addresses and is imported into the VRF of the Management VPN. The other Route Target is used to export the networks form the Management VRF and import them into all other VRFs.

Review Questions

Answer the following questions

- When would you need managed CE router service?
- How do you implement managed CE router service?
- What's the main difference between managed CE router service and usual central services VPN topology?

Chapter Summary

After completing this chapter, you should be able to perform the following tasks:

- Design and implement simple VPN solutions with optimal intra-VPN routing
- Design and implement various routing protocols within VPNs
- Design and implement central services VPN topologies
- Design and implement hub-and-spoke VPN topologies
- Design and implement VPN topology required for managed router services

6

Internet Access from a VPN

Overview

Integrating Internet Access with an MPLS/VPN solution is one of the most common SP business requirements. This chapter provides a good understanding of underlying design issues, several potential design scenarios and some sample configurations.

This chapter contains the following topics:

- Integrating Internet Access with the MPLS VPN Solution
- Design Options for Integrating Internet Access with MPLS VPN
- Leaking Between VPN and Global Backbone Routing
- Separating Internet Access from VPN Service
- Internet Access Backbone as a Separate VPN

Objectives

Upon completion of this chapter, you will be able to perform the following tasks:

- Explain the requirements for Internet Access from a VPN.
- Describe various design models for integrated Internet Access and their benefits and drawbacks.
- Design and implement an MPLS VPN solutions based on these design models.
- Design and implement a Wholesale Internet Access solution.

Integrating Internet Access with the MPLS VPN Solution

Objectives

- Upon completion of this section, you will be able to explain the requirements for combining Internet Access with VPN services.

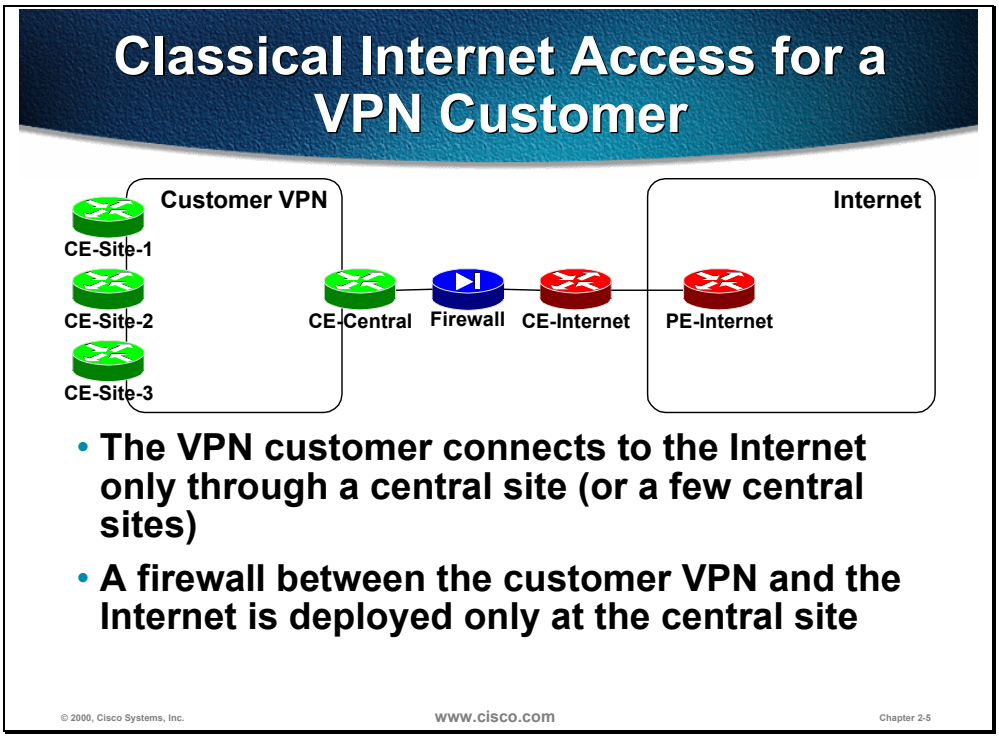

Classical Internet access is implemented through a (usually central) firewall that connects the customer's network to the Internet in a secure fashion. The customer's private network (or Virtual Private Network if the customer is using a VPN service) and the Internet are connected only through the firewall.

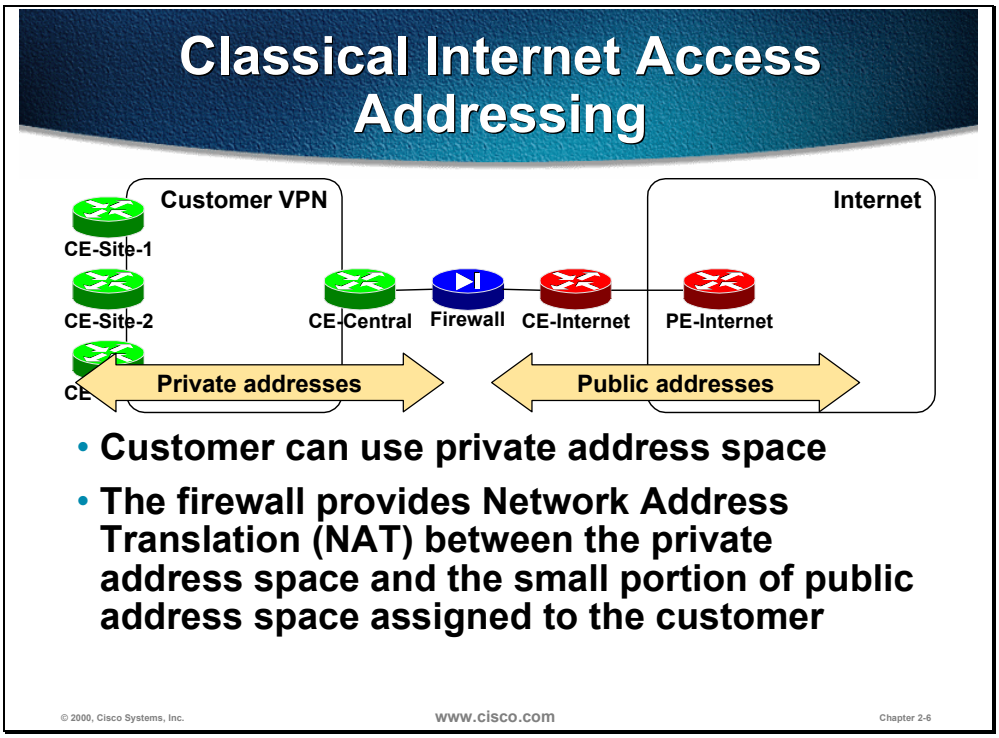

Addressing requirements of this type of connection are very simple:

- The customer is assigned a small block of public address space used by the firewall.
- The customer typically uses private addresses inside the customer network.
- The firewall performs Network Address Translation (NAT) between the customer's private addresses and the public addresses assigned to the customer by the Internet Service Provider (ISP). Alternatively, the firewall might perform an application-level proxy function that also isolates private and public IP addresses.

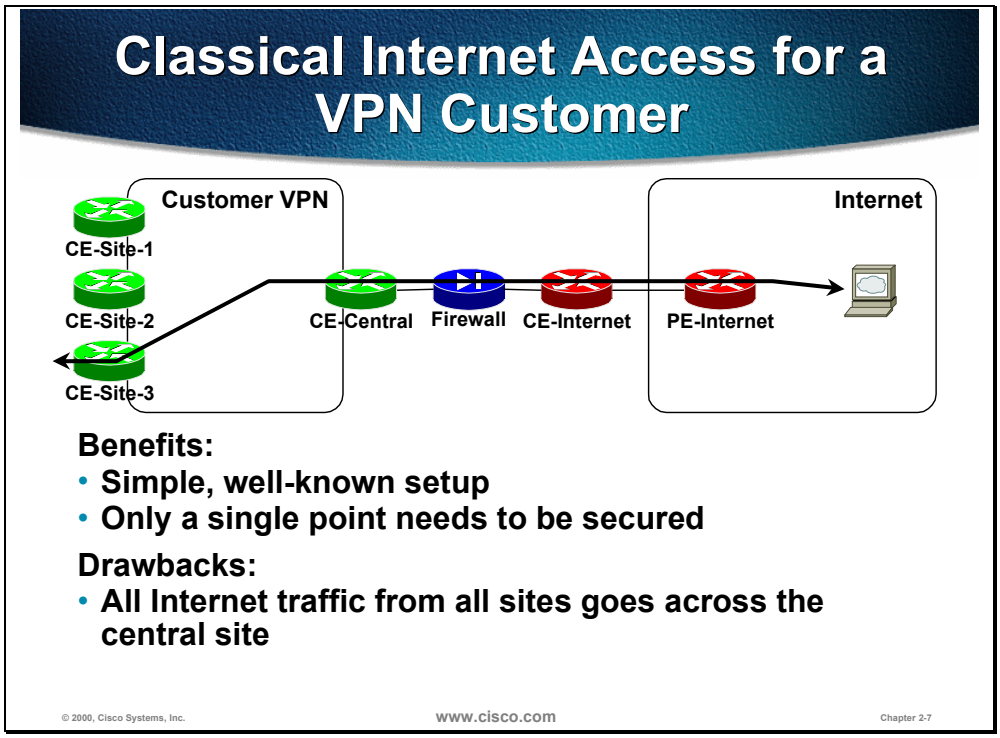

There are a number of benefits associated with this design:

- It is a well-known setup used world-wide for Internet connectivity from a corporate network. Access to expertise needed to implement such a setup is thus simple and straightforward.
- There is only one interconnection point between the secure customer network and the Internet. Security of the Internet access only has to be managed at this central point.

The major drawback of this design is the traffic flow – all traffic from the customer network to the Internet has to pass through the central firewall. While this might not be a drawback for smaller customers, it can be a severe limitation for large organizations with many users, especially when geographically separated.

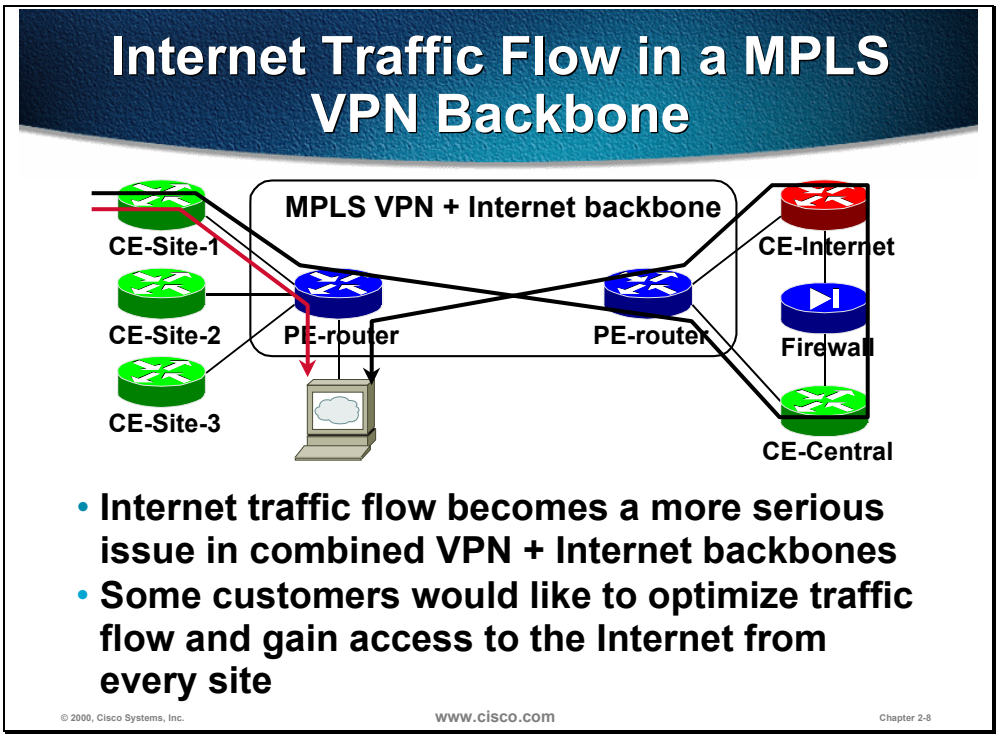

The traffic flow issue becomes even more pronounced when the customer VPN (based on, for example, MPLS VPN service) and the Internet traffic share the same Service Provider backbone. In this case, the traffic from a customer site may have to traverse the Service Provider backbone as VPN traffic, and then return into the same backbone by the corporate firewall, ending up at a server very close to the original site.

Based on this analysis, the drawbacks of the central firewall design can be summarized:

- The link between the central site and the provider backbone has to be over-dimensioned, as it has to transport all of the customer's Internet traffic.
- The provider backbone is over-utilized, as the same traffic crosses the backbone twice, first as VPN traffic and then as Internet traffic (or vice versa).
- Response times and quality of service may suffer since the traffic between the customer site and an Internet destination always has to cross the central firewall, even when the Internet destination is very close to the customer site.

These drawbacks have prompted some large users and service providers to consider alternate designs in which every customer site can originate and receive Internet traffic directly.

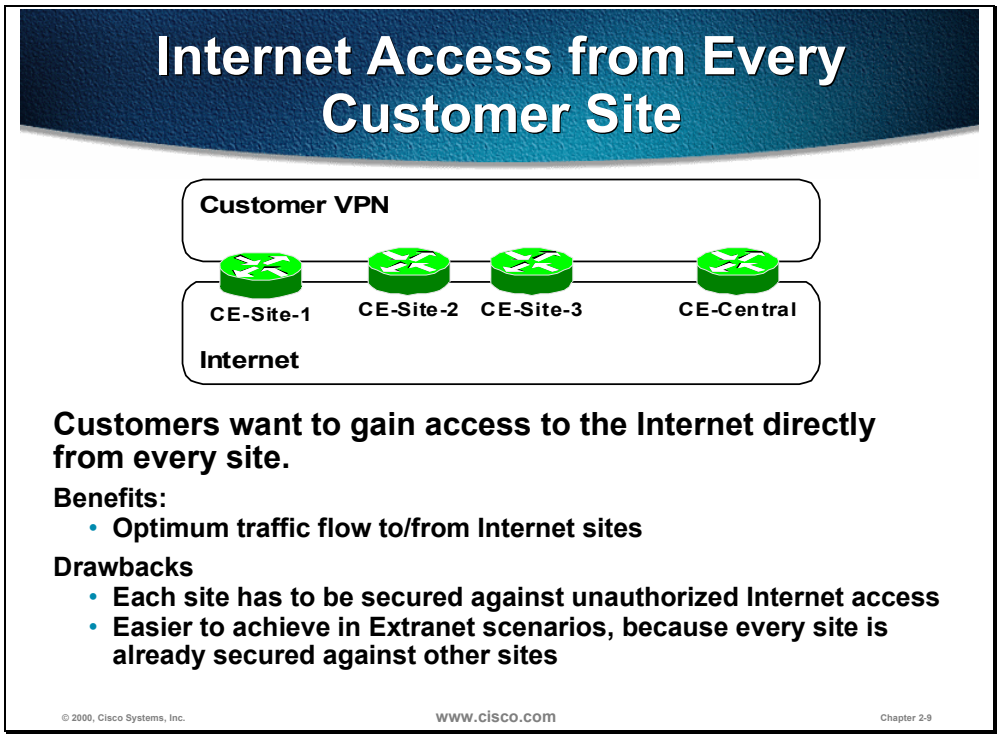

To bypass the limitations of Internet access through a central firewall, some customers are turning toward designs in which each customer site has its own independent Internet access. While this design clearly solves all traffic flow issues, the associated drawback is higher exposure – each site has to be individually secured against unauthorized Internet access. This design is applicable primarily for larger sites (concentrating traffic from close-by smaller sites) or for Extranet VPNs in which each site is already secured against the other sites participating in the Extranet VPN.

Internet Access from Every Site - Addressing

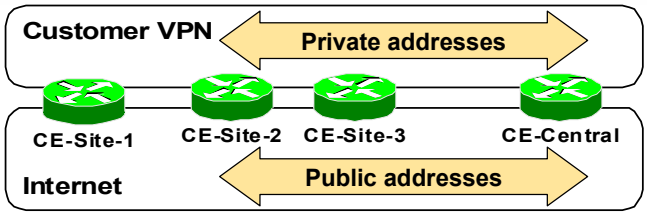

Two addressing options:

- **Every CE router performs NAT functionality – a small part of public address space has to be assigned to each CE router**
- **Customer only uses public IP addresses in the private network - not realistic for many customers**

In order to gain Internet access from every site, each site requires at least some public IP addresses. Two methods can be used to achieve this goal:

- A small part of public address space can be assigned to each customer site. Network Address Translation between the private IP addresses and the public IP addresses needs to be performed at each site.
- If the customer is already using public IP addresses in the VPN, NAT functionality is not needed. Unfortunately, this option is only open to those customers that own large address blocks of public IP addresses.

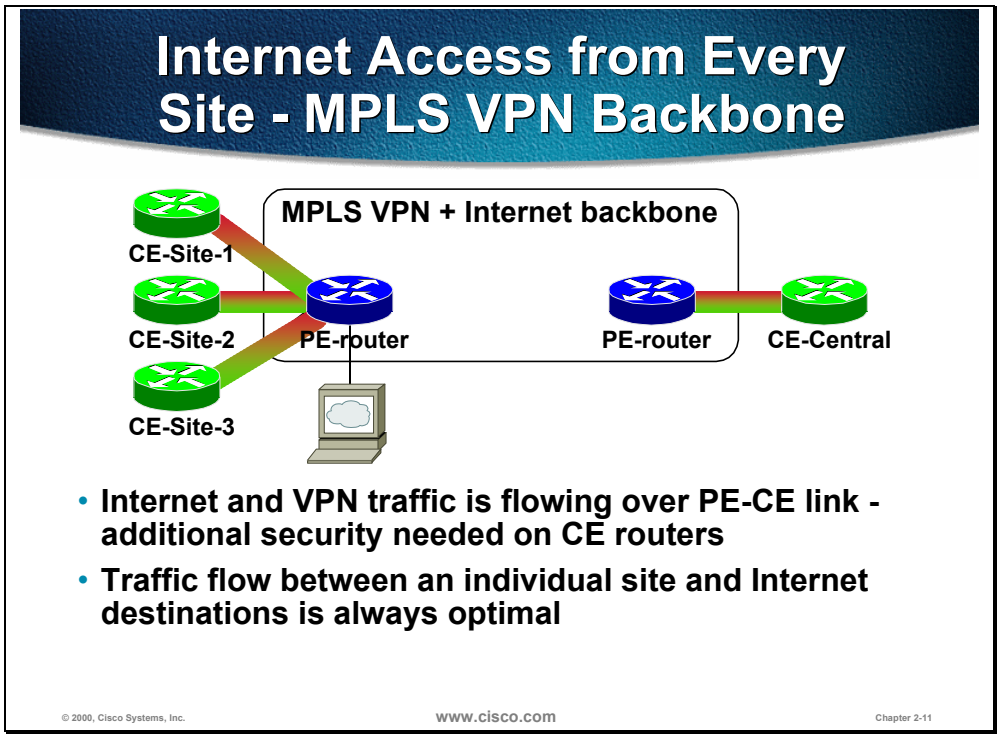

To achieve Internet access from every customer site, each CE router must forward VPN traffic toward other customer sites as well as Internet traffic toward Internet destinations. The two traffic types are usually sent over the same physical link to minimize costs. Switched WAN encapsulation (Frame Relay or ATM) could be used to separate the VPN and Internet traffic onto different virtual circuits or the traffic can share the same logical link as well, resulting in reduced security. On the other hand, the weaker (or more complex) security of this design is offset by optimal traffic flow between every site and Internet destinations.

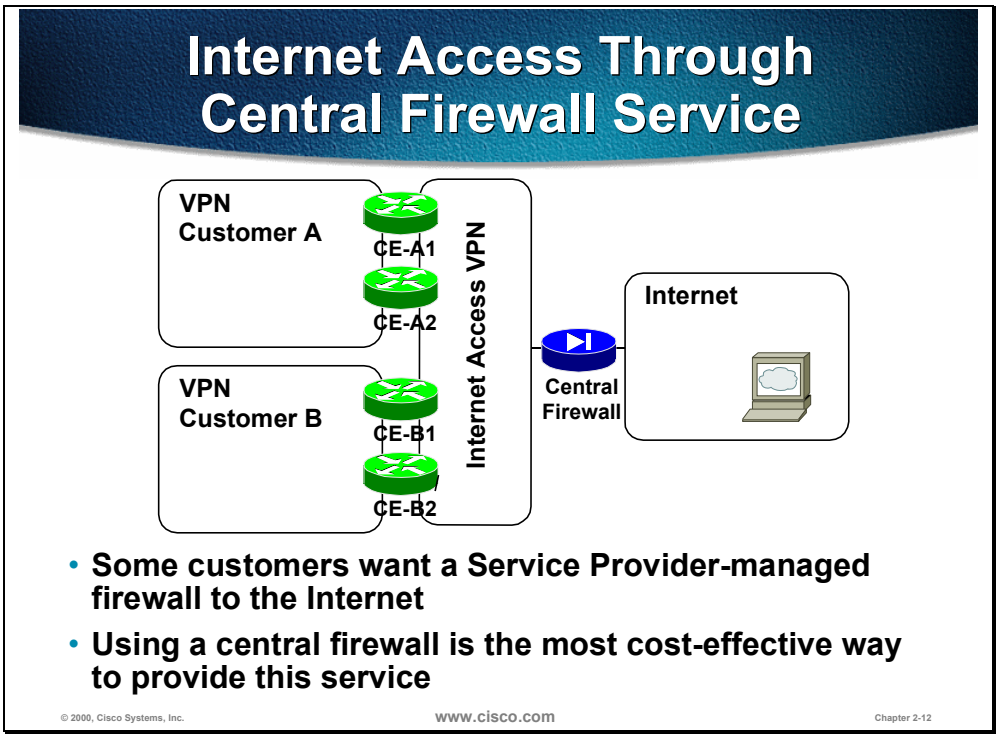

For customers who do not want the complexity of managing their own firewall, a managed firewall service offered by the Service Provider is a welcome relief. These customers typically want the Service Provider to take care of the security issues of their connection to the Internet.

The Service Provider could implement the managed firewall service by deploying a dedicated firewall at each customer site or (for a more cost effective approach) by using a central firewall that provides secure Internet access to all customers.

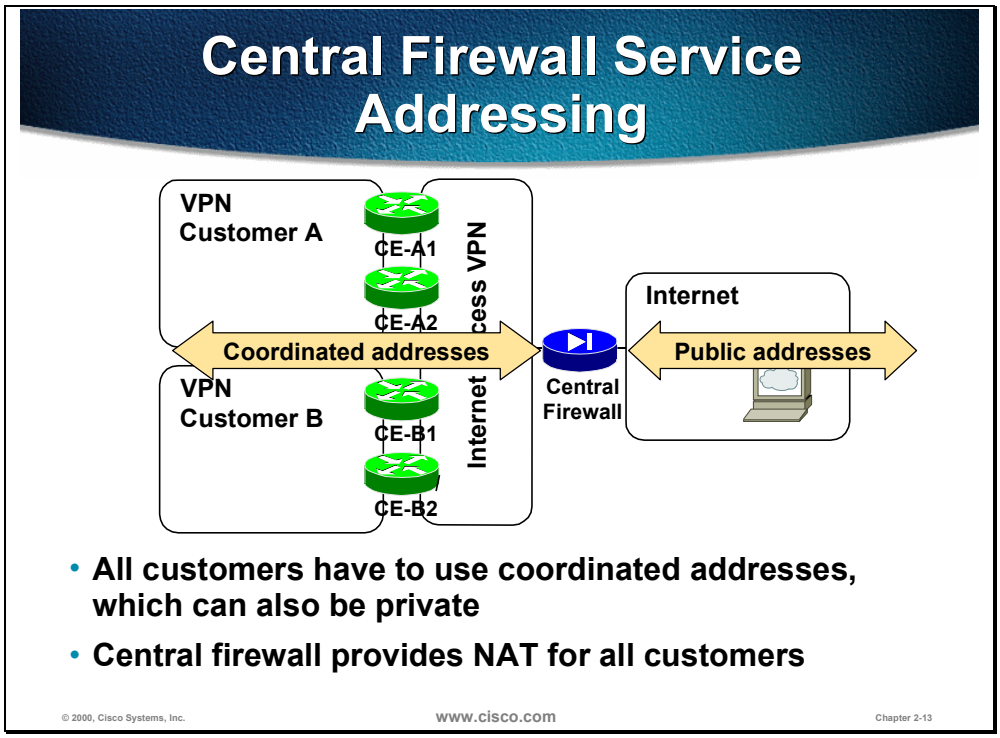

The central firewall, hosted by the Service Provider, has to use public addresses toward the Internet. Private addresses can be used between the central firewall and the individual customers. However, these addresses need to be coordinated between the Service Provider and the customers to prevent routing conflicts and overlapping addresses visible to the central firewall. Customers using central firewall service are thus limited to IP addresses assigned to them by the Service Provider, much in the same way as Internet customers are limited to the public IP addresses assigned by their ISP.

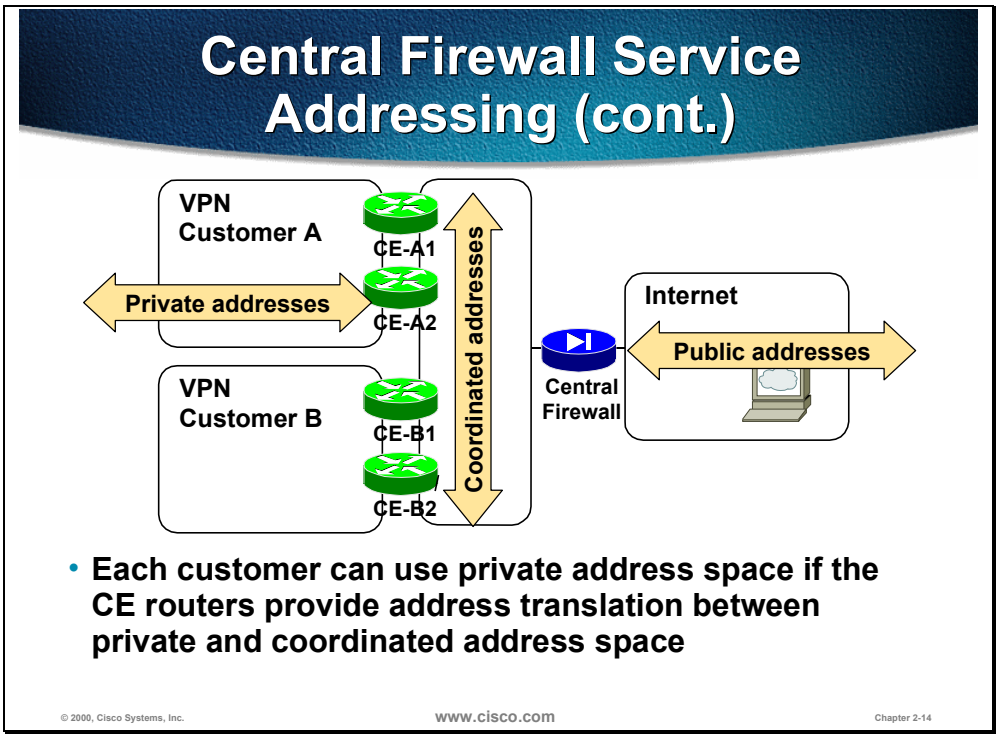

Customers of central firewall service who still want to retain their own private addresses inside their network can use NAT on the CE routers, connecting their private network to the transit network that links customer sites to the central firewall.

Note Service Providers usually use private IP addresses as the address space between the central firewall and the customers. There is always a potential for overlapping addresses between the coordinated address space and the address space of an individual customer. The Customer Edge (CE) device providing NAT functionality therefore has to support address translation between overlapping sets of IP addresses.

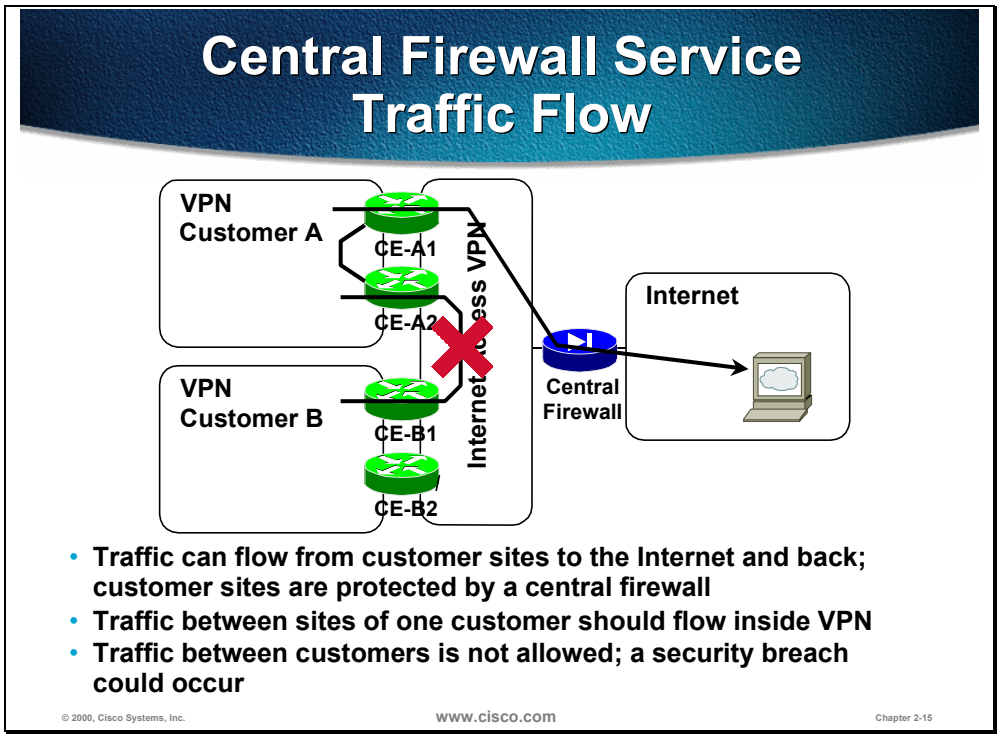

The traffic flow between sites participating in a central firewall service is limited by the security requirements of the service:

- Traffic between the customer sites and the Internet must flow freely, restricted only by the security functions of the central firewall.
- Traffic between sites of an individual customer should never flow across the VPN that links the customer sites with the central firewall. This traffic must flow inside the customer VPN.
- Traffic between customers using the central firewall is not allowed, as the individual customers are not protected from outside access (this is the task of the Service Provider, handled by the central firewall). Inter-customer traffic could lead to potential security problems.

Note The restrictions on inter-customer traffic prevents customers from deploying publicly accessible servers in their networks, as these servers would not be available to other customers of the same service.

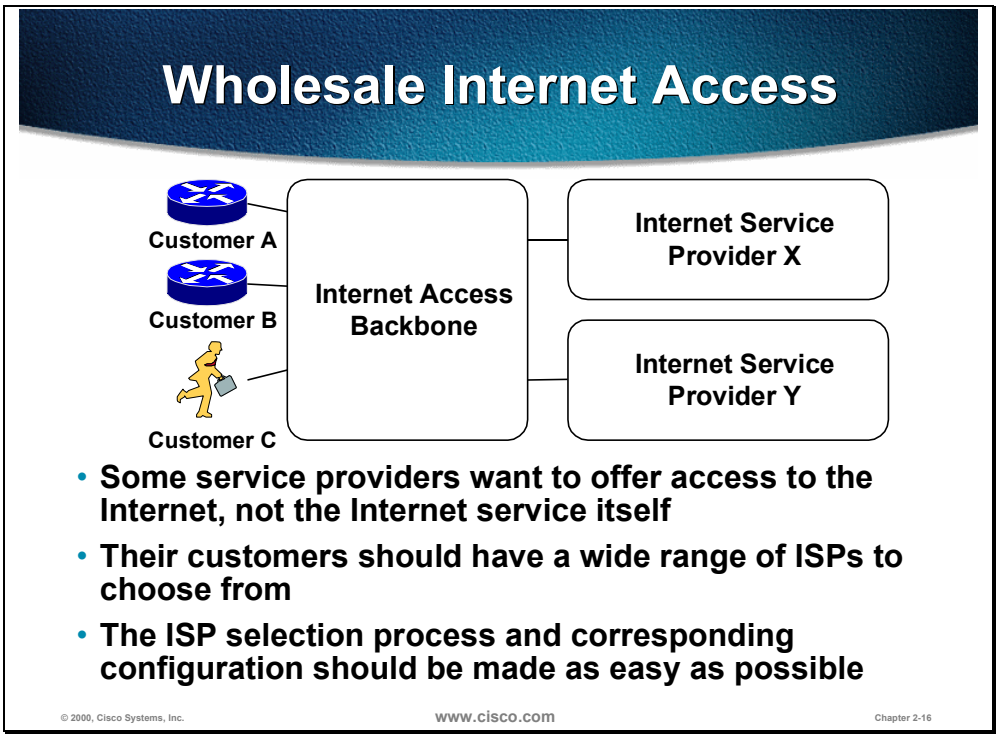

Parallel to Wholesale Dial service (where an ISP uses modem pools of another Service Providers) is the Wholesale Internet Access service, where an ISP uses IP transport infrastructure of another Service Provider to reach the end-users. The business model of this service varies – the end-users might be customers of the Service Provider that owns the transport backbone (for example, a cable operator), who offers Internet access through a large set of ISPs as a value-added service. Alternatively, the Service Provider owning the Internet Access Backbone might act as a true wholesaler, selling transport infrastructure to Internet Service Providers who then charge end-users for the whole package.

When a Service Provider owns the backbone and provides Internet access to customers, the Service Provider usually wants to offer a wide range of upstream ISPs to choose from, in order to satisfy various customers' connectivity and reliability requirements. The selection of upstream ISPs and the corresponding configuration process should therefore be as easy as possible.

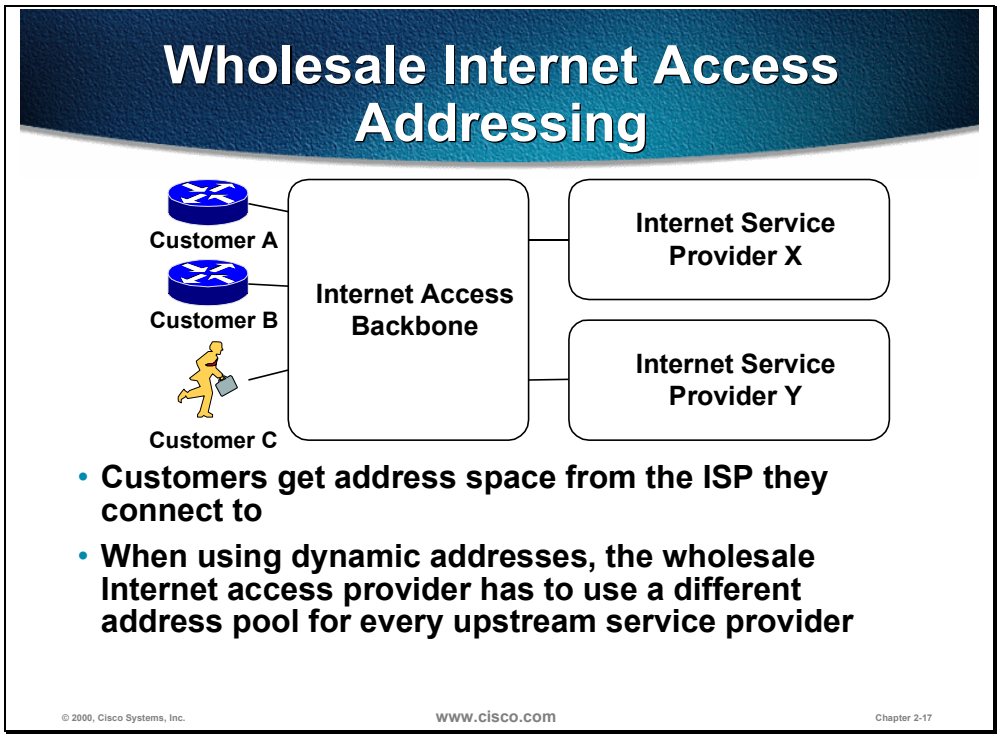

Regardless of the business model used in the Wholesale Internet Access service, the addressing requirements are always the same – the upstream ISP allocates a portion of its address space to the end-users connected to the Internet Access Backbone. The Wholesale Internet Access provider consequently has to use a different address pool for every upstream ISP.

Summary

Traditionally, corporate Internet access was implemented by means of a central firewall located at the customer's central site. Internet traffic from all customer sites would have to pass this central firewall, resulting in tight security.

Some customers find the traffic flow limitations of the central firewall setup too limiting and opt for designs where every site (or major sites) has its own Internet access. The Internet traffic flow of this solution is optimal, but this gain is offset by the increased complexity of managing a firewall at every customer site.

A large number of customers find the task of deploying and managing their own firewall too cumbersome. These customers appreciate managed firewall service from their service provider (or third-party providers). The Internet Service Provider can optimize the costs of providing managed firewall service by deploying a central firewall infrastructure serving many customers.

With the advent of new transport technologies (Cable, DSL, Wireless), the Service Providers deploying these technologies have started looking for new business models that might differentiate them from pure connectivity providers. Wholesale Internet Access with a flexible selection of upstream ISP is one of these innovative options.

Review Questions

- Describe four major customer requirements for Internet access services.
- What are the addressing requirements for classical Internet access service?
- What are the security implications of having Internet access from every VPN site?
- What are the addressing requirements when every VPN site has direct Internet access?
- What are the benefits of giving Internet access to every VPN site as compared to having a central exit point to the Internet?
- What are the benefits of central firewall service?
- What are the addressing requirements of central firewall service?
- How can customers with private address space use the central firewall service?
- What are the benefits of Wholesale Internet Access service?
- Who assigns the customer address space in the Wholesale Internet Access setup?

Design Options for Integrating Internet Access with MPLS VPN

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Identify different design models for combining Internet access with VPN services.
- List the benefits and drawbacks of these models.
- Explain the implications of their usage.

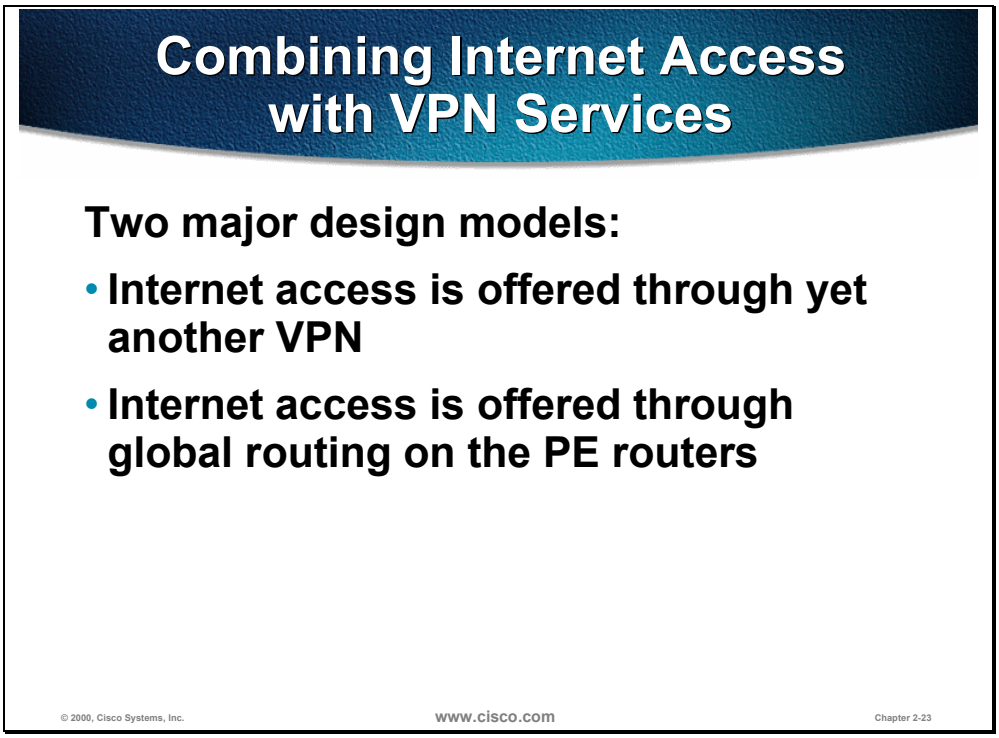

Network designers that want to offer Internet access and MPLS VPN services in the same MPLS backbone can choose between two major design models:

- Internet routing can be implemented as yet another VPN, or
- Internet routing is implemented through global routing on the PE routers.

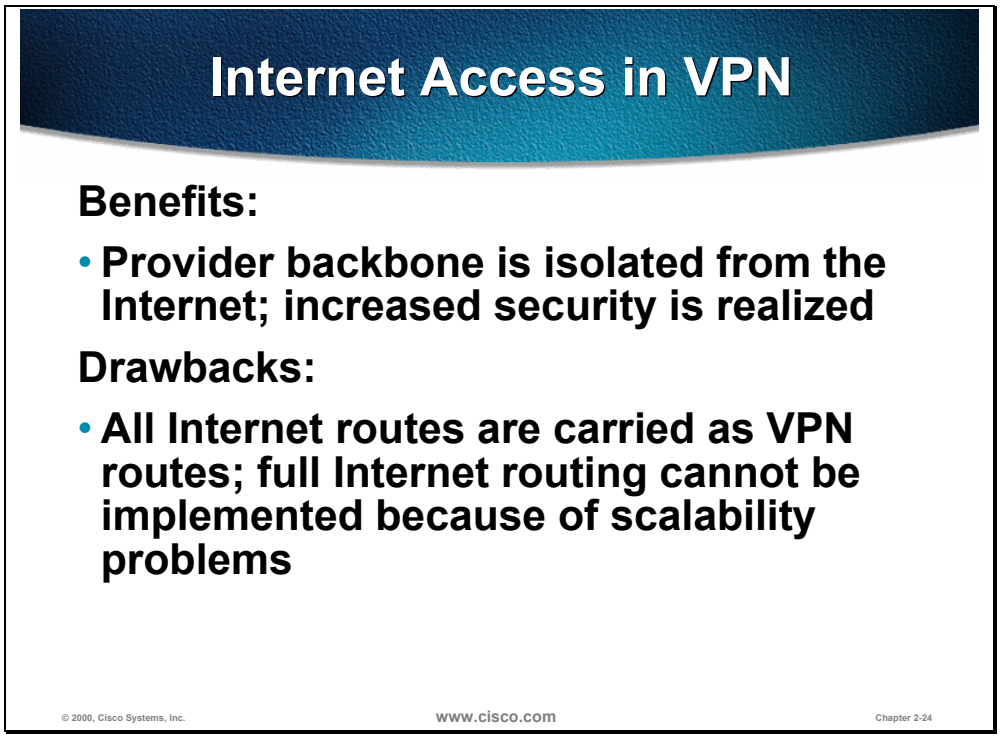

The major benefit of implementing Internet access as a separate VPN is increased isolation between the provider backbone and the Internet, which results in increased security. The flexibility of MPLS VPN topologies also provides for some innovative design options that allow the Service Providers to offer services that were simply not possible to implement with pure IP routing.

The obvious drawback of running the Internet as a VPN in the MPLS VPN architecture is the scalability of such a solution. The Internet VPN simply cannot carry full Internet routing due to scalability problems associated with carrying close to a hundred thousand routes inside a single VPN.

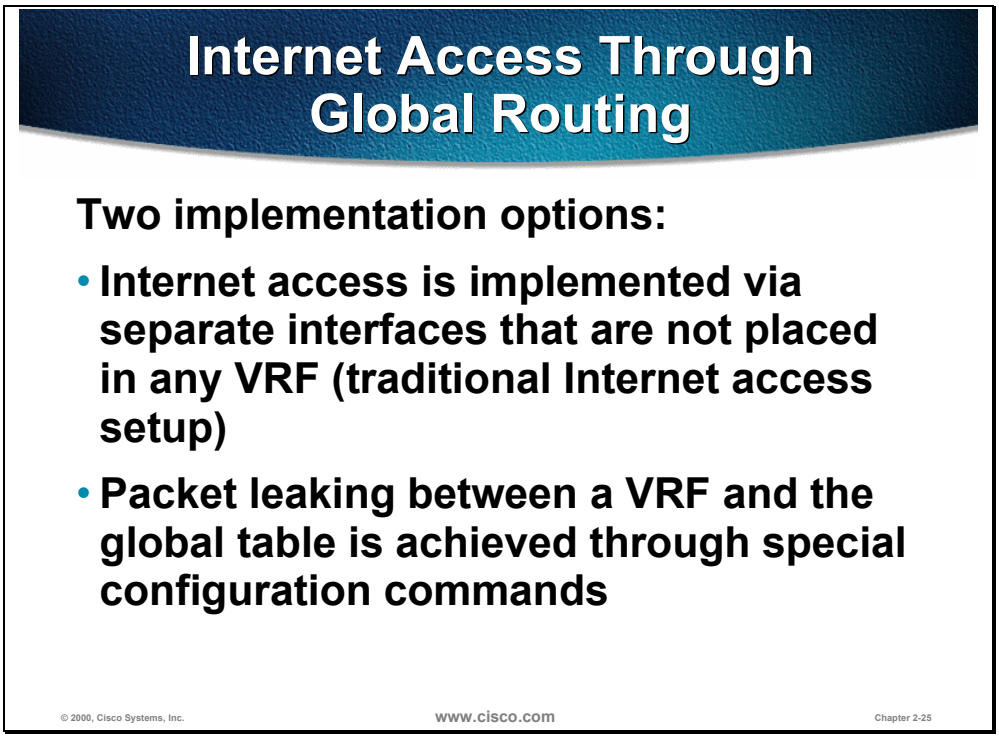

Implementing the Internet access through global routing is identical to building an IP backbone offering Internet services – IPv4 Border Gateway Protocol (BGP) is deployed between the PE routers to exchange Internet routes and the global routing table on the PE routers is used to forward the traffic toward Internet destinations.

VPN customers can reach the global routing table (which is used to forward Internet traffic) in two ways:

- The VPN customer could use a separate logical link for Internet access. This method is equivalent to traditional VPN and Internet access.
- MPLS VPN also provides mechanisms that allow packets originating in a VPN to end in global address space and packets originating in global address space to be forwarded toward a CE router in a VPN.

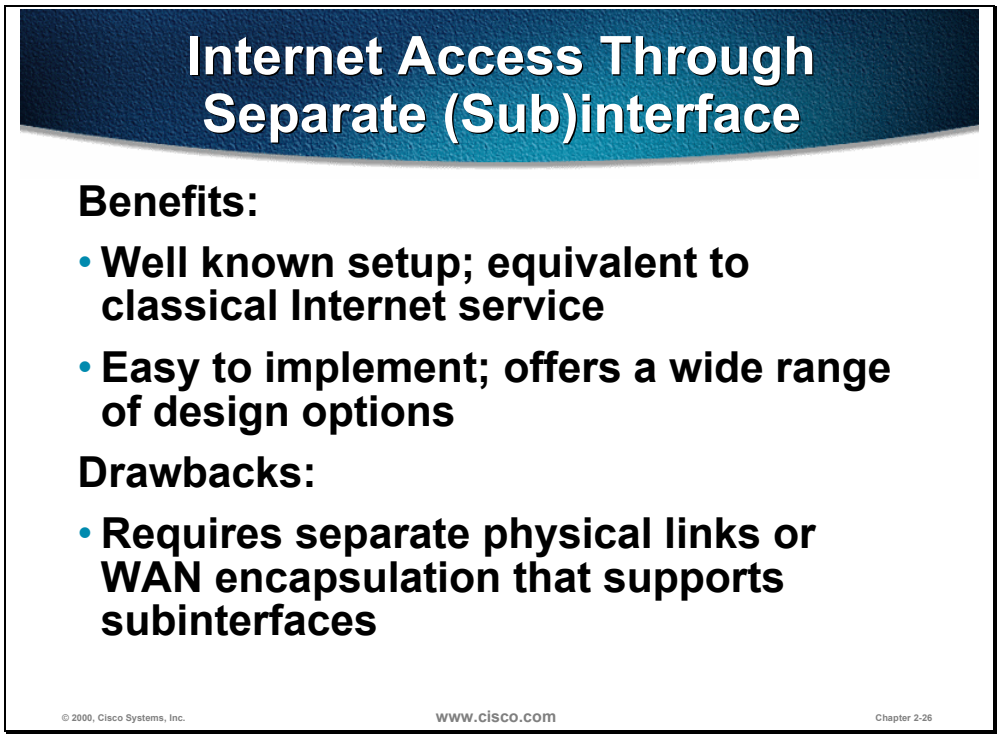

Internet access through separate logical links is easy to set up, because it is equivalent to the classical combination of Internet and VPN service that many customers are using today. This setup is also compatible with all the Internet services required by some customers (for example, the requirement to receive full Internet routing from a Service Provider).

The drawback of this design is the increased complexity, or cost, of the PE-CE connectivity. Separation of Internet and VPN connectivity requires either two separate physical links or a single physical link with WAN encapsulation that supports subinterfaces (for example, Frame Relay).

Note Some customers might be reluctant to change their encapsulation type to Frame Relay as the IP quality of service mechanisms on Frame Relay differ from those provided on point-to-point (PPP) links.

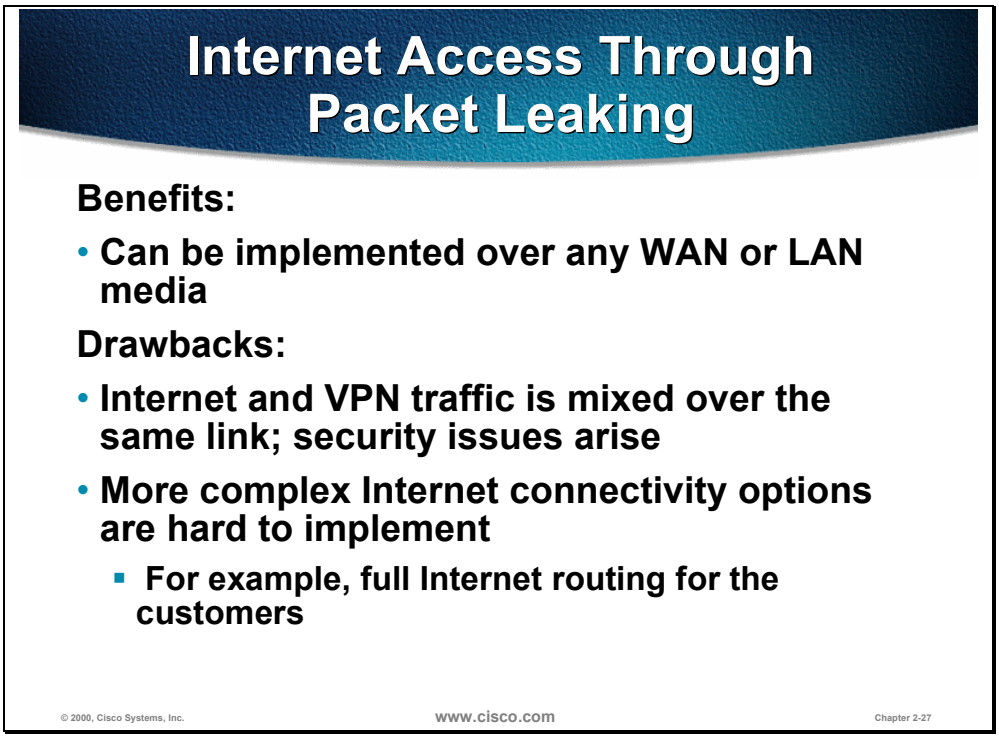

For customers that cannot use Frame Relay encapsulation on the PE-CE link and are not willing to invest into a separate physical link, the packet leaking between VRF and global routing table might be an option. This method can be implemented over any WAN or LAN media, resulting in total access infrastructure flexibility. There are, however, several drawbacks associated with it:

- The Internet and VPN traffic is mixed over the same logical link, resulting in more complex security issues than the more traditional Internet connectivity schemes.
- Some Internet connectivity options (for example, providing full Internet routing to a customer) are harder (although not impossible) to implement.

Summary

There are two major design models you can use for combining Internet access with MPLS VPN services:

- Internet access can be implemented as a separate VPN, or
- Internet access can be implemented through global routing in the PE routers.

Internet access in a VPN is more secure, as there is better isolation between the MPLS VPN backbone and the Internet. MPLS VPN also offers better topology options than pure IP routing. The drawback of this approach is the inability to offer full Internet routing to the customers.

Internet access through global routing is implemented in the same way as a traditional ISP backbone. Customers can be connected to the Internet through separate physical (or logical) links, identical to the traditional way of providing Internet access to the VPN customers.

Alternatively, packet leaking between VRF and global routing table can be used to provide Internet access for customers that are limited by their choice of access method.

Review Questions

- List two major Internet access design models.
- What are the benefits of running an Internet backbone inside a VPN?
- What are the benefits of running an Internet backbone in the global routing table?
- Describe two major implementation options for implementing Internet access in the global routing table.

Leaking Between VPN and Global Backbone Routing

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Design Internet access from VPN that is based on packet leaking between a VRF and a global routing table.
- Identify the benefits and drawbacks of this solution.
- Implement the solution in a MPLS VPN network.

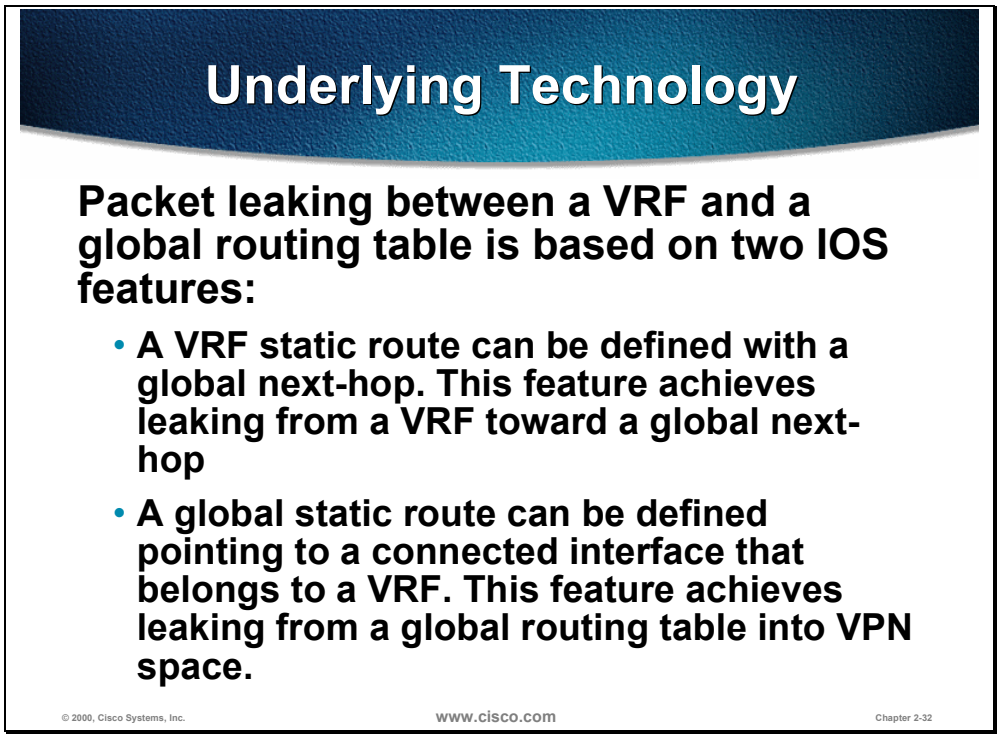

Packet leaking between a VRF and the global routing table is implemented with two IOS mechanisms:

- A static route with a global next-hop can be configured in a VRF. Packets following this static route will end in the global address space at the next-hop router. Traffic originated at a customer site can thus be forwarded into the Internet.
- Global static route can be defined pointing to a connected interface, which belongs to a VRF. This static route is further redistributed into IGP or BGP. Packets originated in the global address space will follow this route (in the global routing table) and will eventually be forwarded toward a CE router. Traffic originating in the Internet can thus be forwarded to the CE router.

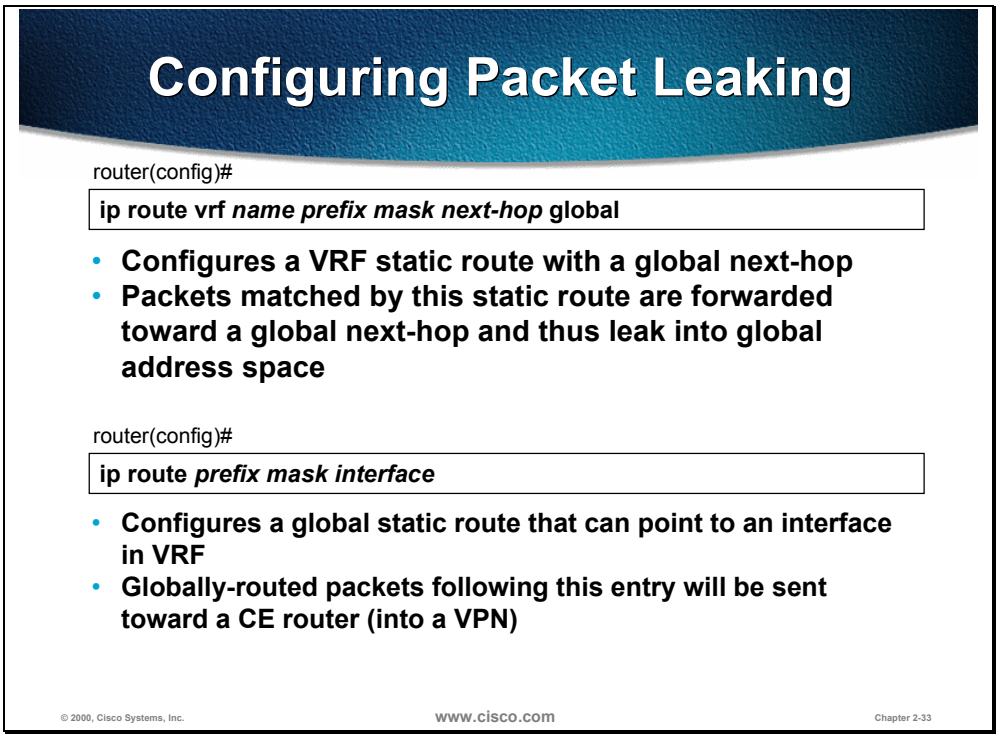

ip route vrf

To establish static routes for a VRF, use the **ip route vrf** command in global configuration mode. 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	Type of network interface to use.
interface-number	Number identifying the network interface to use.
global	(Optional) 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 (tag) value that can be used for controlling redistribution of routes through route maps.

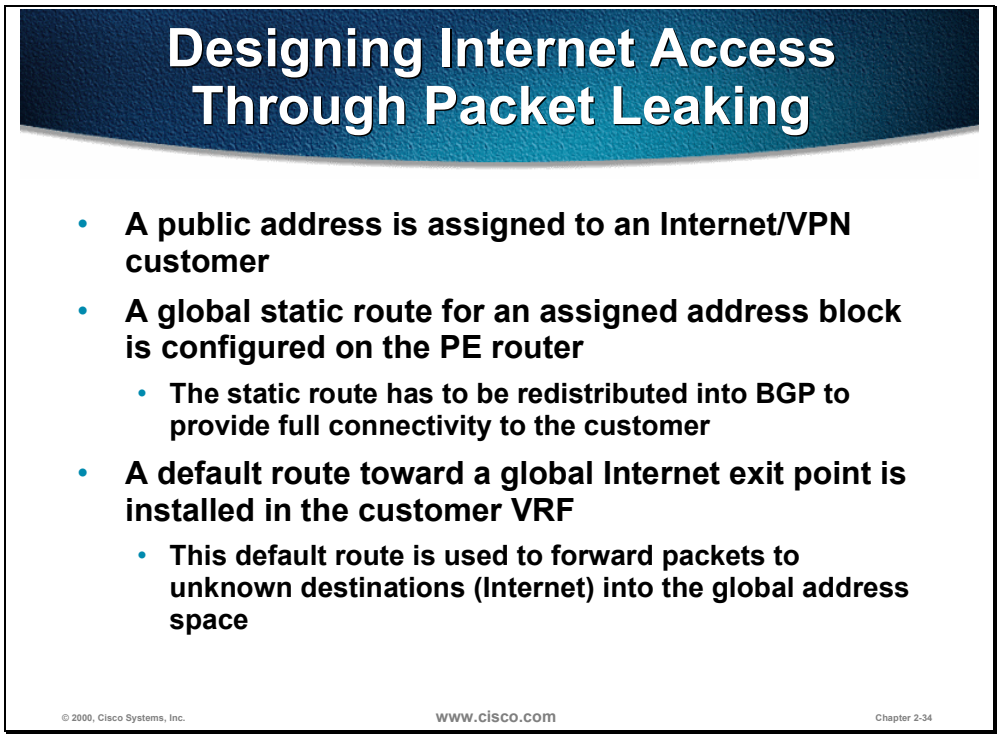

Internet Access through packet leaking is implemented in three steps:

Step 1 A portion of public IP address space is allocated to the customer

A VPN customer, who wants to access the Internet directly without Network Address Translation, needs to use public IP addresses to do so. The customer has to use these addresses within the VPN.

Step 2 Global static route for IP prefix allocated to the customer is configured on the PE router, pointing to the PE-CE link.

The global static route is needed to enable packet forwarding from the global address space toward the customer. This static route needs to be redistributed into the Service Provider's routing protocol (IGP or BGP).

Step 3 Default static route toward an Internet exit point is installed in the customer VRF

This default route is used to forward the packets toward unknown destinations toward a next-hop in global address space. Similar to the previous step, this static route needs to be redistributed into the routing protocol inside the VPN to enable CE routers to reach the Internet.

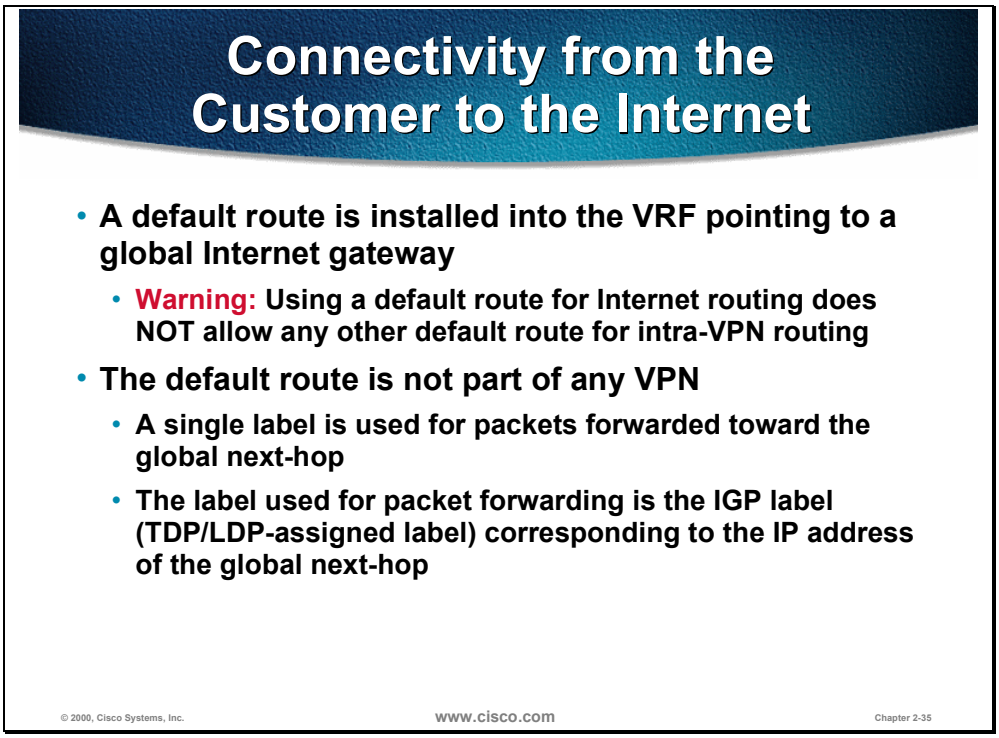

The default route with global next-hop that is used to pass packets from the VPN into the Internet is installed in the VRF on the PE router, preventing any other default routing inside the VPN.

The default route is not part of a VPN, as it has a global next hop. The packet forwarding is also different from standard intra-VPN packet forwarding – the packets received from the CE routers that are using the route with a global next-hop are labeled only with a single label (TDP/LDP-assigned label for the specified next-hop), not with a label stack.

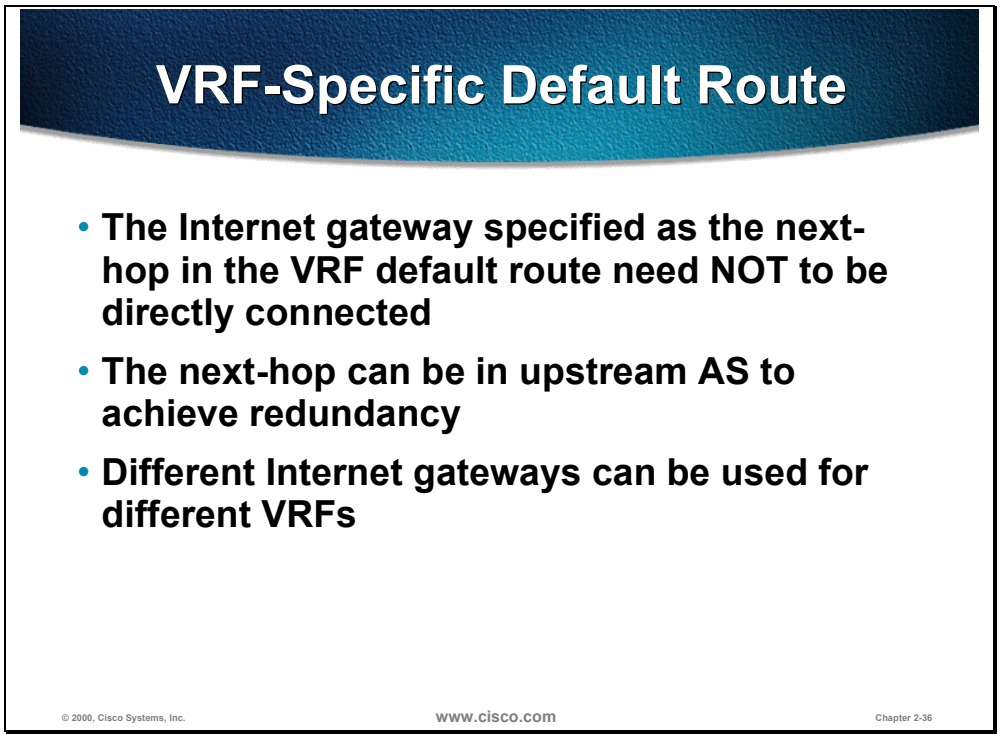

The default route used to reach Internet destinations from a VRF is VRF-specific. Different customers (residing in different VRFs) can therefore use different Internet exit points, even if they reside on the same PE router.

The next-hop (Internet exit point) specified in the default route does not have to be directly connected. Any IP address can be used as the next-hop as long as there is a TDP or LDP label associated with that address. With proper network design, you can use a network in an upstream autonomous system as the next-hop, achieving redundancy between Internet exit points.

Note The next-hop has to be non-local. An IP address on the PE router where the VRF static route is configured cannot be used as a global next-hop.

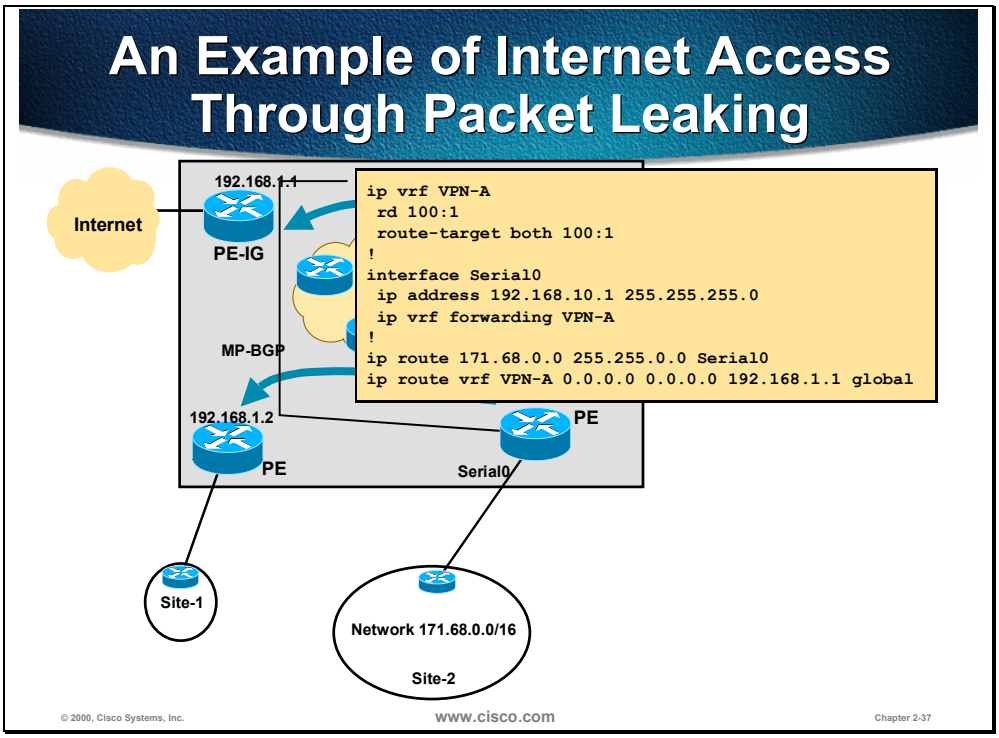

The diagram above shows a typical example of Internet access through packet leaking. A customer VRF (**VPN-A**) is configured on the PE router and an interface is associated with the VRF. A default route is then installed in the VRF, pointing to a global next-hop (PE-IG router). A global route is configured for the customer's IP prefix (172.68.0.0/16), pointing to the PE-CE interface of the PE router.

Note This example does not include redistribution of static routes into the intra-VPN and global routing protocols.

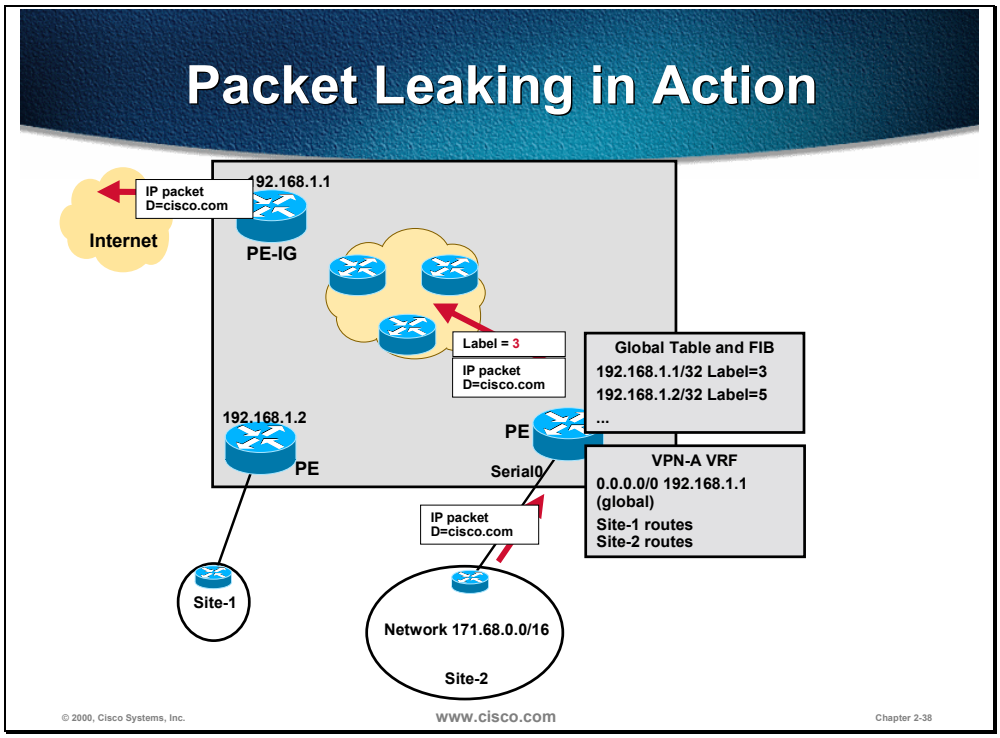

With the static routes configured on the PE router, the routing and forwarding table in the VPN_A VRF contains routes for Site-1 and Site-2 as well as a default route pointing toward the PE-IG. The label in the VRF FIB associated with the default route is copied from the global FIB (label 3 is used in our example, as it is the label associated with 192.168.1.1/32 in the global FIB).

When a packet is received toward a destination not reachable thorough any other VRF route, the default route is used and the packet is labeled with a single label and forwarded toward the PE-IG router.

Usability of Packet Leaking for Various Internet Access Services

In the following pages we'll analyze whether we can implement various Internet Access Services with the packet leaking mechanism.

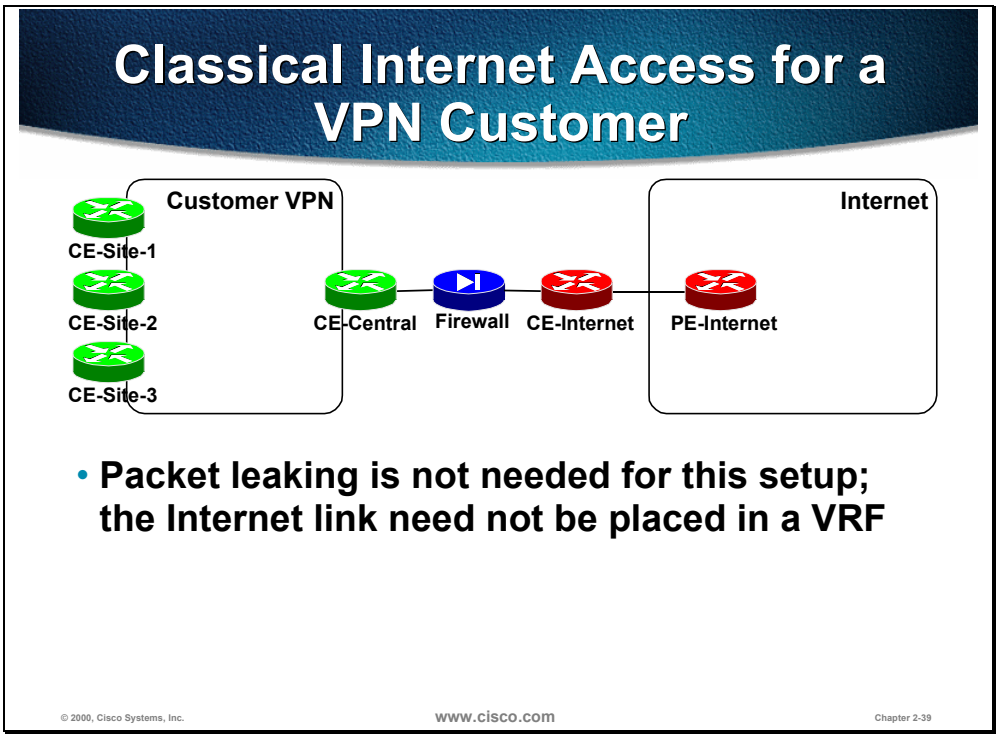

The classical Internet Access service does not need packet-leaking mechanism, as there are always two links between the customer and the provider – a link between CE-Central and the PE router providing VPN services and a link between CE-Internet and the PE router providing Internet services (the same PE router can act in both roles).

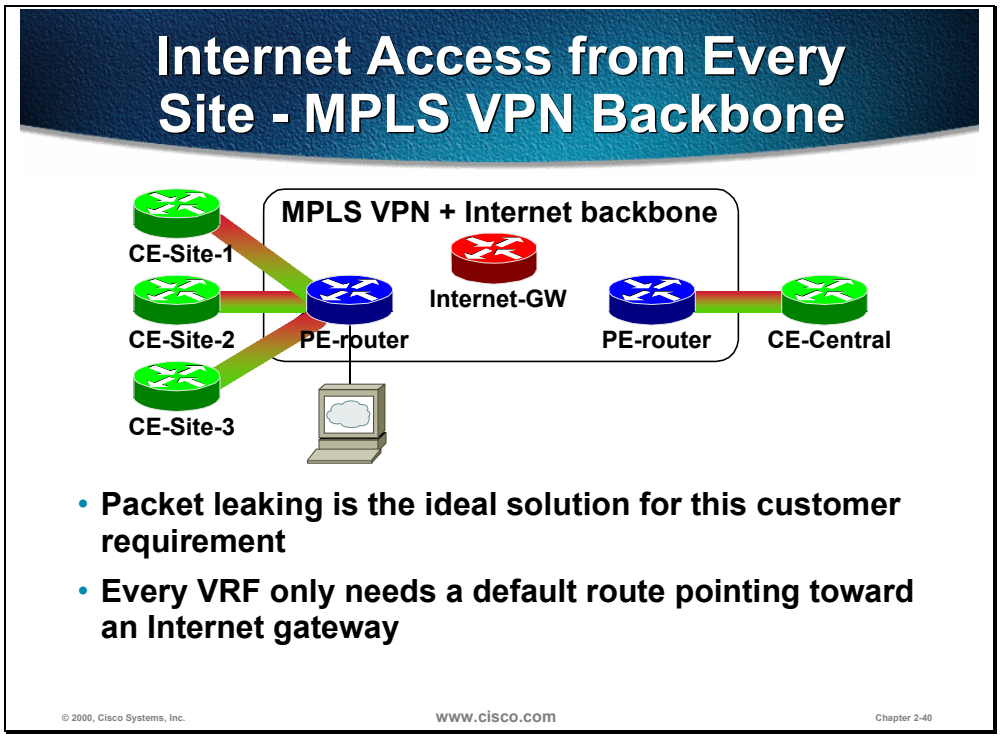

For customers that want to access Internet from every site without incurring the additional costs or complexity of separate VPN and Internet link, the packet leaking is an ideal solution.

To achieve optimum routing between the customer sites and Internet destinations, the default route pointing toward an Internet exit point needs to be installed in every VRF (if the default route would only be installed in one VRF, all packets from the customer sites would have to traverse that PE router). The default routes could use the same Internet gateway, but this setup might result in suboptimal routing for geographically dispersed customers. For large geographically dispersed customers each default route should use a next-hop address of an Internet router closest to the PE router.

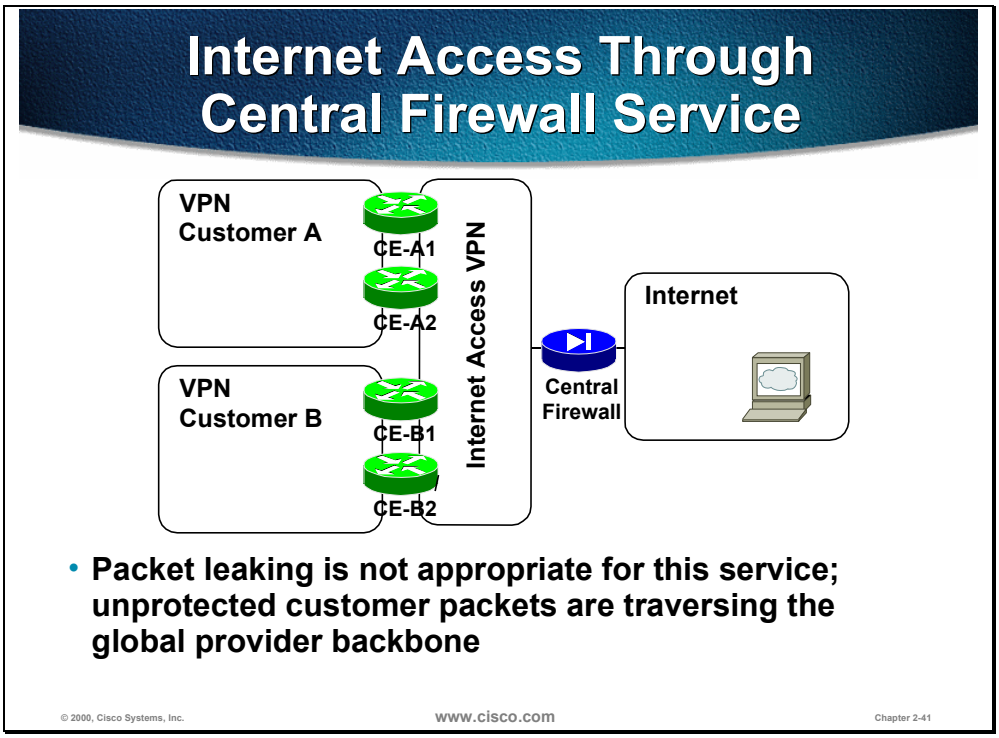

The central firewall service cannot be implemented with packet leaking. Unprotected customer packets traversing the infrastructure between the VPN customers and the central firewall (the Internet Access VPN in the diagram above) would be routed in the global address space together with other Internet traffic, resulting in unacceptable risk.

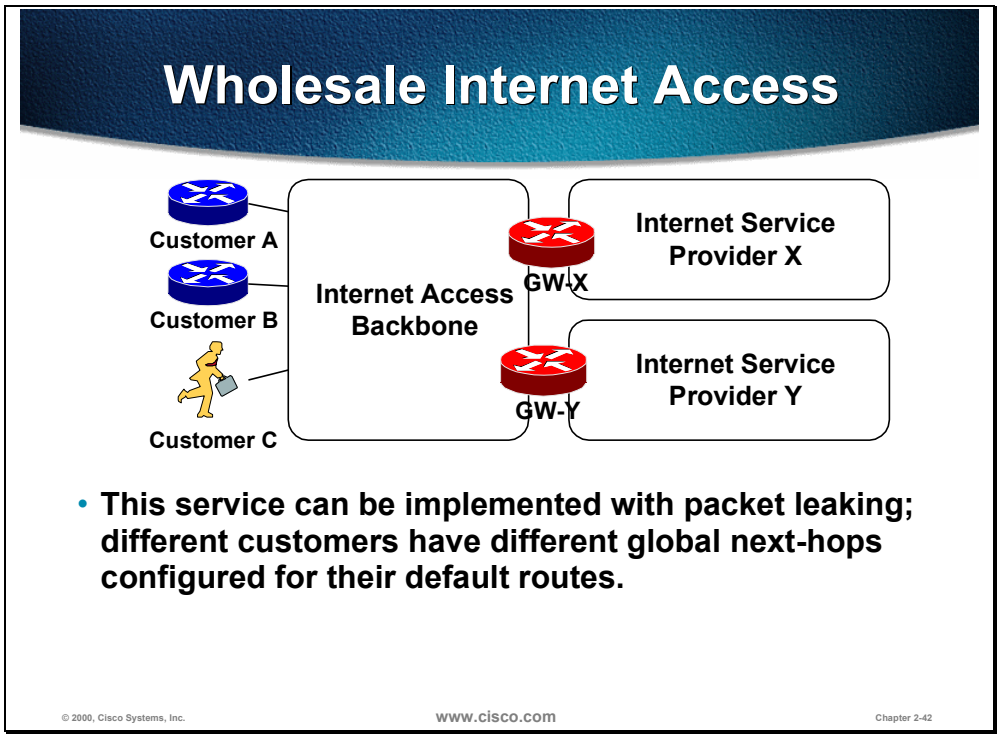

Packet leaking can also be used to implement Wholesale Internet Access service – the next hop of the per-VRF default route indicates the upstream ISP the customer wants to use.

Redundant Internet Access with Packet Leaking

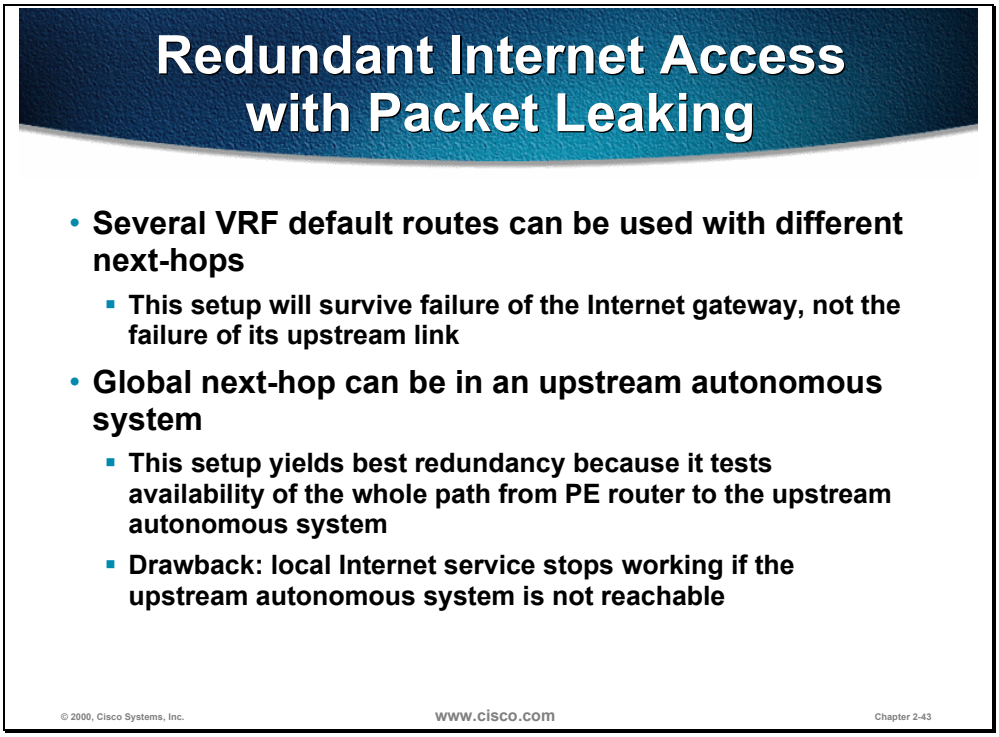

There are two methods that can be used to ensure redundant Internet access implemented through packet leaking mechanism:

- Several default routes can be installed in the VRF, associated with different global next-hops. This setup will survive the failure of the next-hop (when the next-hop is not present in global routing table, the per-VRF default route is not used), but not any other failure (for example, the failure of upstream link between the Internet gateway and upstream ISP).
- The next-hop specified in the per-VRF default route could use a network from an upstream autonomous system. This setup yields best redundancy, as it provides protection for the whole path between the PE router and upstream autonomous system. The drawback of this approach is that the availability of Internet service depends on the availability of a network in an upstream AS. If the upstream AS fails, the customer has no Internet connectivity whatsoever, even though the local Internet destinations are still available.

The best redundancy is provided by a combination of both mechanisms:

- A default route with next-hop in a neighboring AS is used as the primary default route
- A second (floating) default route with a next-hop in the Service Provider's network is used as a backup.

Note The redundancy mechanisms outlined in this section only work well if the PE router has no default route in its global routing table.

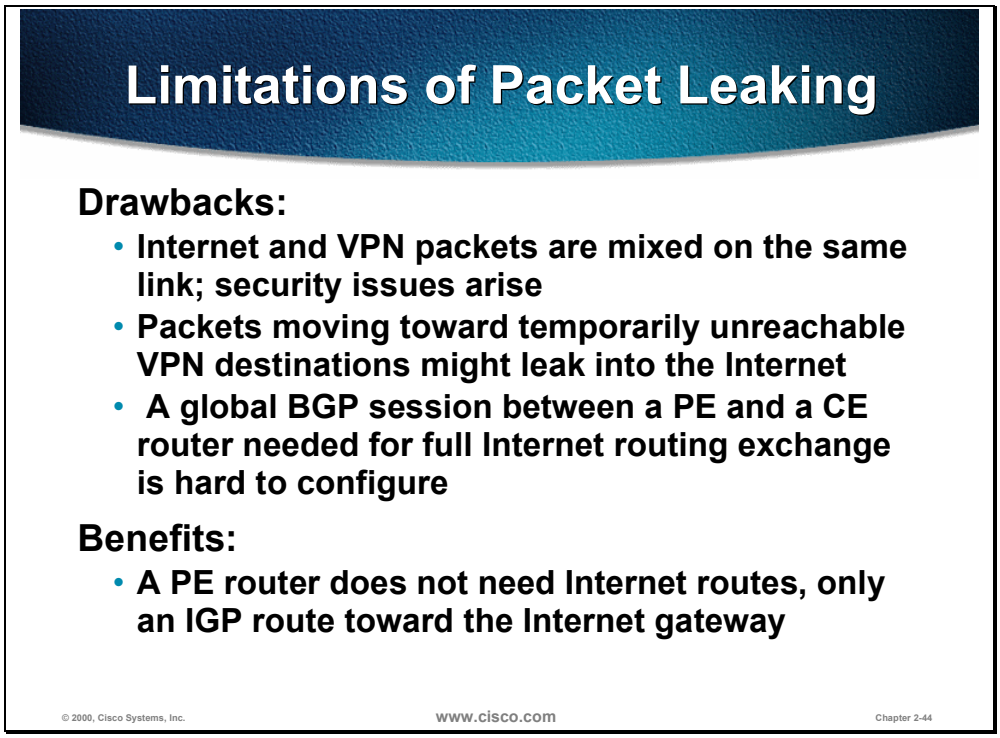

There are several drawbacks associated with the packet leaking mechanism:

- Internet and VPN packets are mixed on the same PE-CE link, resulting in potential security vulnerabilities.
- The VRF default route will be used to forward packets sent toward VPN destinations that are currently not available. These packets will therefore end in the global address space, resulting in decreased privacy of customer's traffic.
- It is difficult to implement Internet BGP sessions between the PE router and the CE router for customers that need full Internet routing.

On the other hand, the packet leaking mechanism significantly reduces the routing overhead placed on the PE router, as it needs no Internet routes (or even a default route). The only entry in the routing table on the PE router needed for successful packet leaking from a VPN into the Internet is the IGP route toward the Internet gateway.

Summary

In this section, you've seen the first mechanism that can be used to give MPLS VPN customers Internet access – the packet leaking between a VRF and the global address space. The packet leaking is implemented using two mechanisms in Cisco IOS:

- Leaking from VRF into the global address space is configured using a per-VRF static route with a global next hop.
- A global static route pointing toward a PE-CE interface is used to forward packets from global address space toward a CE router.

The packet-leaking mechanism is well suited for customers that need Internet access from every site and for Wholesale Internet Access services.

Review Questions

- Which IOS mechanisms are used to implement packet leaking between a VRF and a global address space?
- How is the leaking from a VRF into the global address space accomplished?
- How do you configure leaking from global address space toward a CE router?
- How is packet leaking used to implement Internet access service for VPN customers?
- What label is used to forward packets toward a global next-hop?
- What are the benefits of Internet access based on packet leaking?
- Which Internet access services can be implemented with packet leaking?
- Which Internet access services cannot be implemented with packet leaking?

Separating Internet Access from VPN Service

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Design an Internet access service where the Internet access is totally separate from MPLS VPN service.
- Identify PE-CE requirements for this solution.
- Implement the solution in a MPLS VPN network.

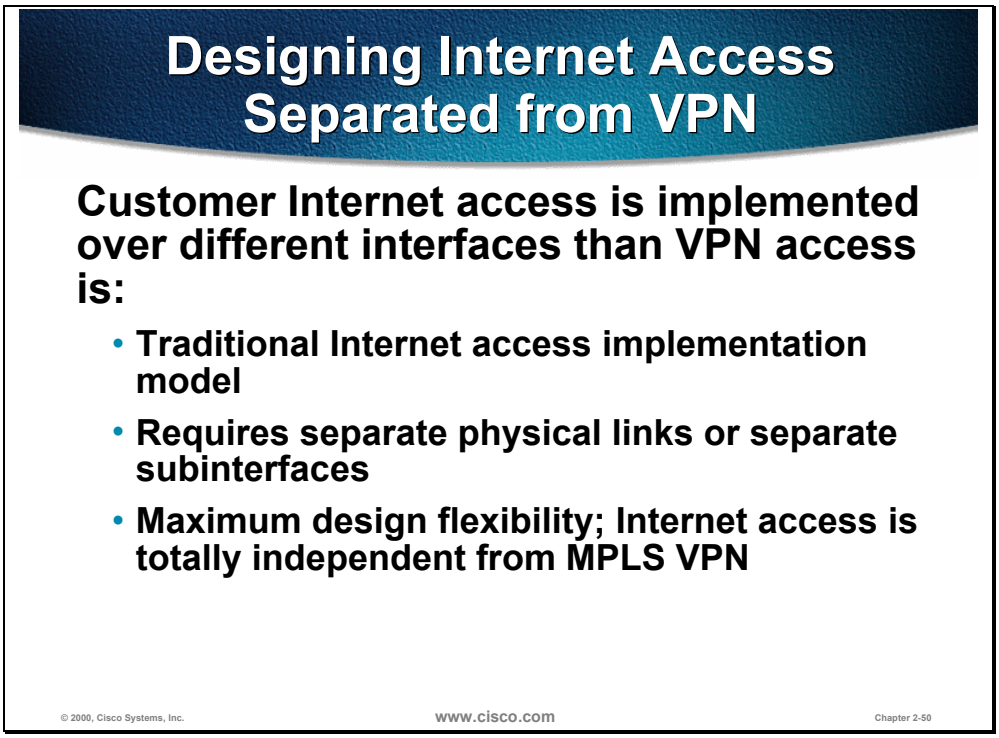

Internet access can always be implemented with the traditional implementation model, with two links between the customer's site(s) and the Service Provider network – a VPN link and an Internet link. The two links can be implemented with one physical link if you use a layer-2 encapsulation that supports subinterfaces (Frame Relay, ATM or VLAN).

The traditional Internet access implementation model gives maximum design flexibility, as the Internet access is completely separated from the MPLS VPN service. Nevertheless, the limitations of traditional IP routing prevent this implementation method from being used for innovative Internet Access solutions such as Wholesale Internet Access.

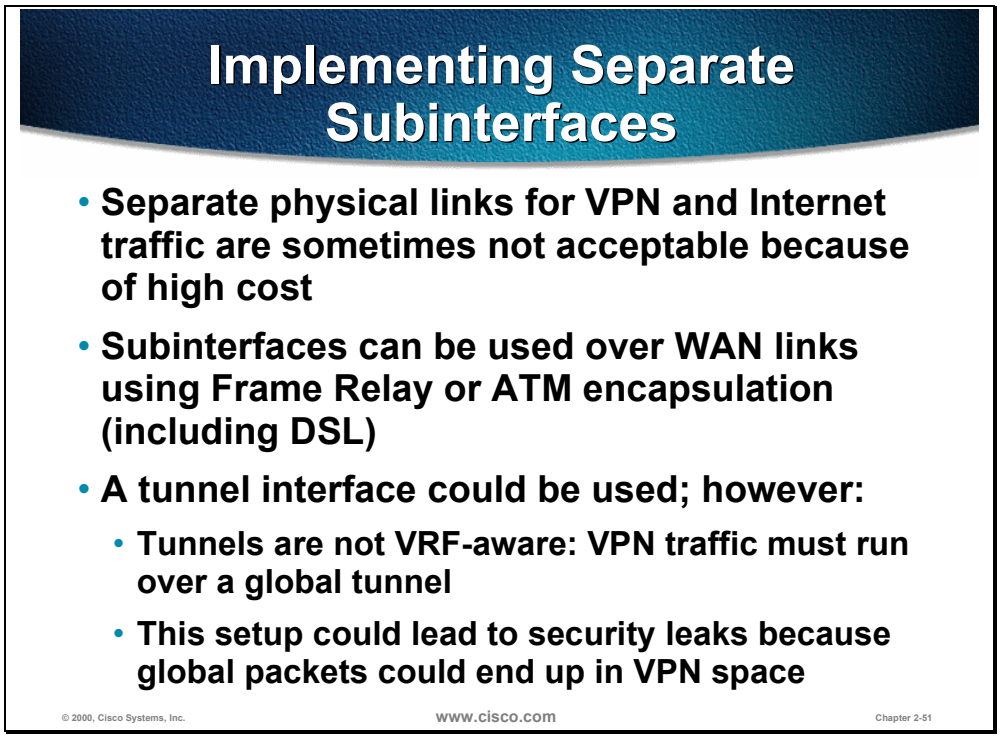

In situations where the cost factor prohibits separate physical links for VPN and Internet traffic, subinterfaces can be used to create two logical links over a single physical link. Subinterfaces can only be configured on WAN links using Frame Relay or ATM encapsulation (including xDSL) and on LAN links using any VLAN encapsulation (ISL or 802.1q). For all other encapsulation types, a tunnel interface could be used between the CE and the PE router.

However, the use of tunnel interfaces is strongly discouraged for security reasons:

- A tunnel interface on the PE router is not VRF-aware. The endpoints of the tunnel have to be in global routing table – the VPN traffic must be tunneled across an Internet interface.
- It's also very easy to spoof GRE tunnels (if the tunnel key is configured, and the key is known). An intruder from the Internet could easily generate traffic that could appear as if it were coming over the GRE tunnel from the CE router and would therefore be forwarded into the customer's VPN.

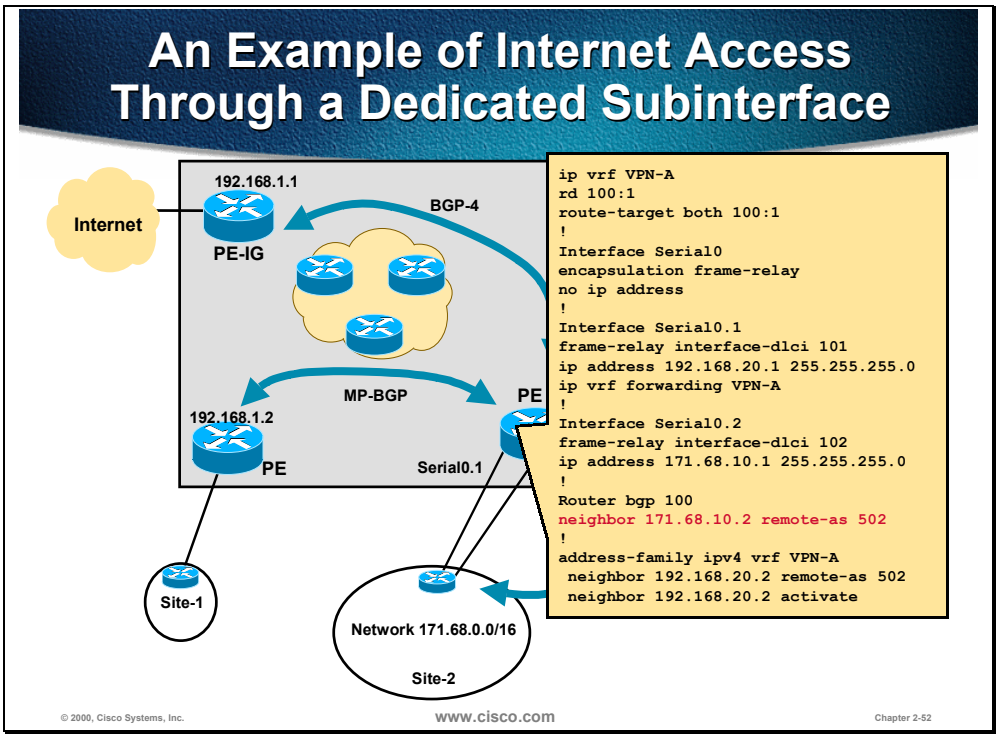

The example above illustrates the configuration needed to implement Internet Access through a dedicated Frame Relay interface. The following configuration steps are performed:

- The customer VRF (**VPN-A**) is created.
- Frame Relay encapsulation is configured on the PE-CE link (**Serial 0**).
- VPN subinterface (**Serial 0.1**) is created and associated with DLCI **101**.
- Internet subinterface (**Serial 0.2**) is created and associated with DLCI **102**.
- CE router is configured as a BGP neighbor in both the global BGP process and inside the VPN in the VRF **VPN-A**.

Note Allowas-IN feature would have to be configured on the PE router if the customer is propagating individual site routes to the Internet through BGP.

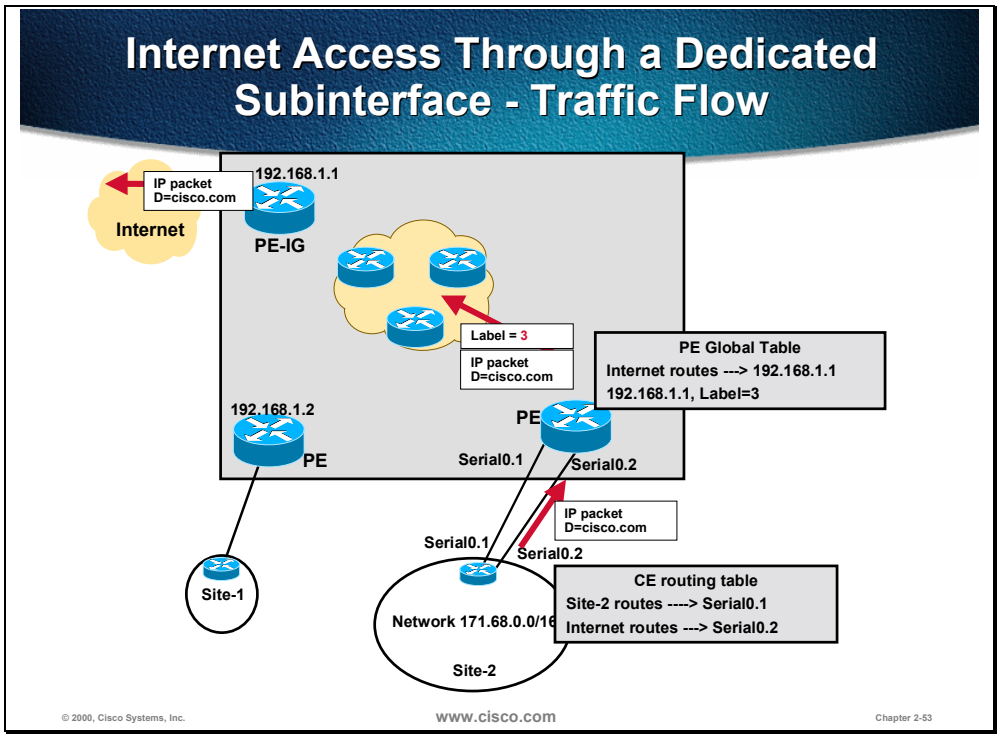

The Internet traffic flow in this setup is identical to the traditional Internet traffic flow – when a packet is received from the CE router through the Internet subinterface, a lookup is performed in the global FIB on the PE router and the packet is forwarded toward BGP next-hop.

Usability of Separated Internet Access for Various Internet Access Services

In this section we'll analyze whether we can implement various Internet Access Services with the packet leaking mechanism.

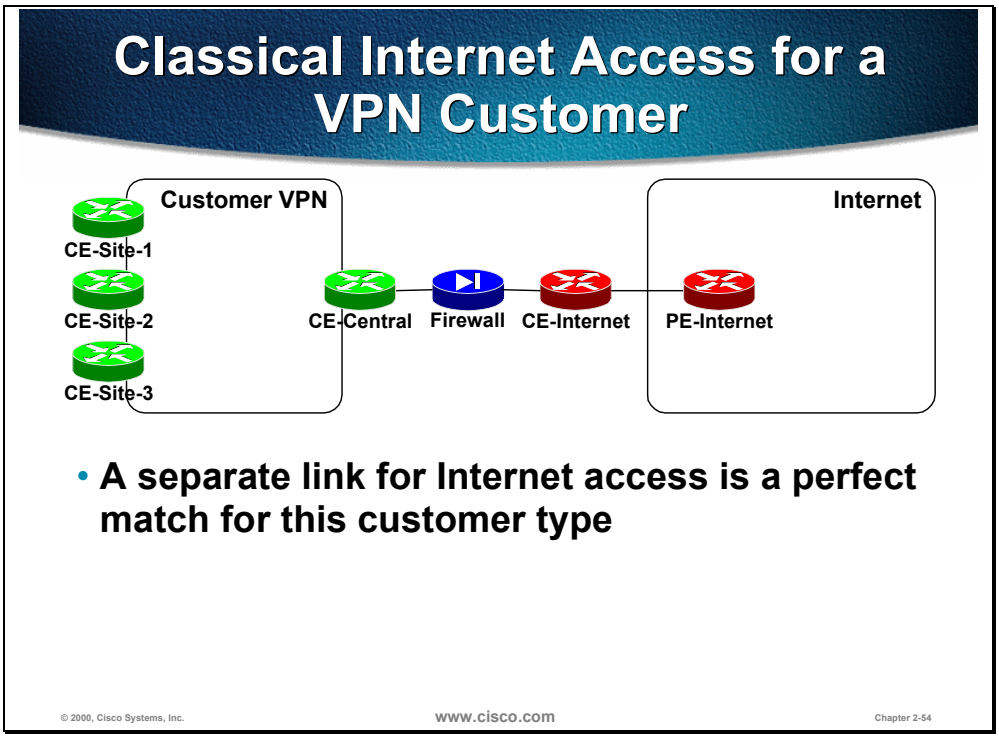

Classical Internet Access setup for a VPN customer is based on a separated Internet Access design model. This design model is thus a perfect match for the customers looking for Classical Internet Access service.

Internet Access from Every Customer Site

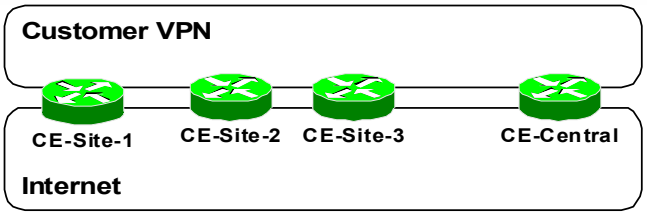

- **Using separate link(s) for Internet access will lead to a complex setup for this customer type**
- **Every CE router needs two links (or subinterfaces) to its PE router**

For customers that need Internet access from every site, two physical (or logical) links between every CE router and the PE routers might prove to be too complex or too expensive.

Summary

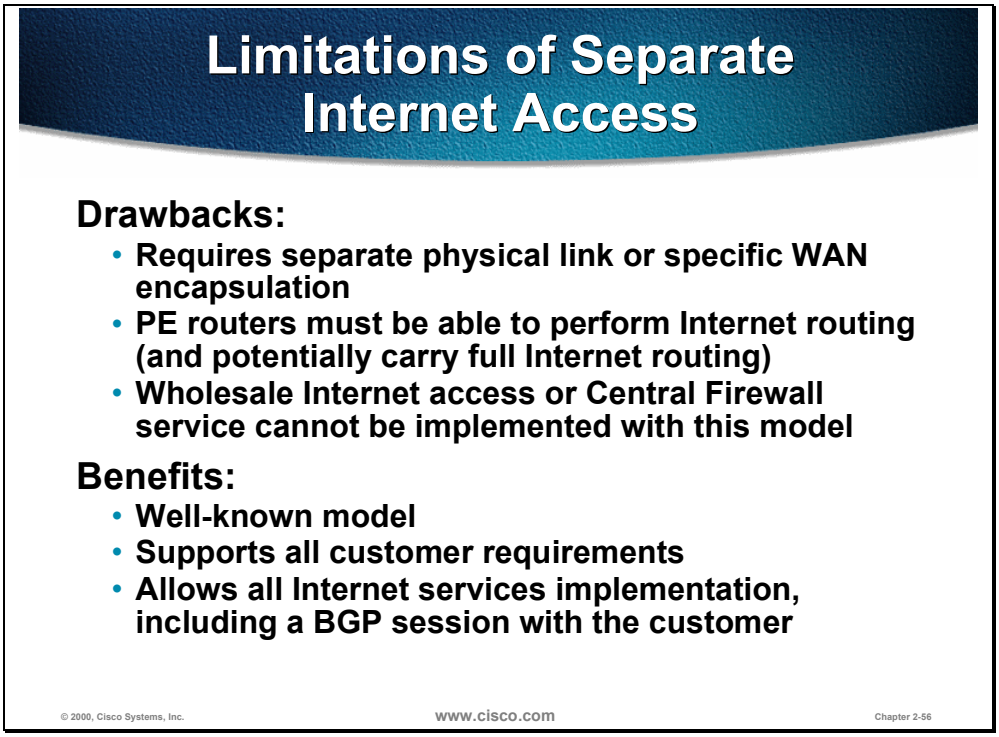

The benefits of Separate Internet Access design model are obvious:

- It is a well-known and widely understood model.
- It supports all customer requirements, including multi-homed customer connectivity with full Internet routing.

The drawbacks of this model are:

- It requires two dedicated physical links between the PE and the CE router or specific WAN/LAN encapsulations that might not be suitable for all customers.
- PE routers must be able to perform hop-by-hop Internet routing and use either default route to reach the Internet or carry full Internet routing table.
- Advanced Internet Access services (centralized managed firewall service or wholesale Internet access service) cannot be realized with this model at all.

Review Questions

- What is the effect of MPLS VPN technology on implementing Internet access through a separate (sub)interface?
- Which WAN encapsulation types can be used to avoid using two physical links?
- What are the benefits of using a separate (sub)interface for Internet access?
- Which Internet access services cannot be implemented within this model?

Internet Access Backbone as a Separate VPN

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Design Internet access solutions where the Internet access is provided through a separate VPN.
- Identify the scaling issues of this design.
- Implement the design in a MPLS VPN network.

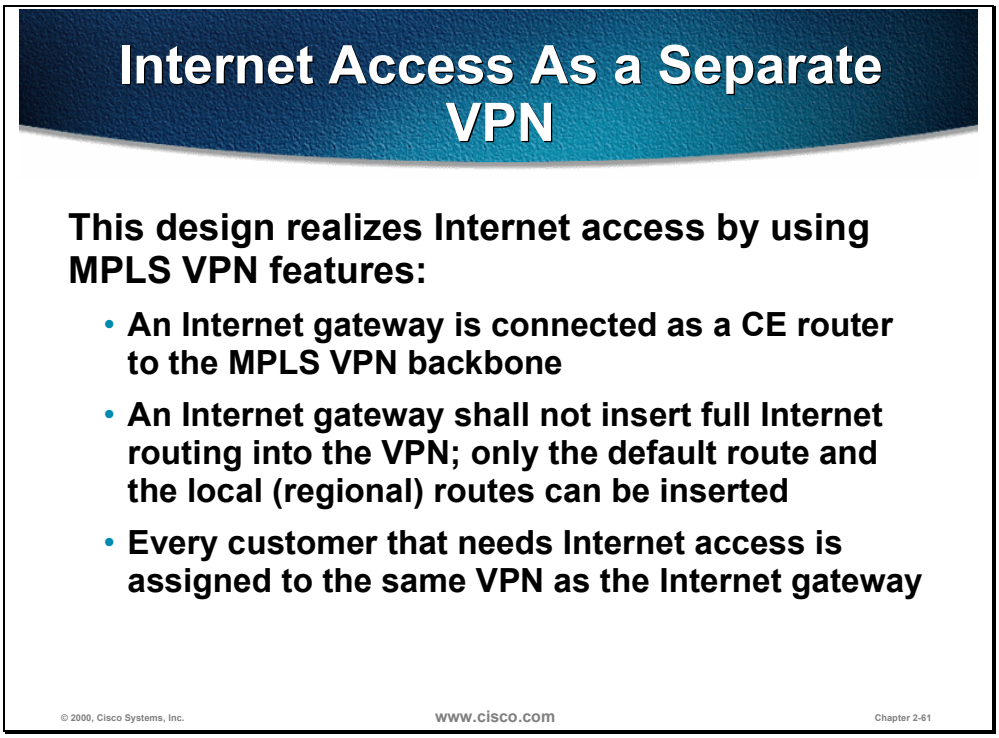

MPLS VPN architecture suggests an obvious solution to Internet Access for VPN customers – define the Internet as yet another VPN and use various MPLS VPN topologies to implement various types of Internet access. Under this design model, the Internet gateways appear as CE routers to the MPLS VPN backbone and the customer's Internet access is enabled by combining Internet VPN with the customer VPN in the customer's VRFs (overlapping VPN topology).

The Internet VPN should not contain the full set of VPN routes, as that would make the solution completely non-scalable. The Internet gateway routers (CE routers) should announce a default route toward the PE routers. To optimize local routing, the local (or regional) Internet routes shall also be inserted in the Internet VPN.

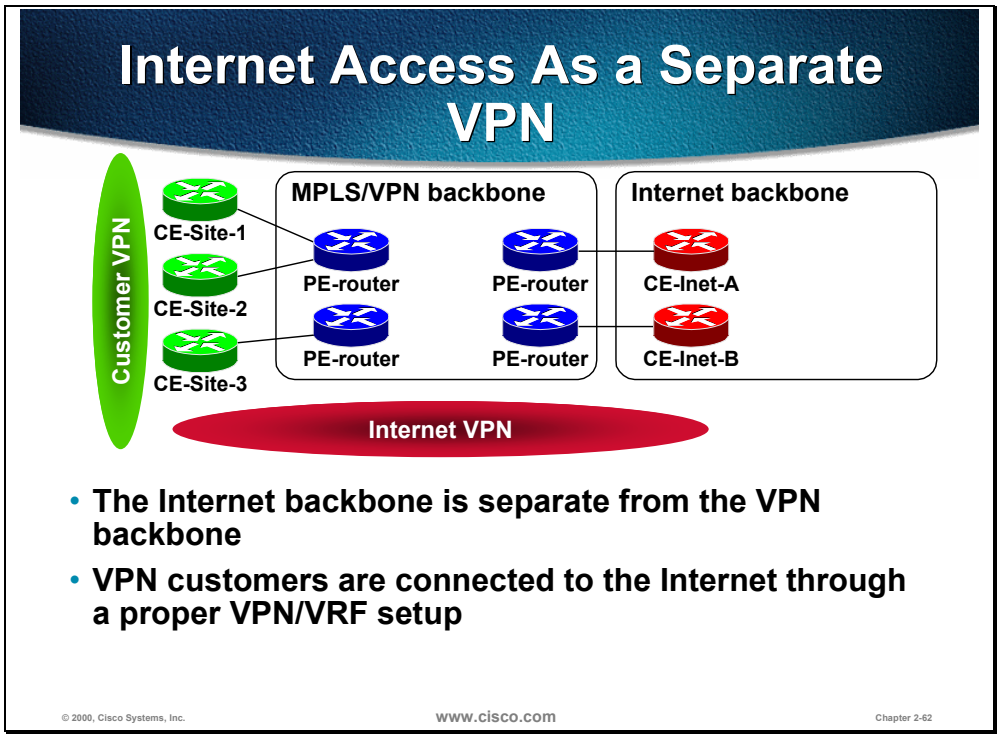

When implementing Internet access as a separate VPN, the Internet backbone is separate from the MPLS VPN backbone, resulting in increased security of the MPLS VPN backbone (for example, Internet hosts can only reach PE routers, but not the P routers). The VPN customers are connected to the Internet simply through proper VRF setup.

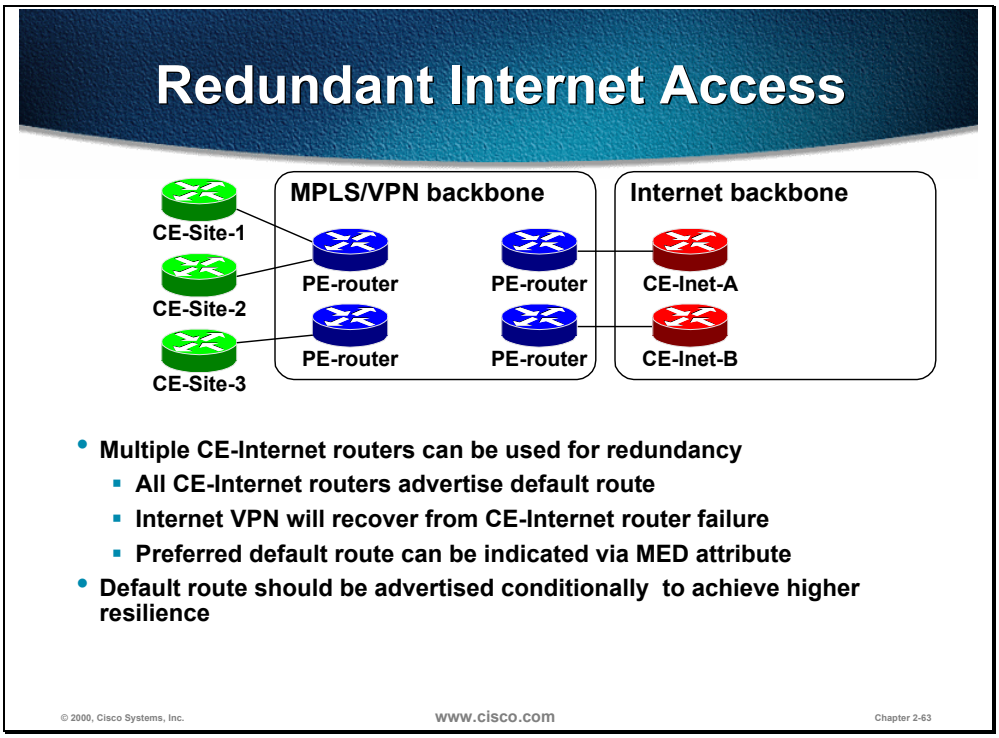

Redundant Internet access is easy to achieve when the Internet service is implemented as a VPN in the MPLS VPN backbone:

Step 1 Multiple Internet gateways (acting as CE routers) have to be connected to the MPLS VPN backbone to ensure router and link redundancy.

Step 2 All Internet gateways advertise the default route to the PE routers, resulting in routing redundancy.

The Internet gateways also announce local Internet routes. As these routes would be announced with different BGP attributes (most notably MED), the PE routers will select the proper CE-Inet router as the exit point toward those destinations.

The MED attribute can also be used to indicate the preferred default route to the PE routers. In this setup, one CE-Inet router acts as a primary Internet gateway and the other CE-Inet router(s) acts as a backup.

Step 3 The redundancy established so far covers the path between customer sites and the CE-Inet routers. A failure in the Internet backbone might break the Internet connectivity for the customers if the CE-Inet routers announce the default route unconditionally. Conditional advertisement of the default route is therefore configured on the CE-Inet routers – they announce only the default route to the PE routers if they can reach an upstream destination.

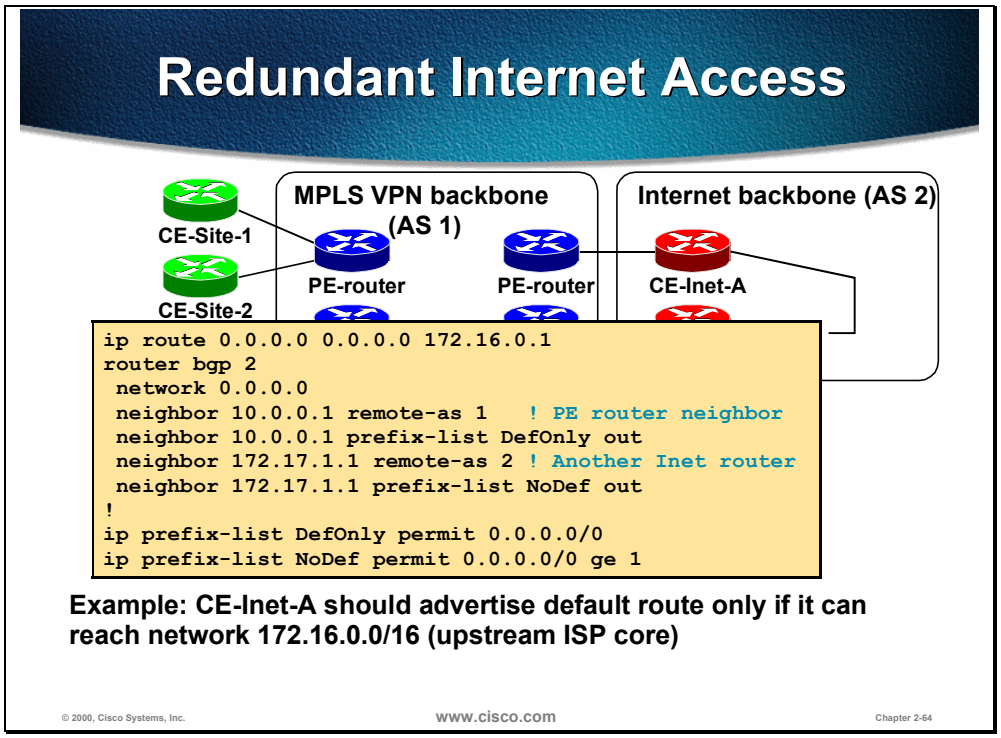

The example contains sample configuration of a CE-Inet router with conditional default route advertisement. The CE-Inet-A router will only advertise the default route to the PE-router if it can reach the network 172.16.0.0/16.

The following steps are used to configure this functionality:

Step 1 A static default route is configured toward a next hop in network 172.16.0.0. If the network 172.16.0.0 is not reachable, this static route will not enter the IP routing table.

Step 2 The default route origination is configured in the BGP routing process with the **network** command. The default route will only be originated in BGP if it is present in the IP routing table (which, based on the previous step, means that the network 172.16.0.0/16 is reachable.)

Step 3 Prefix lists are used to filter BGP routing updates – the default route is only sent to the PE routers, not to the other Internet routers.

Usability of Internet in a VPN Solution for Various Internet Access Services

Next we'll analyze whether we can implement various Internet Access Services with the Internet-in-a-VPN solution.

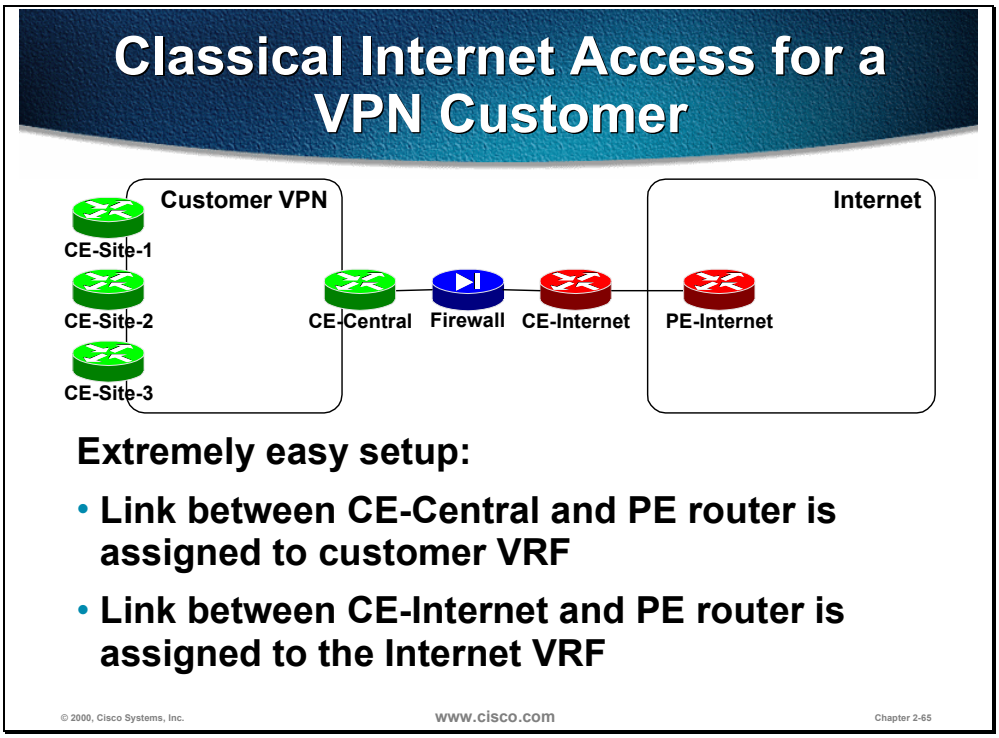

The classical Internet access model can be easily implemented with the Internet configured as a VPN over MPLS VPN backbone – the link between a PE router and the CE-Internet router is assigned to the Internet VRF, the link between a PE router and the CE-Central router is assigned to the customer VRF. The External BGP (EBGP) multihop session can be configured between the Internet gateway (CE-Inet router in the previous diagram) and the CE-Internet router in this diagram to give full Internet routing to the customer.

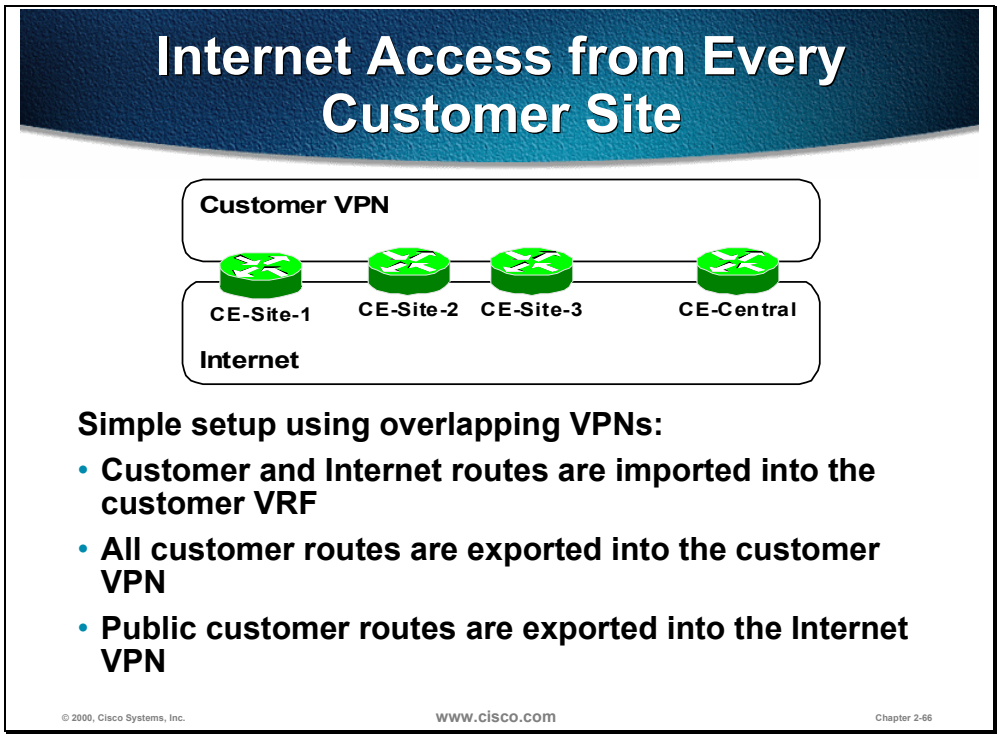

Internet access from every customer site is best implemented with an overlapping VPN solution:

- Customer routes are marked with a customer-specific (**Customer**) route target.
- Internet routes are marked with a special (**Internet**) route target.
- Customer sites that need to reach Internet are placed in a separate VRF. **Customer** and **Internet** routes are imported into this VRF and the routes exported from this VRF are marked with **Customer** and **Internet** route targets.

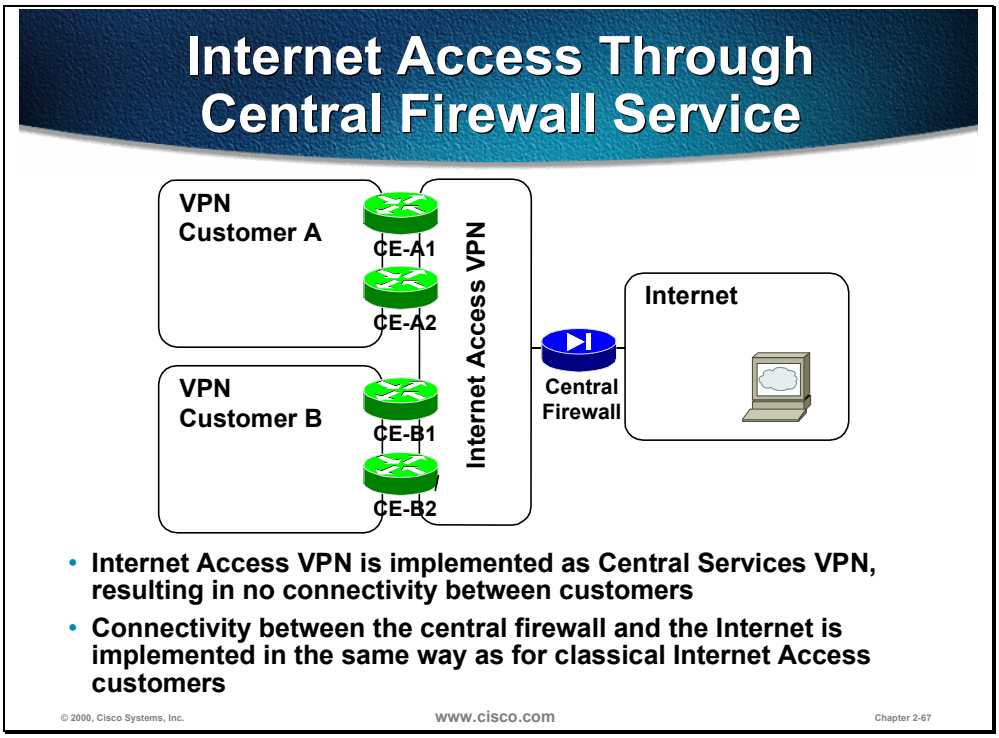

The central managed firewall service should be implemented with the Central Services VPN topology, with the central firewall being the server site and all customer CE routers residing in client sites. For customers with their own VPNs implemented over the same MPLS VPN backbone, the topology that overlaps customer VPN and Central Services VPN should be used.

The Central Services VPN prevents direct exchange of traffic between client sites, resulting in satisfying security for the customers of this service.

Connectivity between the central firewall and the Internet is implemented in the same way as the Internet access for classical Internet customers. If the Internet is configured in a VPN, the public interface of the firewall is connected to an interface on a PE router, which is placed in the Internet VRF.

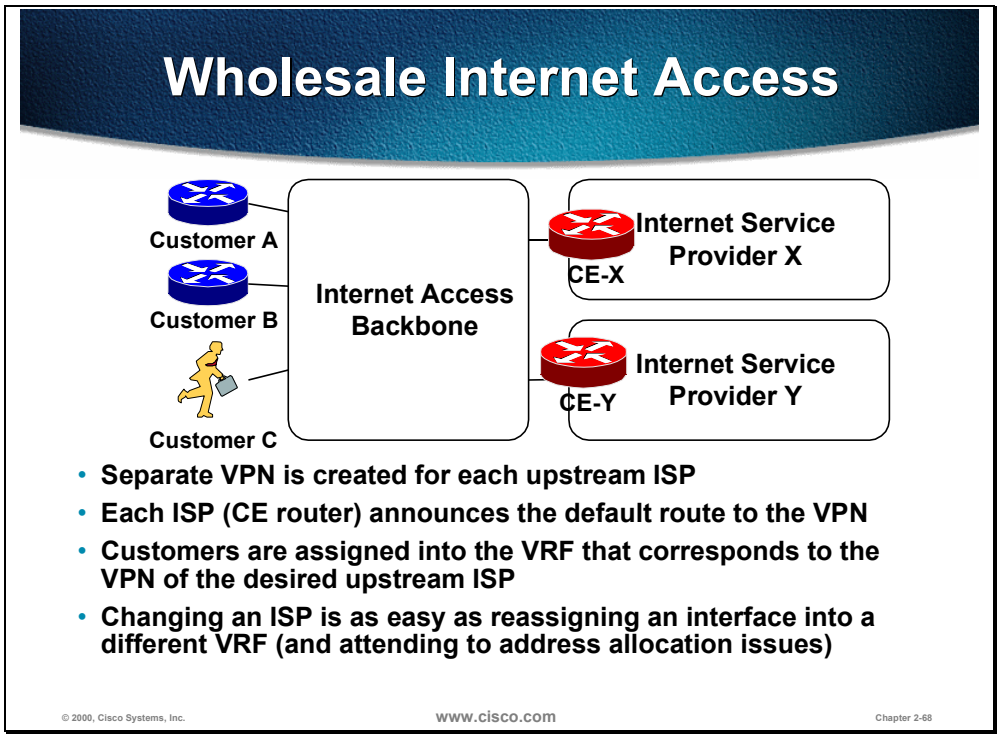

Wholesale Internet Access is implemented by creating a separate VPN for every upstream ISP. The Internet gateway of the upstream ISP (acting as a CE router toward the MPLS VPN-based Internet Access Backbone) announces a default route, which is used for routing inside the VPN.

Customers are tied to upstream service providers simply by placing the PE-CE link into the VRF associated with the upstream service provider. Changing an ISP becomes as easy as reassigning the interface into a different VRF and attending to address allocation issues. For customers using access methods supporting dynamic address allocation (for example, dial-up or cable), the new customer IP address from the address space of the new ISP is assigned automatically.

Summary

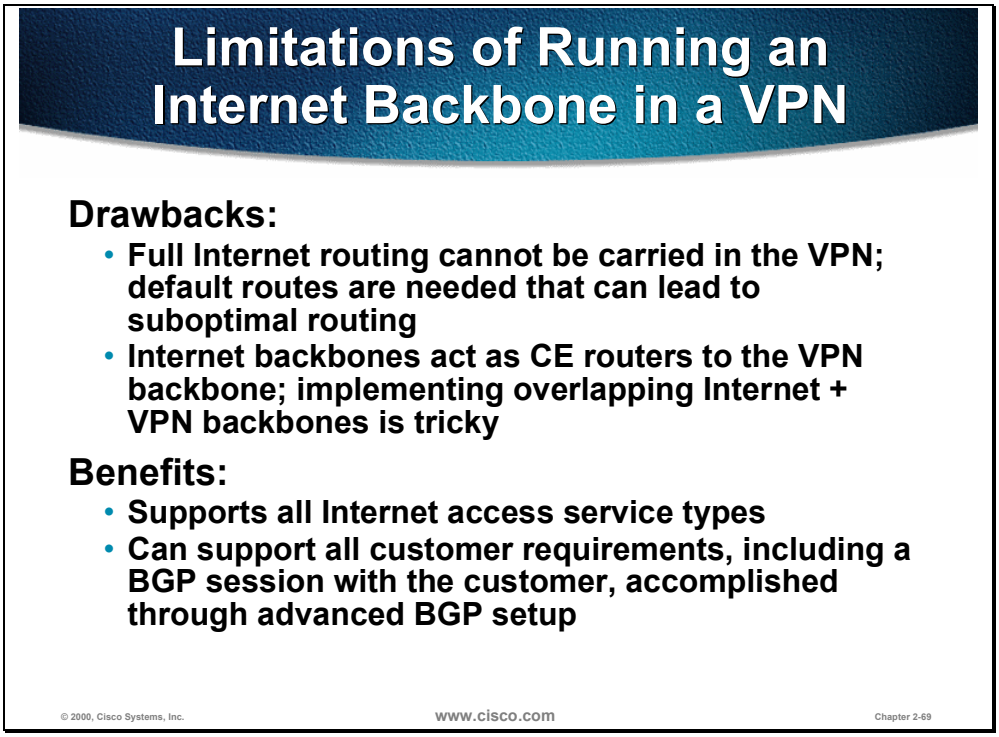

Internet access implemented as a separate VPN has a few drawbacks:

- Full Internet routing cannot be carried inside a VPN, and therefore default routing toward the Internet gateways has to be used, potentially resulting in suboptimal routing.

Note With the future MPLS VPN extensions, called **recursive VPN** or **Carrier's Carrier** model, even full Internet routing can be propagated across a VPN.

- The Internet backbone is positioned as a customer toward the MPLS VPN backbone. If the Service Provider runs Internet service and MPLS VPN service on the same set of routers, the interconnection between the two services requires special considerations.

The benefits of this design far outweigh the limitations:

- This design model supports all Internet Access services, ranging from traditional Internet access to innovative services like Wholesale Internet Access.
- It also supports all customer requirements, including full Internet routing on the customer routes.

Note A multihop EBGP session needs to be established between the customer router and the Internet gateway to propagate full Internet routing to the customer's router.

Review Questions

- What is the basic idea behind providing Internet Access through a VPN?
- Which Internet access services can be implemented by running the Internet in a separate VPN?
- How would you implement redundant Internet access when running the Internet in a VPN?
- What are the limitations of this design?

Chapter Summary

After completing this chapter, you should be able to perform the following tasks:

- Describe the requirements for Internet access from a VPN.
- Describe various design models for integrated Internet Access and their benefits and drawbacks.
- Design and implement MPLS VPN solutions based on these design models.
- Design and implement a Wholesale Internet Access solution.

7

MPLS VPN Design Guidelines

Overview

This chapter discusses various design guidelines for the MPLS/VPN backbone.

It includes the following topics:

- Backbone and PE-CE addressing scheme
- Backbone interior routing protocol selection and design
- Generic route distinguisher and route target allocation schemes
- End-to-end convergence issues

Objectives

Upon completion of this chapter, you will be able to perform the following tasks:

- Select a proper addressing scheme for the MPLS/VPN backbone.
- Select the optimal Interior Gateway Protocol.
- Develop comprehensive Route Distinguisher and Route Target Allocation Schemes.
- Design BGP in the MP-BGP backbone.
- Optimize overall network convergence.

Backbone and PE-CE Link Addressing Scheme

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Decide when to use numbered or unnumbered links.
- Decide when to use public or private IP addresses.
- Develop an addressing scheme within the backbone and between the PE and CE routers.

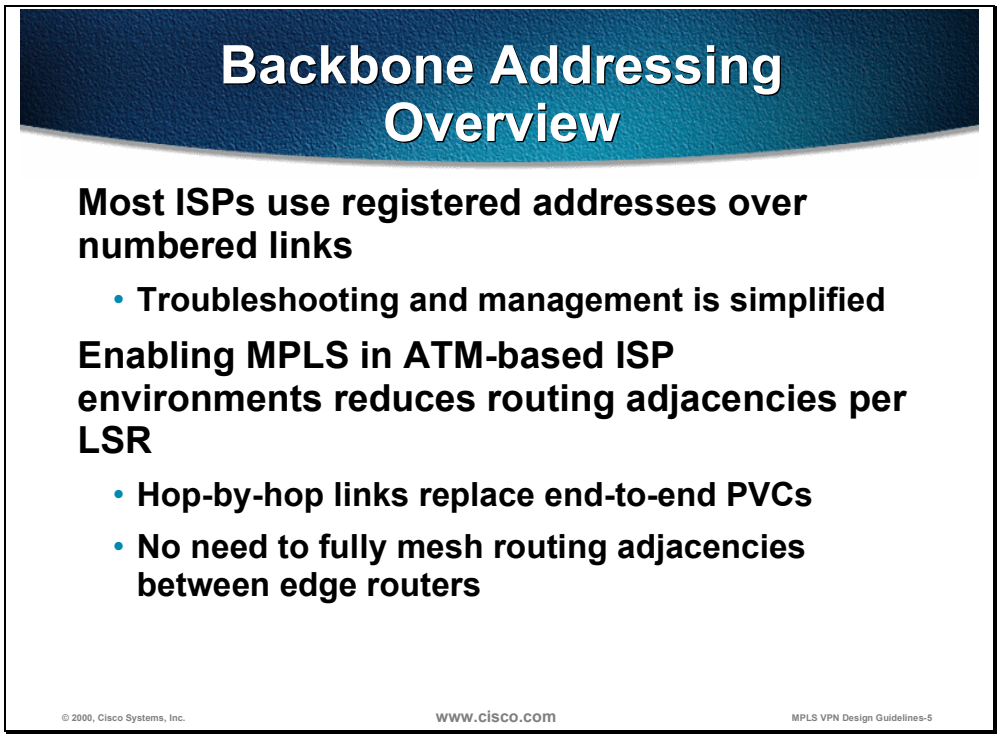

Most service providers use registered IP addresses to simplify management and to prevent traceroute across the autonomous system to show private addresses that are not accessible from outside the AS.

These IP addresses, while necessary for proper ISP backbone operation, are nonetheless wasted. The situation is even worse in ATM environments where the Service Providers have to establish a large number of point-to-point circuits across the ATM backbone, each circuit consuming an IP subnet.

Enabling MPLS in an ATM environment saves address space by removing a number of point-to-point virtual circuits that require small subnets of registered addresses. In addition MPLS seamlessly provides a full mesh between ATM-LSRs without having IP adjacencies between routers. Instead, an IP adjacency is formed between routers and MPLS-capable ATM switches.

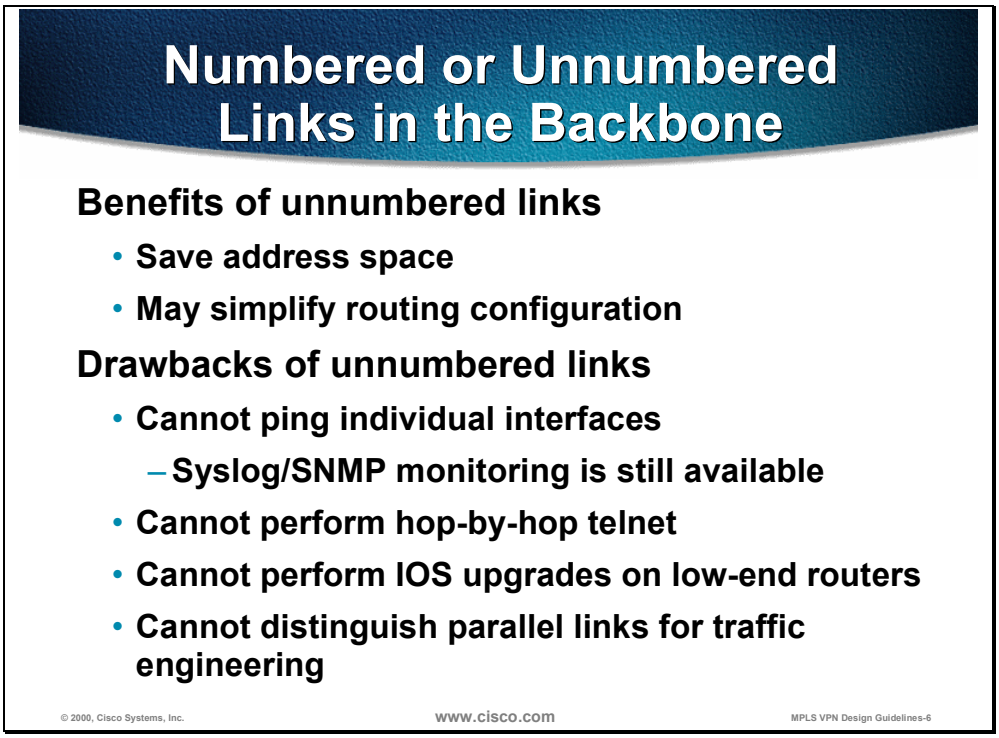

Using unnumbered interfaces results in a router having more interfaces with the same IP address. The IP address of a loopback interface is usually used on other interfaces to save address space and simplify the configuration. The downside of this approach is that the WAN interfaces on a router no longer have their own address and are therefore unreachable to ping, traceroute or telnet. However the ISP will still be able to telnet and ping the loopback address of the individual routers.

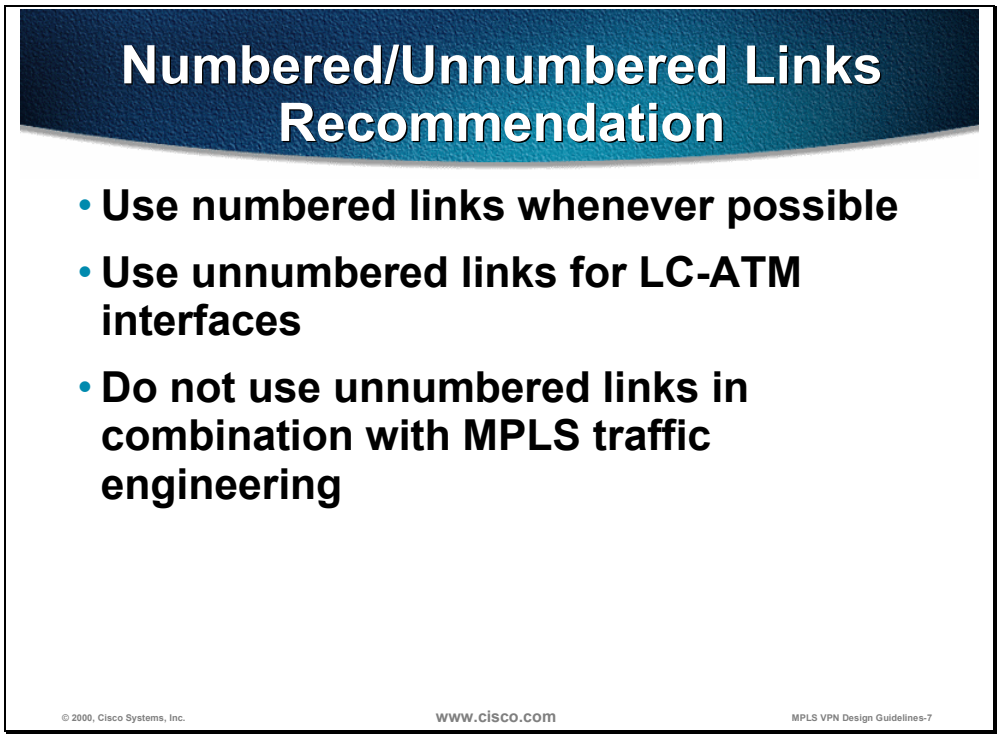

There are more benefits when using numbered interfaces. Numbered addresses should be used whenever possible except for IP adjacencies within MPLS-enabled ATM networks. In these cases, unnumbered interfaces are recommended. On the other hand, unnumbered interfaces are strongly discouraged when you use MPLS traffic engineering.

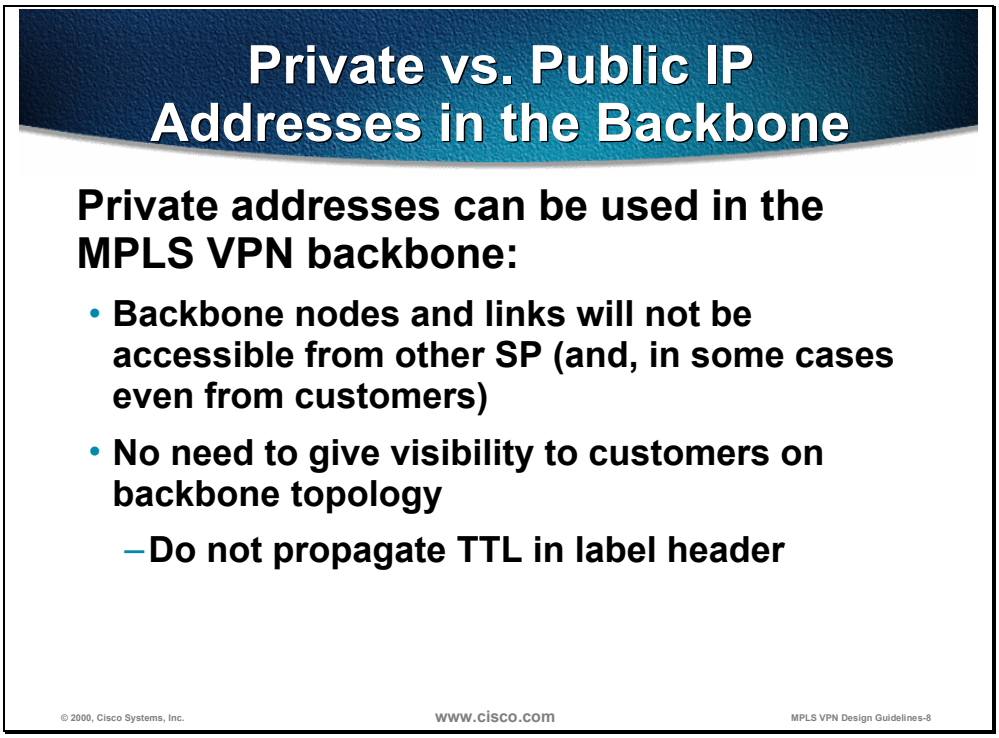

A Service Provider can decide to use private IP addresses in the MPLS core when the TTL propagation is disabled. Traceroute across a network where TTL propagation is disabled will only show the IP addresses of edge (border) routers. Core addresses, therefore, will neither be shown in traceroute nor will they be reachable from outside of the AS.

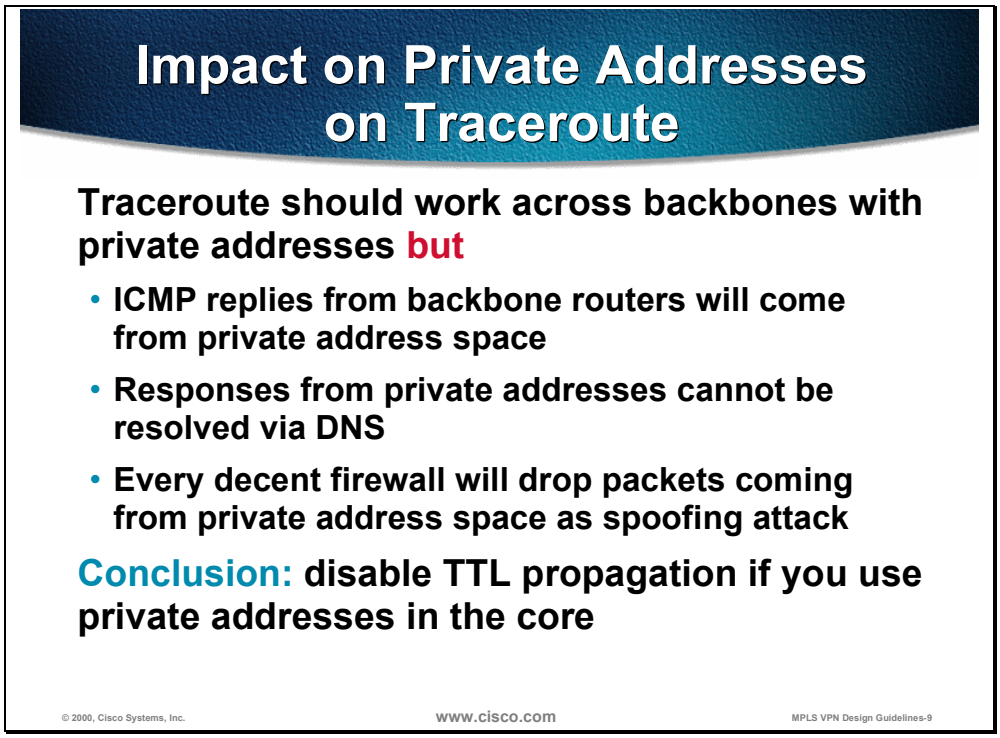

If TTL propagation is disabled, registered addresses are only used on edge (border) routers. Only these routers can send ICMP TTL-Exceeded messages. All other routers can use private IP addresses except on interfaces connecting to edge routers.

If, however, private addresses are used everywhere in the core, traceroute will show a private IP address as the source address of the ICMP reply packet. Such an address cannot be resolved via DNS. Furthermore, if traceroute is initiated from behind a firewall, it is quite likely that the return ICMP messages originating from a private IP address will not be allowed through.

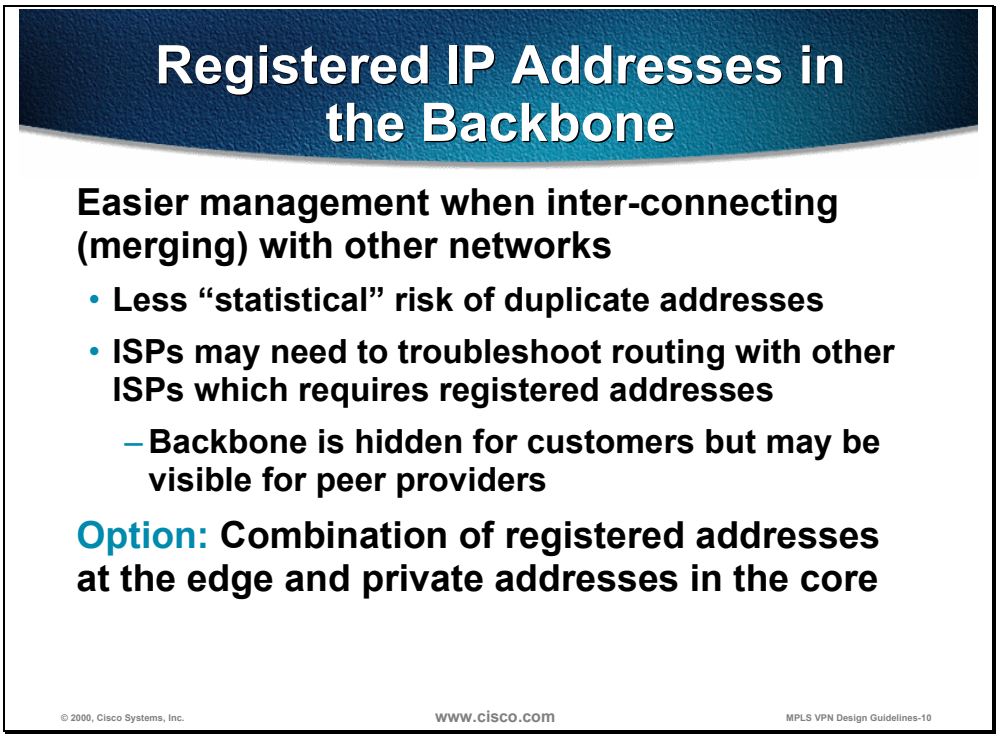

Using registered addresses is the most common practice in today's Service Provider networks.

Using registered addresses at the edge, private addresses in the core, disabling TTL propagation and only propagating labels for BGP next-hop addresses, will have the following results:

- Outside users (administrators of other ASs) can use traceroute to troubleshoot a path. They will see edge routers with registered IP address in traceroute. They will not see core routers but will be able to determine the AS where the problem is located.
- Internal users (local administrators) can use traceroute to private or registered IP addresses of LAN and WAN interfaces. Traceroute will show all core routers because those destinations are not labeled. They will be able to identify the router/link where the problem is.

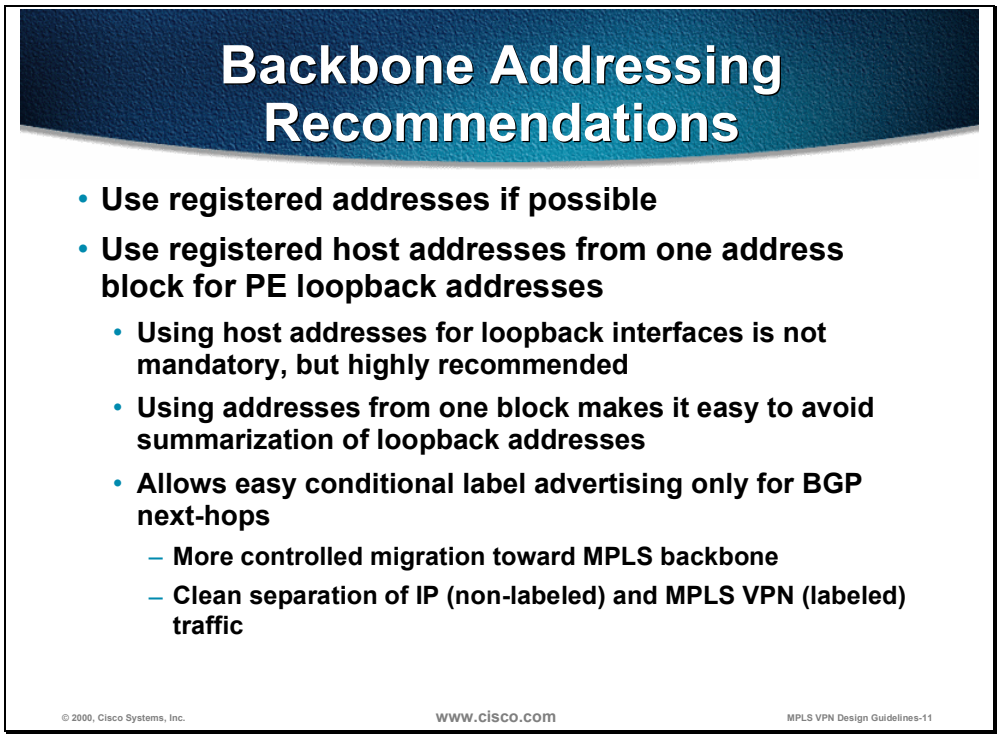

Using registered addresses only is preferred but the option of using registered and private addresses as described on the previous page can be used when running low on IP addresses.

A block of registered IP addresses should be used for loopback interfaces that are used for BGP. One host address from that block should be applied to every PE router to make it easier to exclude those addresses from summarization or to select them for labeling.

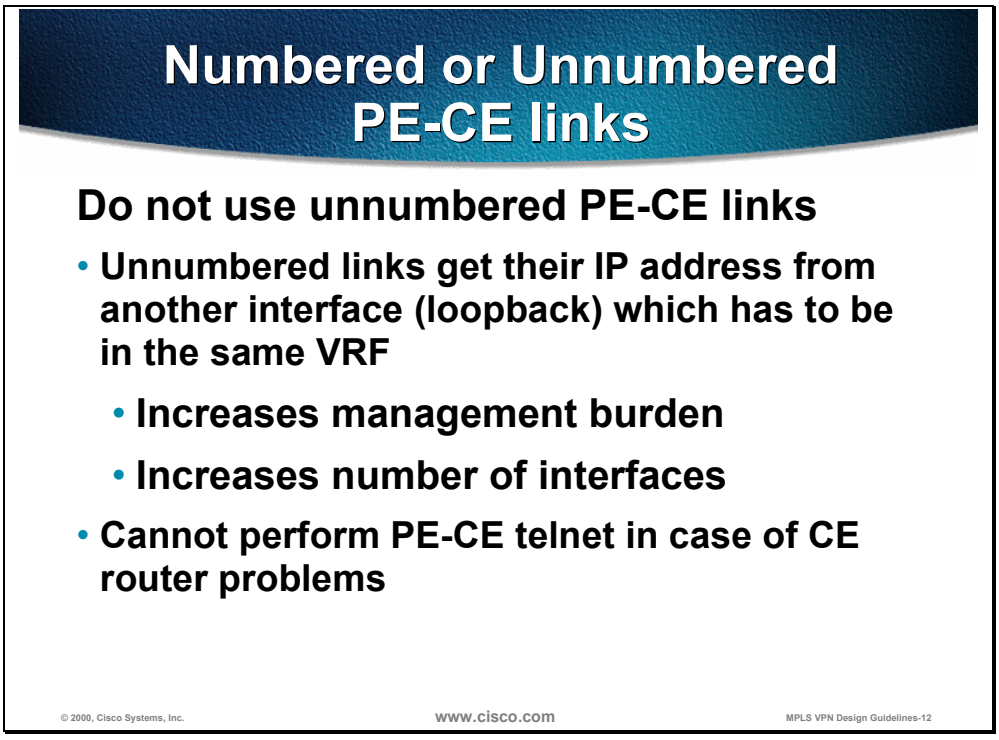

Using unnumbered VRF interfaces requires at least one loopback per VRF. Troubleshooting is more difficult since no interface is reachable either by using ping or telnet.

Using numbered VRF interfaces simplifies management and troubleshooting because every interface has its own address and can, therefore, be accessed by using ping or telnet.

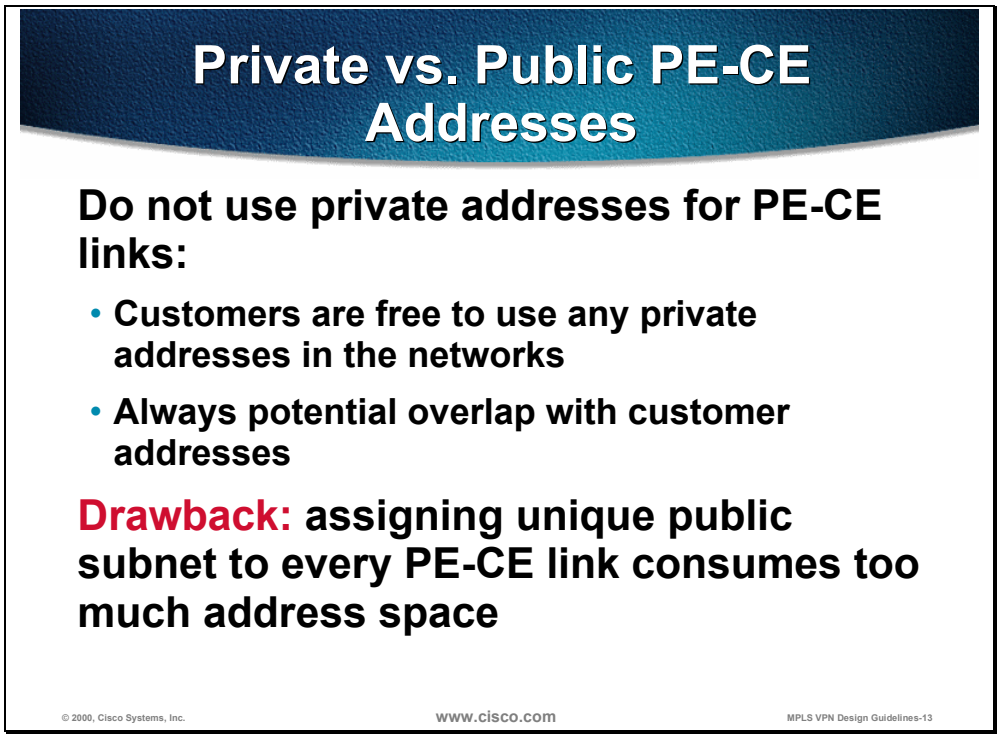

Using private addresses on PE-CE links can result in a conflict with the IP addresses used in the customer network, as the customers might already use the block of private IP addresses assigned to the PE-CE links by the Service Provider somewhere else in the customer network. There are two possible ways to prevent IP address duplication:

- Use a block of registered IP addresses for every VRF.
- Use a block of private addresses taken from the customer's address space (assigned by the customer). This approach requires tighter administrative coordination between the Service Provider and the customer.

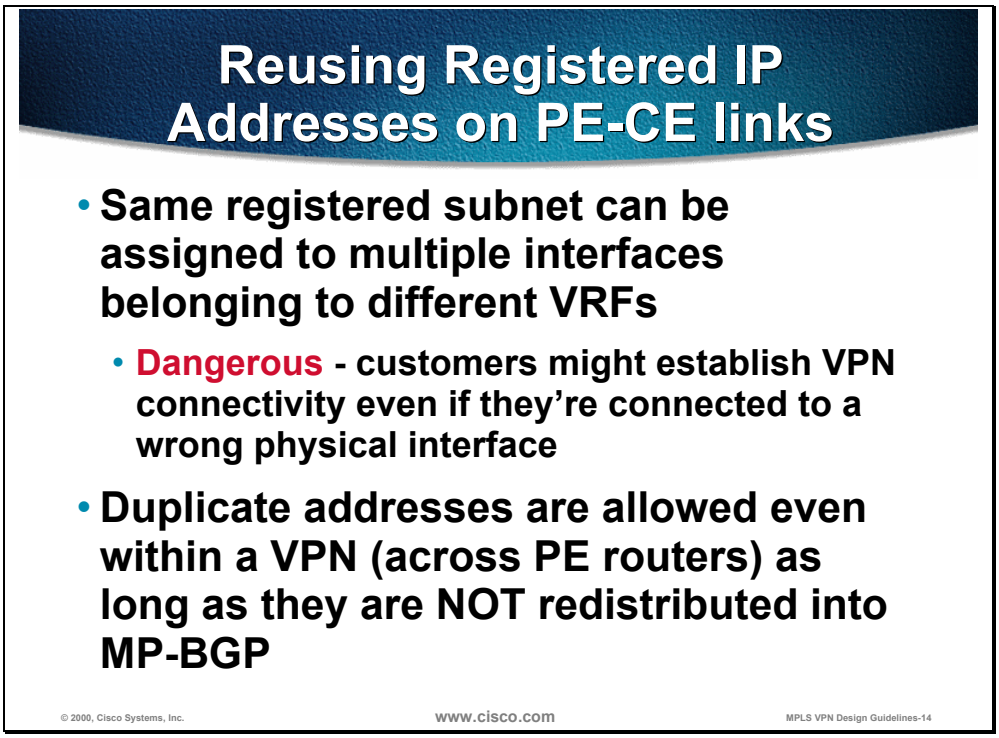

To reduce IP address consumption when registered addresses are used, reuse addresses on links belonging to different VRFs or different PE routers.

There are several options:

- Unique block of registered IP addresses for every VPN. This solution requires a large number of IP addresses.
- One block of addresses for all VPNs. If the same block is used in different VPNs, redistribute connected subnets into MP-BGP to provide reachability of all interfaces within the VPN. There will be a conflict of addresses if two or more VPNs are interconnected. This option is also dangerous from an operational perspective – if a customer site is connected to a wrong interface, the CE-router might still be able to establish connectivity with the PE-router.
- One block of addresses for all PE routers. Addresses are unique on every PE router but they are not unique within a VPN. This means that connected networks should not be redistributed into MP-BGP. The result is that PE-CE links are not reachable across several hops. If two VPNs are exchanging routing information, ensure that customers' networks are unique.
- One block of addresses for all VRFs. Addresses are not unique within a VPN nor are they unique on the PE router. This option requires the least IP addresses and has the same drawbacks as the previous option.

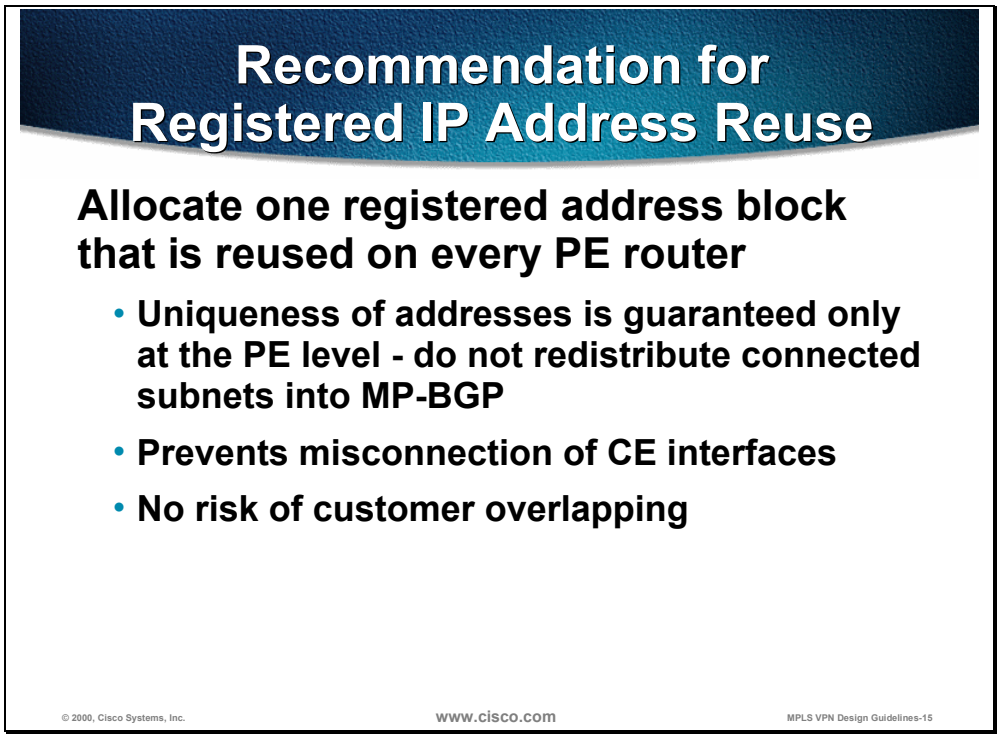

The recommended solution takes a block of registered addresses (enough to accommodate all the interfaces on the largest PE router in the network). Those addresses are reused for every PE router. They are, however, unique on a PE regardless of the VRF to which the interface belongs.

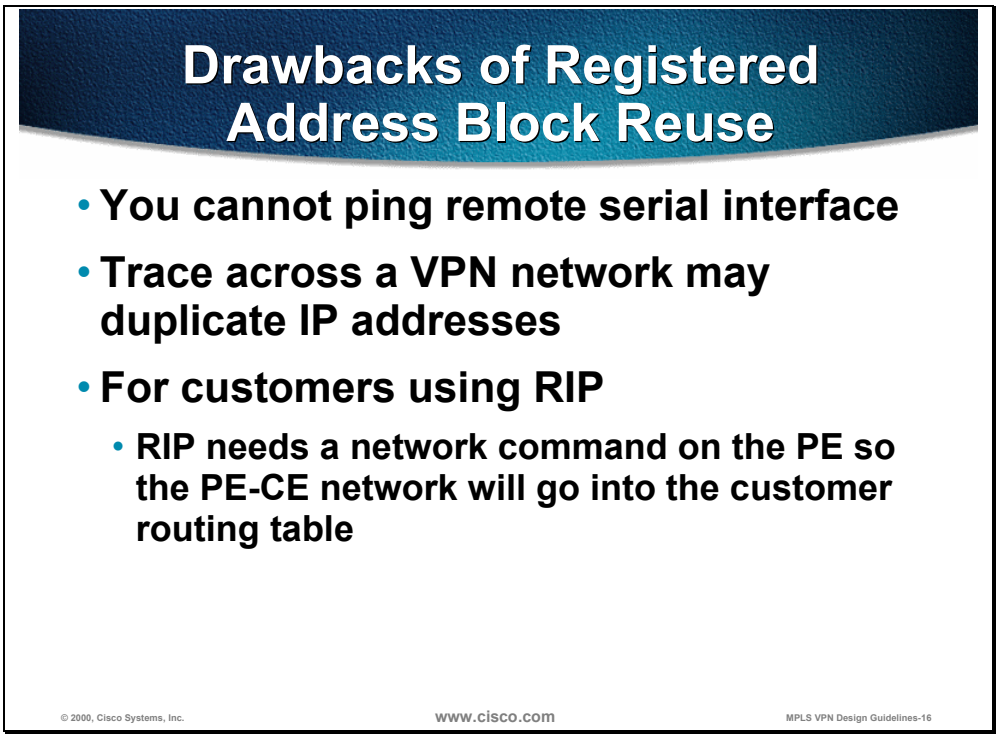

When IP addresses are reused on PE-CE links they should not be redistributed into MP-BGP. Those addresses are then unreachable and cannot be pinged from remote locations. The other result is that the same address may appear several times when performing traceroute to different destinations reachable through different PE routers.

Summary

This section described a variety of possibilities when designing IP addressing of PE-CE links.

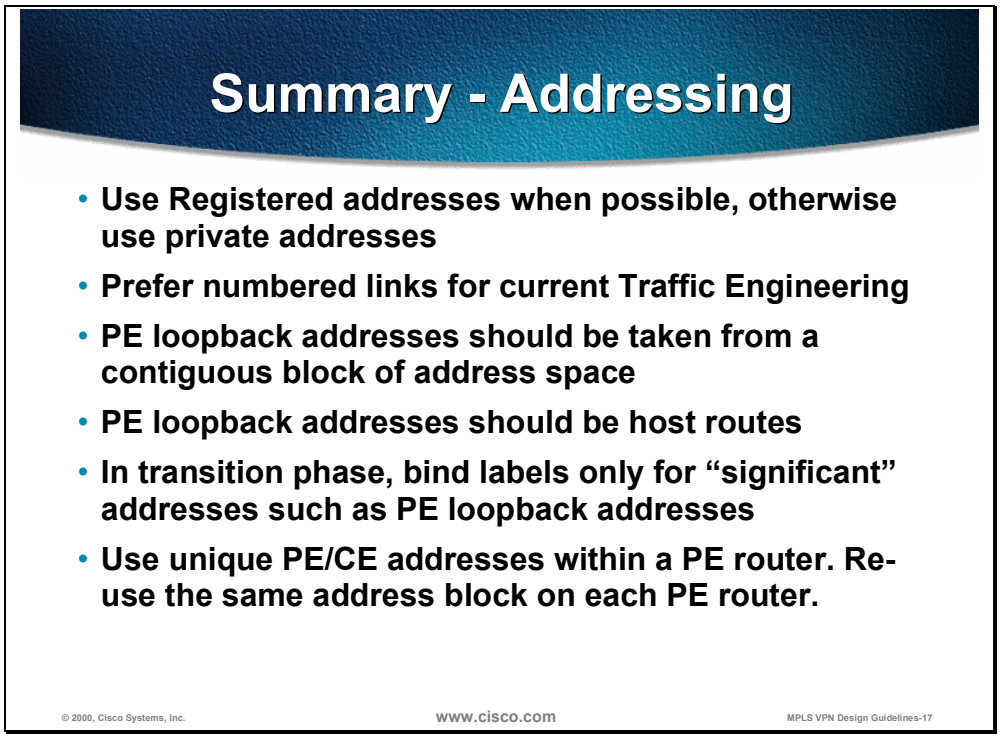

The preferred solution is to use numbered interfaces with registered addresses whenever possible. One can, however, user private addresses in the core and reuse registered addresses on PE-CE links to minimize the number of registered addresses needed for designing an MPLS/VPN network.

Review Questions

Answer the following questions:

- What are the drawbacks of using unnumbered links?
- Where should you use unnumbered links in the MPLS backbone?
- Where would you use unnumbered links between PE and CE routers?
- Why would you use private address space in your IP backbone?
- What are the drawbacks of using private address space in your IP backbone?
- How would you hide the private address space from your customers?
- What is the impact of using private backbone addresses on traceroute?
- Why should you allocate PE loopback addresses from a separate address block?
- Why should you use registered addresses for PE-CE links?
- Why is the reuse of registered addresses between VRFs not advisable?
- When can you reuse registered addresses in the same VPN between PE routers?

Backbone IGP Selection and Design

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Select the proper IGP to run in the backbone.
- Design the selected IGP to meet MPLS/VPN requirements.

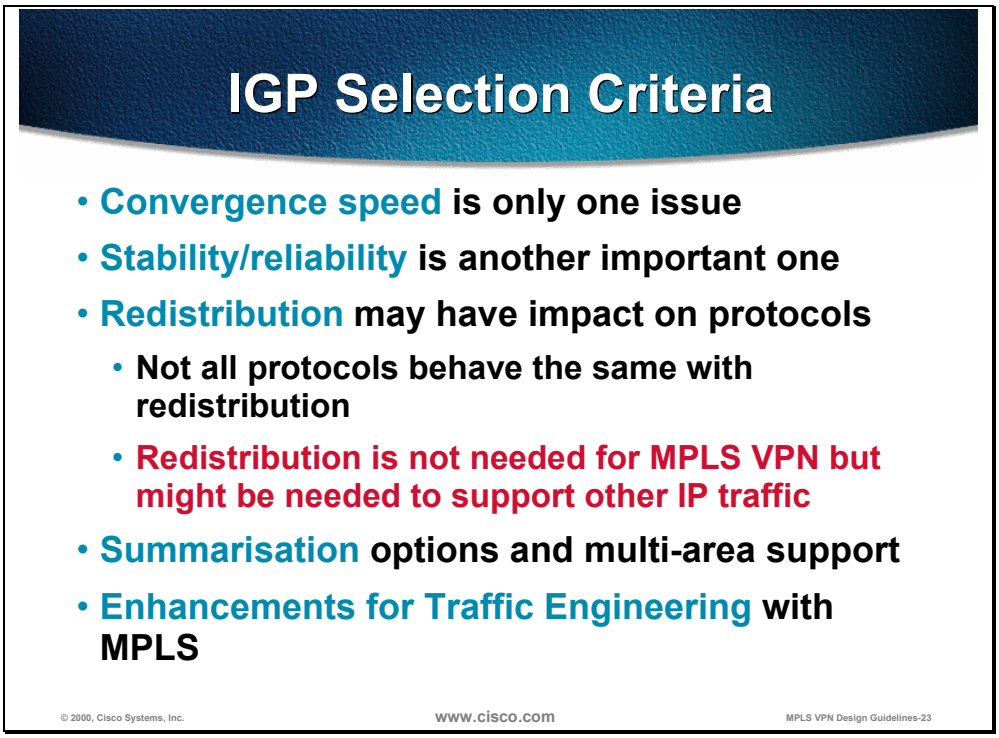

An MPLS/VPN network is generally not affected by the IGP that is used in the core. The criteria for choosing IGP are the same as for any Service Provider network.

IGP should be a balance of fast convergence, stability and scalability. Stability and scalability are also improved by the ability of summarizing networks. Summarization options and multi-area support are also very important selection criteria.

The only constraint when choosing IGP is if MPLS Traffic Engineering (MPLS TE) is planned for the network. In that case IS-IS and OSPF are the only available routing protocols supporting TE.

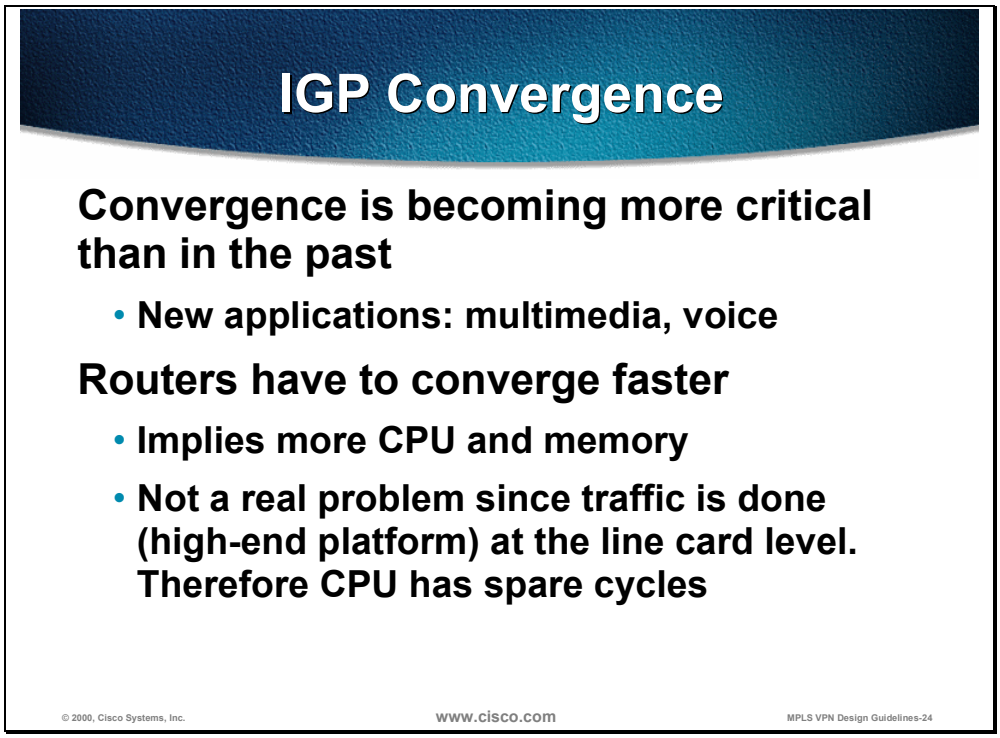

IGP convergence speed is only one in a number of factors that affect convergence across an MPLS/VPN network. Choosing the right IGP may improve overall convergence. Since most high-end routers distribute the switching task to VIPs or line cards, there is enough CPU power left for routing protocol calculations without impact on switching performance.

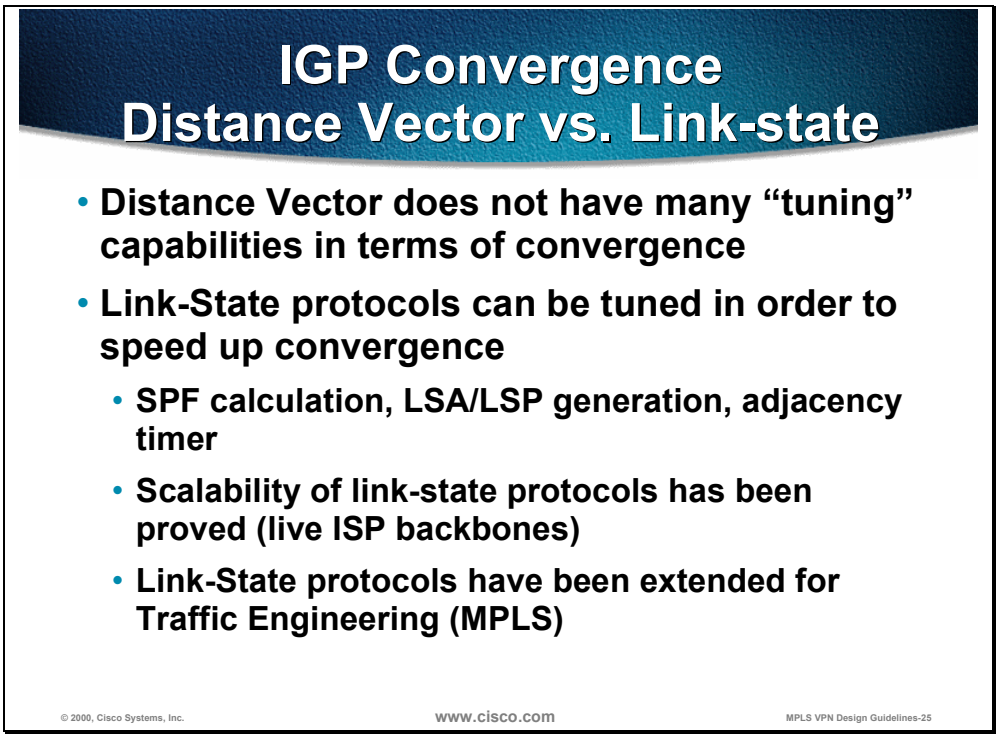

When comparing well-known distance-vector and link state protocols, there are more benefits in using the latter one. Although link-state protocols typically require more CPU, they have more tuning options to set up the protocol to the needs of a specific network.

Link-state protocols also contain the topology of the network, which is required for MPLS Traffic Engineering. IS-IS and OSPF (both link-state protocols) have been extended to support the requirements of Traffic Engineering. If the need to implement Traffic Engineering in the future is foreseen, it is better to initially use one of these two protocols in the MPLS backbone.

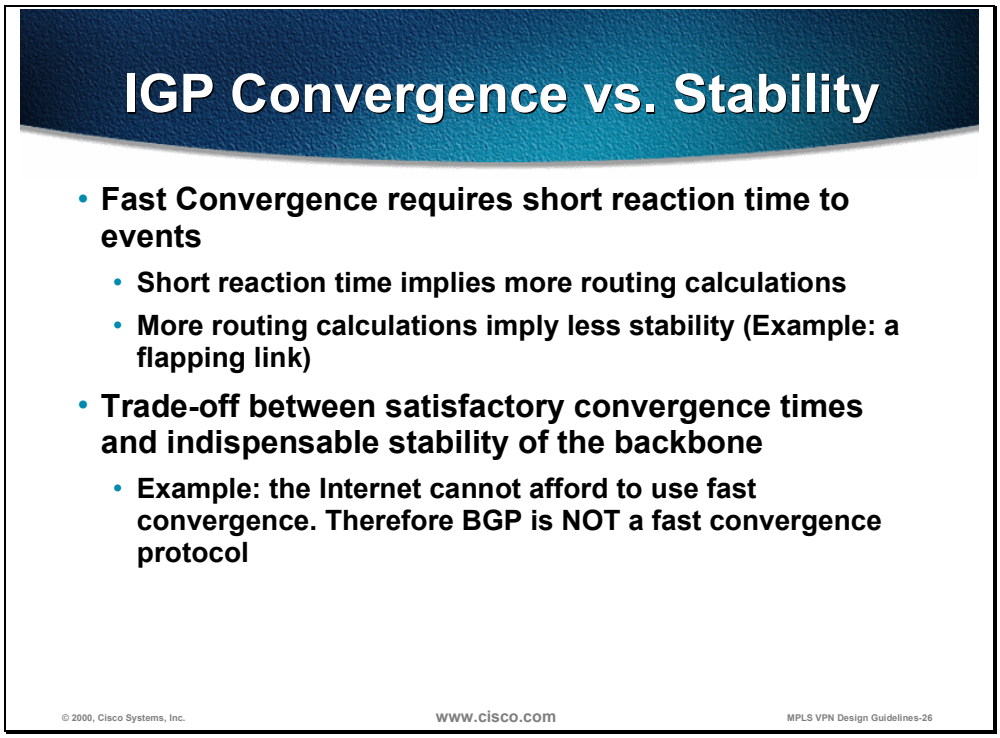

When striving to maximize convergence the result may be a very unstable network. For instance, assume a router immediately sends an update when something changes and the receiving router immediately forwards the information and recalculates the best paths. However, if a number of updates are being sent, the router will recalculate its routing table each time it receives an update. In this example it is also quite likely that the CPU will need much more time to perform all calculations than if it waited to receive more updates and then perform the calculations. A flapping link, as another example, will cause recalculations every time it flaps. Deliberately slowing convergence (i.e. not recalculating best paths immediately) will have a positive effect on stability of the network since there is less chance of routers' CPUs being overloaded.

This is especially important for routers in the Internet where there are a large number of networks and different paths (at the time of writing there were almost 100.000 networks and up to 500.000 different paths in some exchange points). This is the reason why BGP, which is used by Service Providers for interdomain routing, is intentionally slowed down. When changes are happening in the network (there is hardly ever a moment in the Internet when they are not), BGP will send updates every 5 seconds to its internal neighbors and every 30 seconds to its external neighbors. A link that is flapping once a second will appear to be flapping at a maximum rate of once every 30 seconds to someone on the other side of the globe. These mechanisms, however, are not used for IGPs where the number of networks is smaller and a faster convergence is needed.

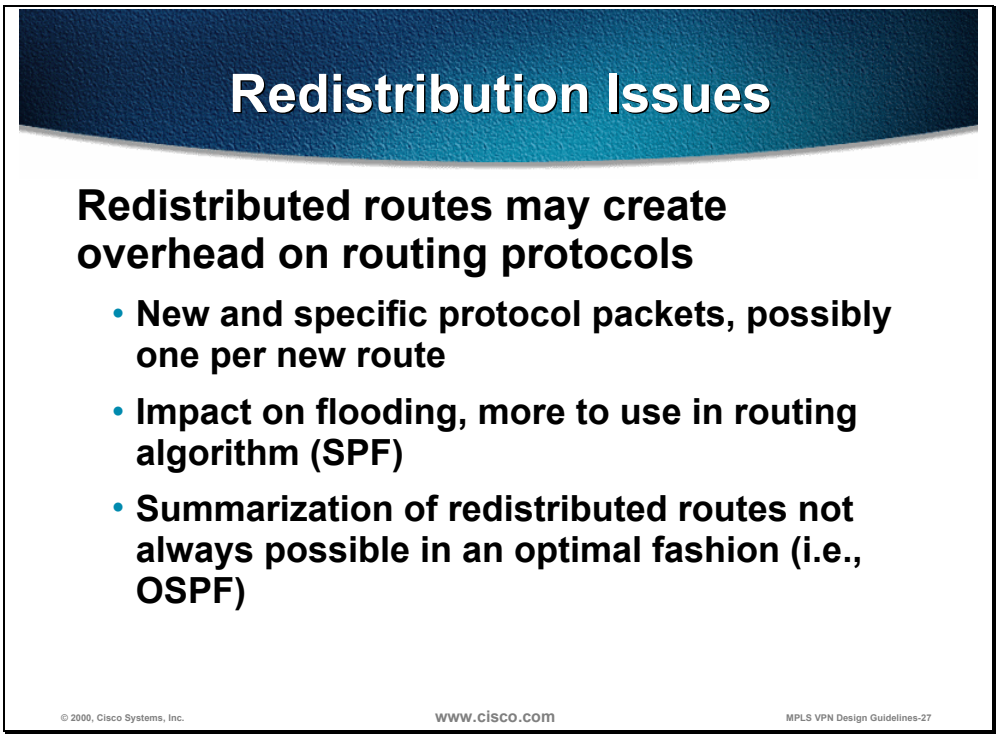

Using redistribution is usually regarded as a quick way to insert routing information into the IGP database and send it to router's neighbors. The result may be too much routing information in the memory of the core routers and the calculations of best paths may take longer because of that. Most protocols, however, allow for subsequent summarization of routing information. The only exception is OSPF where redistributed networks may not always be summarized

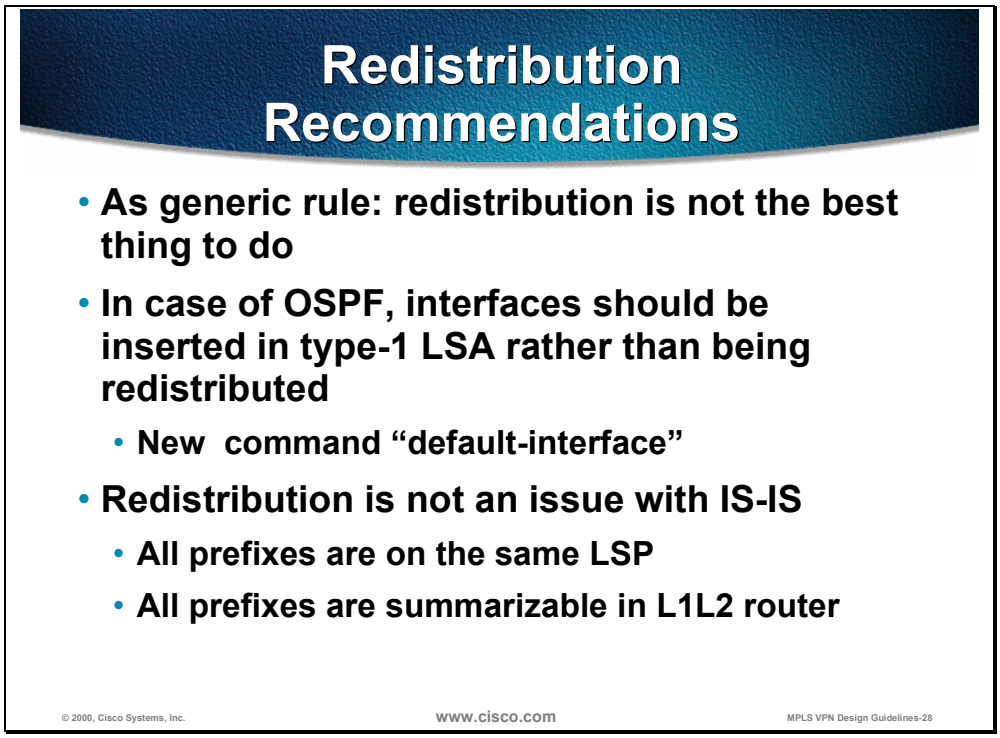

For the reasons shown on the previous slide, redistribution should be avoided when possible. If, however, redistribution of connected subnets into a routing protocol is necessary, they should be included in the routing protocol definition. In this case, the **passive-interface default** command should be used to prevent IGP from running on any interface where it has not been explicitly enabled.

When including a connected subnet in an OSPF routing process, OSPF creates type-1 Link State Advertisements (LSAs) that can later be summarized regardless of the type of area where they originate. There are no such drawbacks when using other IGPs such as IS-IS or EIGRP.

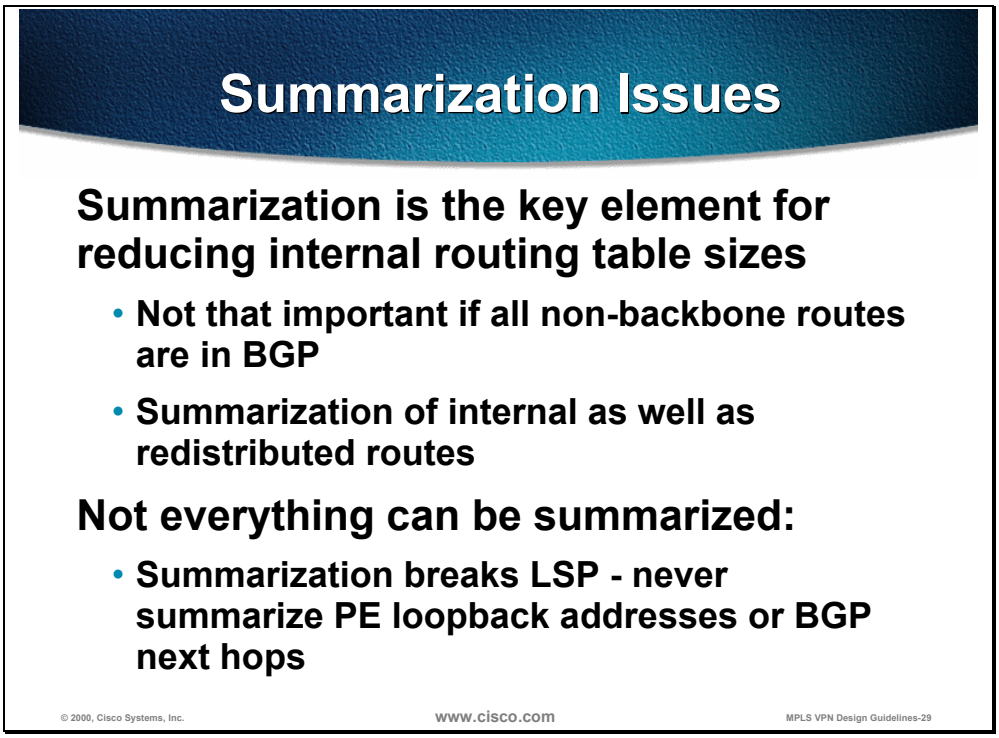

Having good summarization capabilities is an important feature of IGP. There are a growing number of extremely large networks in the world that have scalability problems because they have too many networks and too many routers. Having a good IP addressing scheme is a necessity to minimize the amount of routing information and to make the network more stable (i.e. flapping links are hidden in summaries and do not cause constant recalculations). When using OSPF, ensure that redistributed networks are also being summarized.

There is a very important issue to consider when using summarization in an MPLS/VPN network. VPNs only work if the MPLS core provides unbroken Label Switched Path (LSP) between all PE routers. Summarizing addresses of loopback interfaces, which are used for MP-BGP peering, will cause the LSPs to those loopbacks to break in two and that subsequently causes VPNs to break apart. Therefore, always exclude loopback addresses from summarization in backbone IGP.

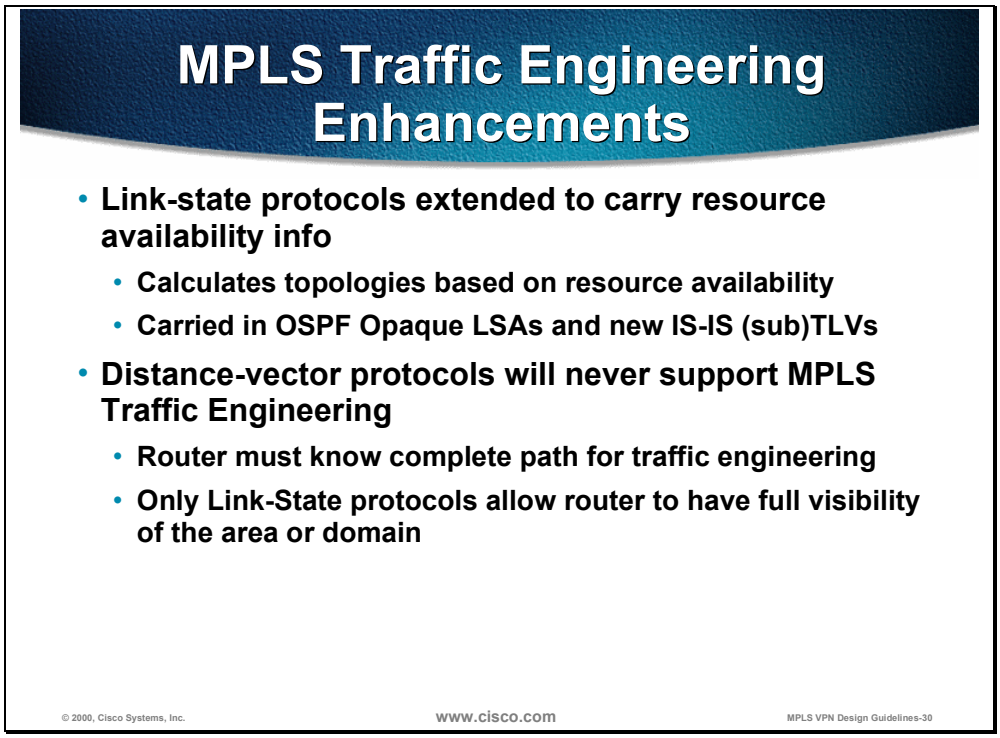

For the purpose of implementing a Traffic Engineering mechanism OSPF and IS-IS were extended to carry some additional information (available resources and constraints of links in the network). These are the only two protocols that already carry the information about individual links and hold the entire topology of an area in its database.

When using Traffic Engineering, therefore, the only choice of protocol is between OSPF and IS-IS. There will never be an implementation of EIGRP to support Traffic Engineering because it simply does not carry the link information and holds no real topology information.

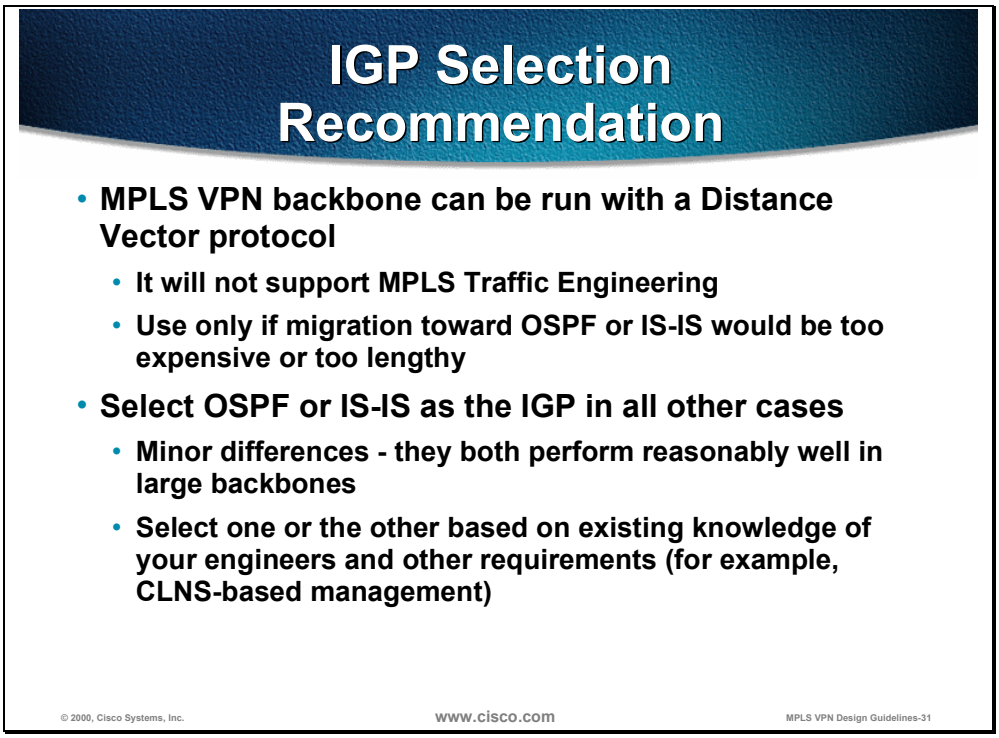

Although MPLS and MPLS VPNs work with any IGP (OSPF, IS-IS, IGRP, EIGRP, RIPv2) only OPSF and IS-IS support Traffic Engineering. Choosing one of these two protocols may be the best decision even if Traffic engineering is not presently planned – it may be in the future.

The choice between the two protocols is usually based on the user's familiarity with one over the other, as their performance is similar.

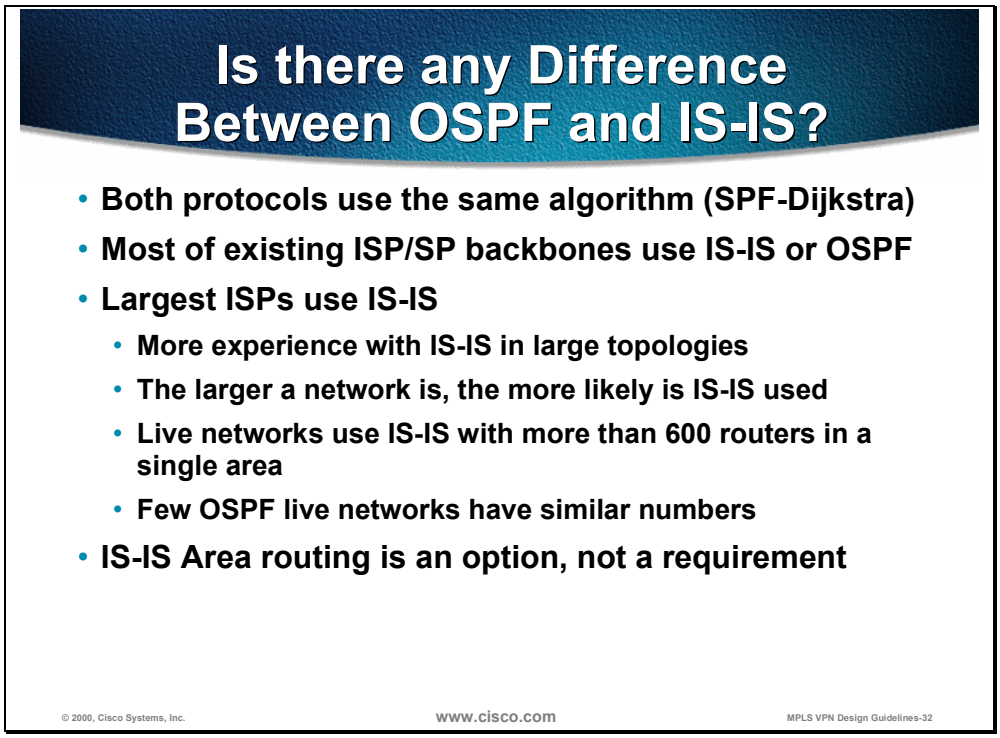

The slide above shows that there are hardly any differences between the two protocols although there are more large networks using IS-IS than OSPF.

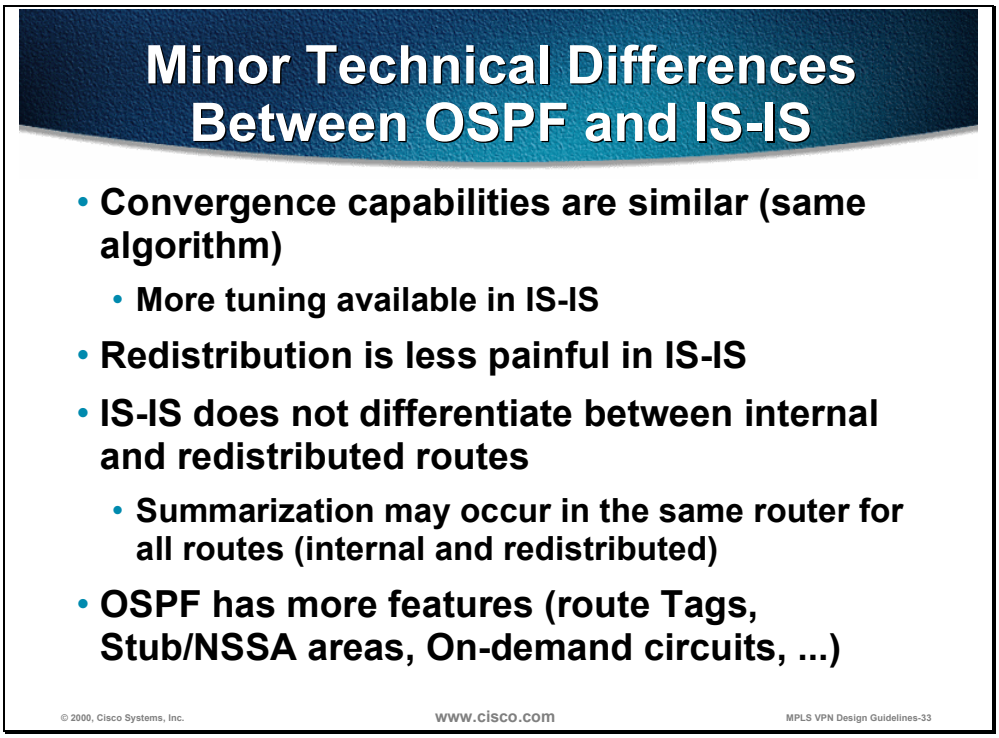

In considering Cisco IOS configuration, IS-IS has more tuning options and is not affected by combining redistribution and summarization. OSPF, on the other hand, has more features.

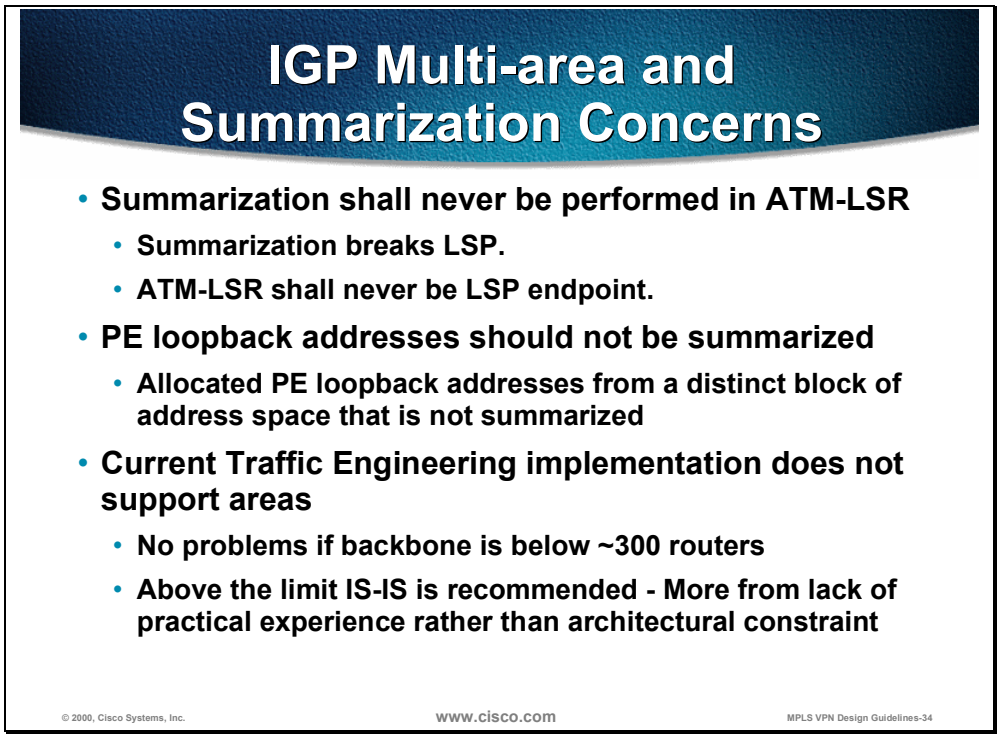

When performing summarization, remember not to summarize PE loopback addresses that are used as BGP next-hop addresses. Do not perform summarization on ATM-LSRs because it breaks the LSP and ATM-LSRs are not capable of IP forwarding.

Traffic Engineering requires a full overview of the topology of the network where Traffic Engineering is to be used. Currently this is only possible if there is only one area in the OSPF or IS-IS implementation.

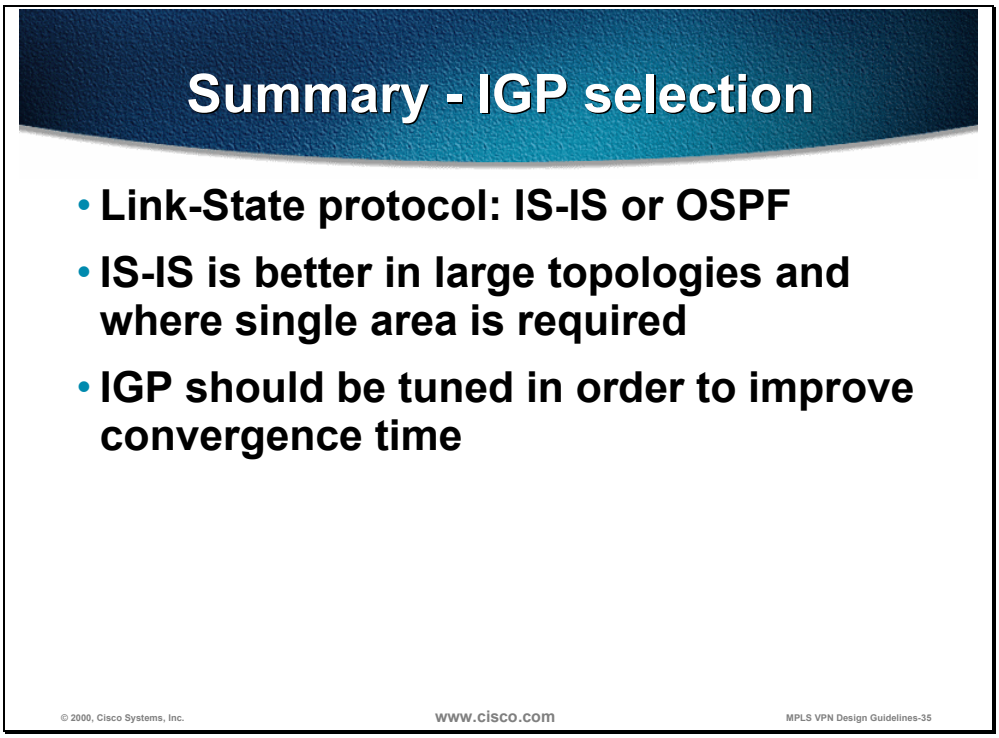

This section described major factors to be taken into account when selecting the right IGP for an MPLS/VPN backbone. These factors are:

- Convergence vs. stability
- Impact of redistribution
- Scalability (summarization capabilities) and multi-area support
- Support for Traffic Engineering

The choice is usually between OSPF and IS-IS.

Review Questions

- List three IGP selection criteria.
- What is the impact of higher convergence speed on network stability?
- How can you tune OSPF convergence?
- How can you tune IS-IS convergence?
- What is the difference between OSPF and IS-IS route redistribution?
- Where can you summarize redistributed routes in OSPF?
- Where can you summarize redistributed routes in IS-IS?
- How do you avoid redistribution of connected interfaces when using OSPF?
- Which routing protocols support MPLS Traffic Engineering?
- Why is MPLS TE not supported by EIGRP?
- When can you use EIGRP as the IGP protocol in your MPLS/VPN backbone?
- What is the impact of route summarization on MPLS/VPN?
- Why is IS-IS recommended for extremely large networks?

Route Distinguisher and Route Target Allocation Schemes

Objective

Upon completion of this section, you will be able to develop generic Route Distinguisher (RD) and Route Target (RT) allocation schemes

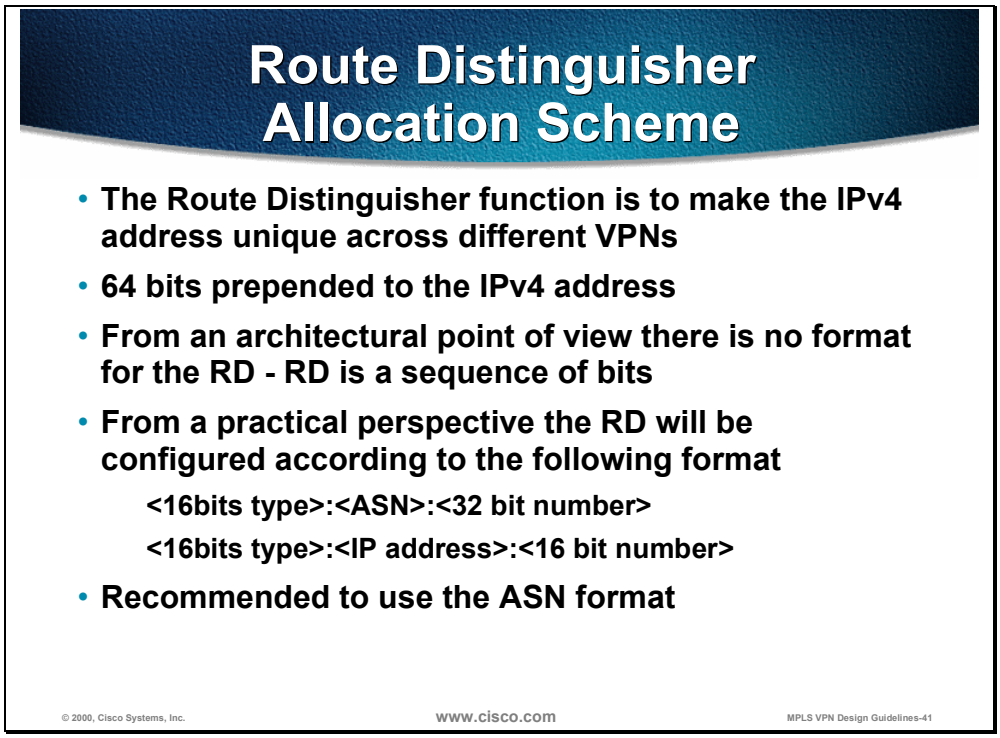

MPLS/VPNs support overlapping addresses in different VPNs. On the other hand, PE routers run one single instance of RIP and BGP. To make sure BGP can distinguish between network 10.0.0.0 belonging to VPN A and the same network belonging to VPN B (which is in reality a different network, as it belongs to private address space of another customer), an additional value is required – Route Distinguisher (RD).

Route Distinguisher can be specified in two different formats:

- A 16-bit Autonomous System number followed by a colon and an arbitrary 32-bit number (e.g. 1:100)
- A 32-bit unique (registered) IP address followed by a colon and a 16-bit arbitrary number (e.g. 1.2.3.4:100)

Combining a Route Distinguisher with the IPv4 address creates a unique prefix (VPN IPv4 address):

- AS format example: 1:100:10.0.0.0/8
- IP format example: 1.2.3.4:100:10.0.0.0/8

The two networks 10.0.0.0 in our example are now different:

- VPN A: 100:100:10.0.0.0/8
- VPN B: 100:200:10.0.0.0/8

A routing protocol such as BGP will no longer recognize the two networks as the same and will forward both networks to its neighbors.

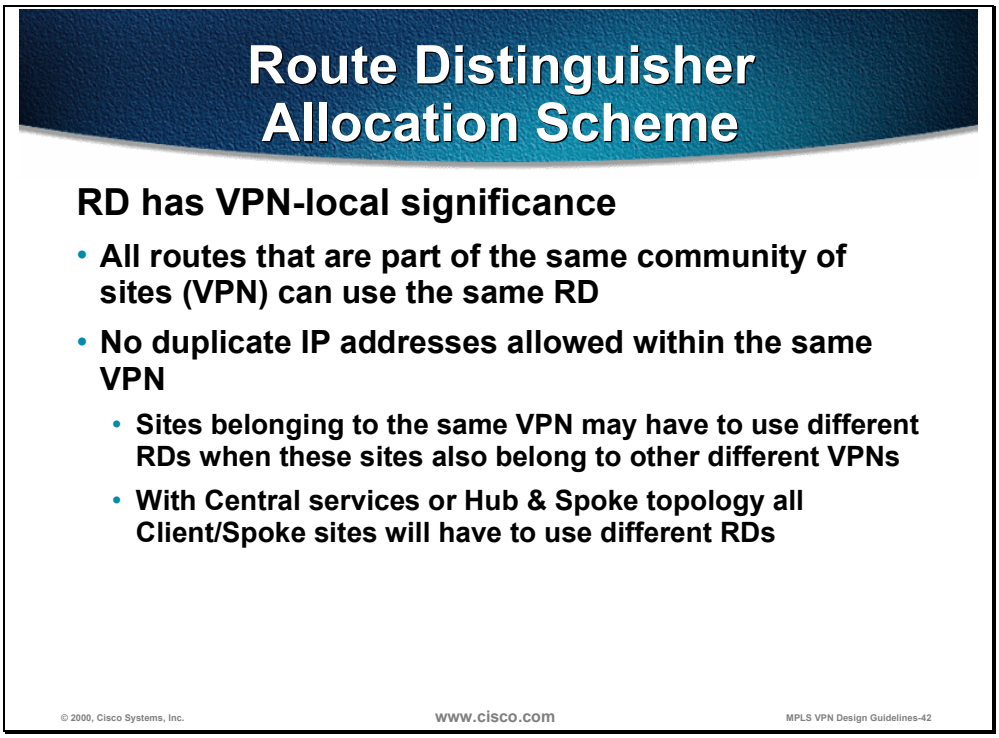

For VRFs that are used for the same VPN, one can use the same Route Distinguisher on all PEs. When using more than one VRF for the same VPN on one PE router, it is necessary to use more Route Distinguisher values. This is the case when more complex VPN designs are used, such as overlapping VPNs, Central Services VPN, Management VPN, Hub&Spoke topology.

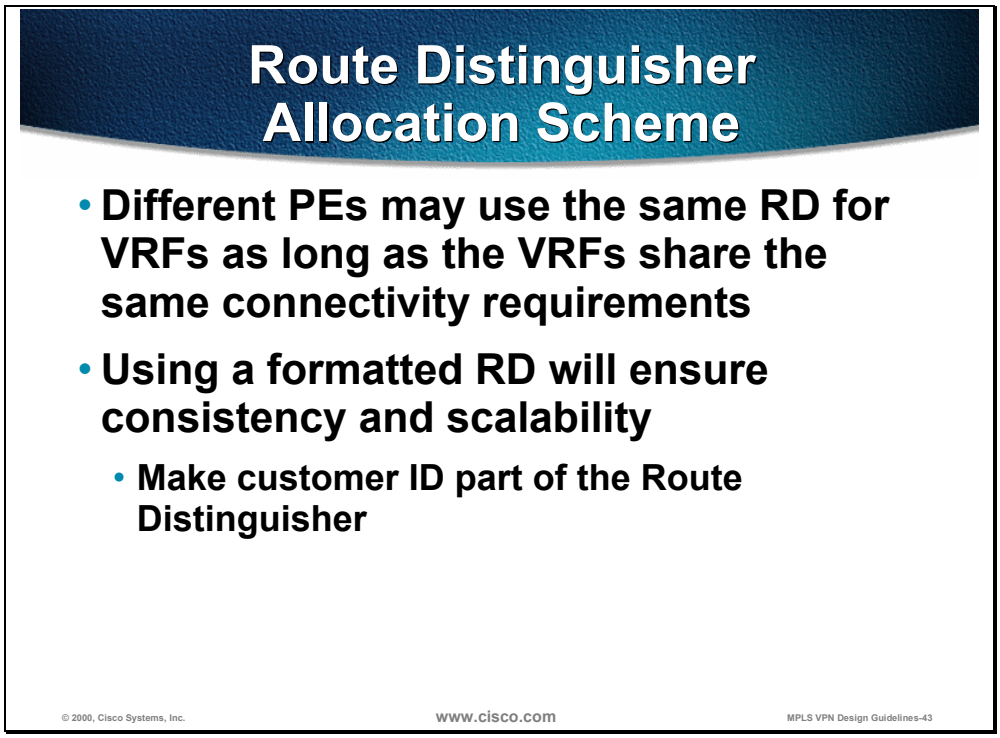

As previously stated, RD values can and should be reused only if the VRFs on different PE routers are used for the same VPN. RDs should be globally unique in all other cases.

Having a good numbering scheme for RD values means that there are initially more values reserved for each customer VPN and that the Customer ID is part of the RD. Here is a sample spreadsheet that can be used for RD numbering:

	Customer	Route Distinguisher range
Internal use	Management VPN	100:100-100:199
	Internet VPN	100:200-100:299
	…	…
Customers	Global Motors	100:1000-100:1099
	Bolts&Nuts	100:1100-100:1199
	…	…

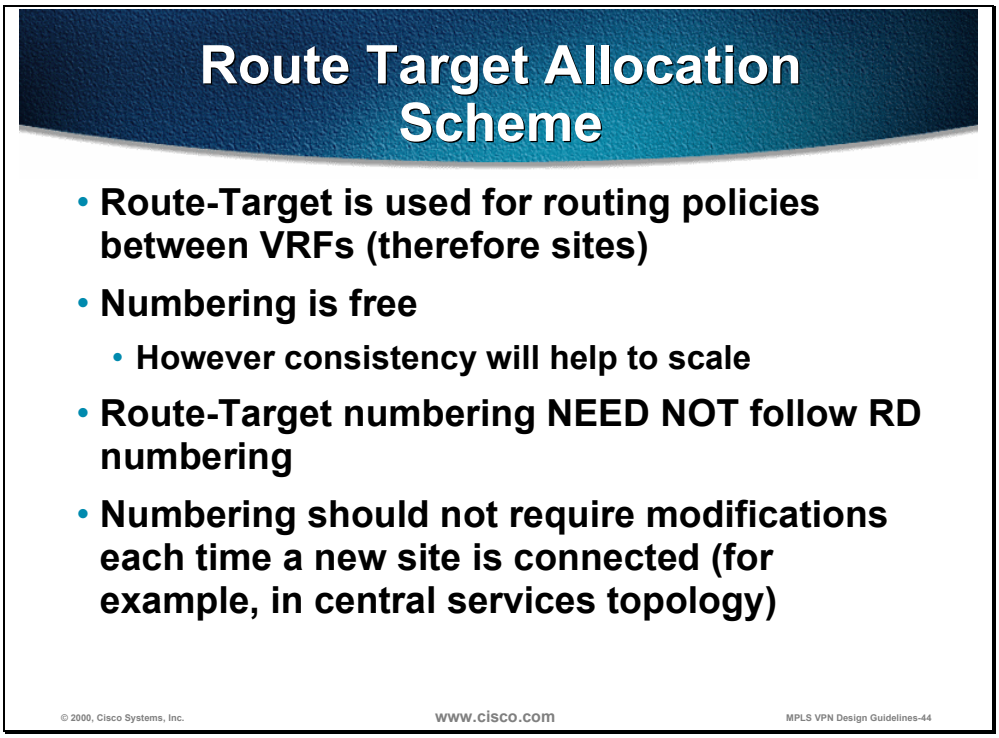

Although Route Target is used for different purposes, the same numbering scheme can be used. A range of numbers should be reserved for each VPN. The previous example has been expanded to include the RT numbering scheme:

	Customer	Route Target range	Route Distinguisher range
Internal use	Mgmt. VPN	100:100-100:199	100:100-100:199
	Internet VPN	100:200-100:299	100:200-100:299
	…	…	…
Customers	Global Motors	100:1000-100:1099	100:1000-100:1099
	Bolts&Nuts	100:1100-100:1199	100:1100-100:1199
	…	…	…

Summary

This section described the Route Distinguisher and Route Target numbering options and made recommendations for their allocation. A numbering plan for Route Targets and Route Distinguishers should be a part of any MPLS/VPN design document. A good numbering scheme may ease troubleshooting in an MPLS/VPN network.

Review Questions

- What is the function of the route distinguisher?
- Can you reuse the same route distinguisher on different PE routers?
- Is there any topology where every site requires a different value of route distinguisher?
- What is the function of the route target?
- Do you have to make the route target equal to the route distinguisher?

End-to-End Convergence Issues

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Explain the difference between overlay VPN convergence and MPLS/VPN convergence.
- List the elements of end-to-end convergence in the MPLS/VPN network.
- Optimize individual elements of MPLS/VPN convergence.

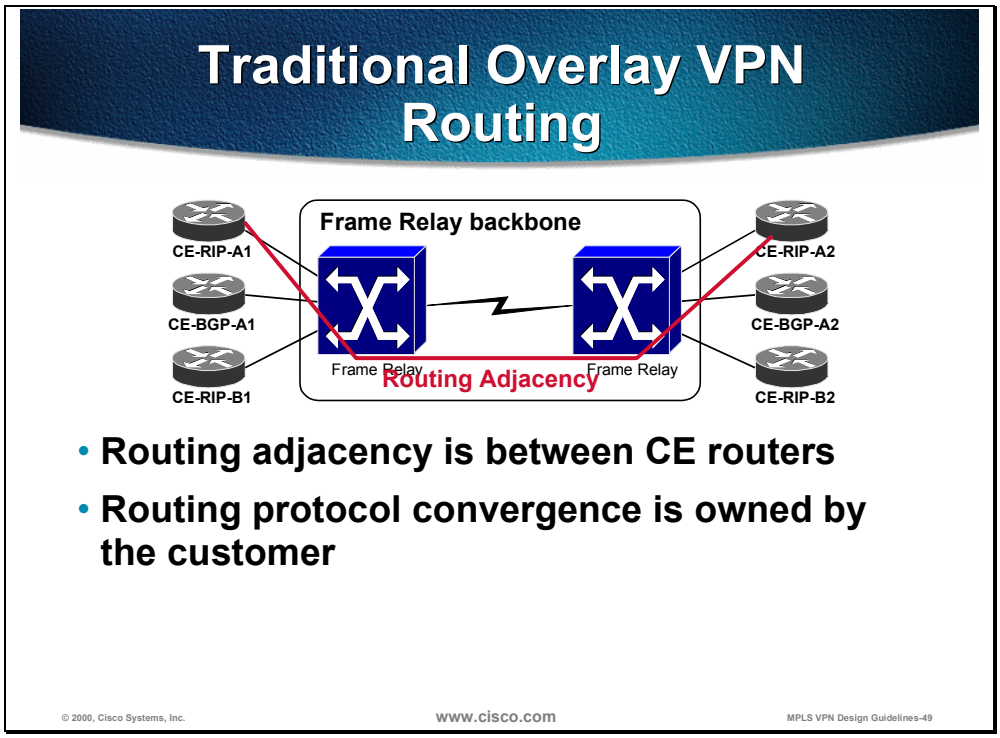

In traditional overlay VPNs a PE device works on layer 2 or 1 and does not delay IP routing updates flowing between CE devices. Routing convergence is therefore influenced by the routing protocol running directly between the CE devices. There is, however, an impact of the Layer 2 Service Provider infrastructure on convergence. The infrastructure of the SP should be able to inform CE devices of a failure in its network. This is usually done through signaling on Layer 2 or 1 but it can take time or it may not happen at all (frame relay keepalives, for example, work between a router and a switch; failure somewhere in the frame relay network may not be signaled to the CE router).

Traditional Overlay VPN Convergence

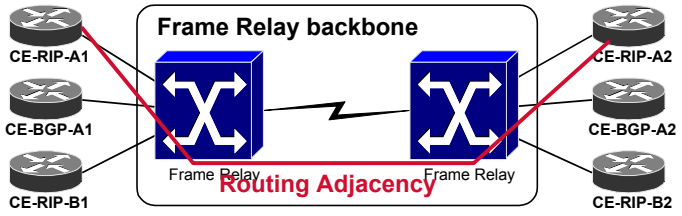

Elements of overlay VPN convergence:

- **Neighbor loss discovery (usually not immediate but based on dead timer)……..up to 40 seconds**
- **Propagation of changed routing information……...few seconds**
- **Topology recomputation……5 - 15 seconds**

All these elements can be tuned resulting in very fast convergence

The slide above shows the estimated time it takes from an actual failure to the moment the routers discover something is wrong, propagate changes to their neighbors and recalculate the shortest paths. The whole process may take anywhere between a second to over a minute, depending on the type of failure, Layer 2 infrastructure and the chosen IGP.

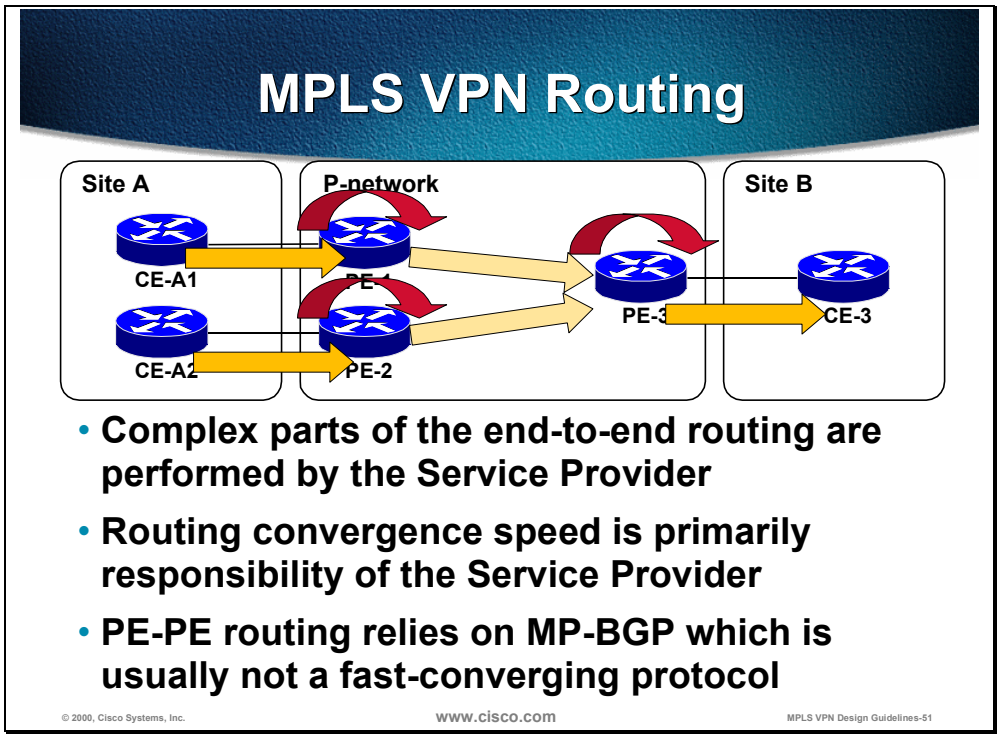

With MPLS/VPN networks the Service Provider is actively involved in network layer of the customer network (layer 3 of the OSI model). The customer's CE devices are no longer peering with other CE devices; they exchange routing updates with the provider devices (PE-routers). There can also be a sequence of PE devices exchanging the customer's routing information between customer sites.

The customer's routing information is redistributed into BGP, which is inherently slower than most IGPs. The overall convergence is, therefore, also influenced by multi-protocol BGP convergence speed and its fine-tuning. This topic is covered in the remainder of this section.

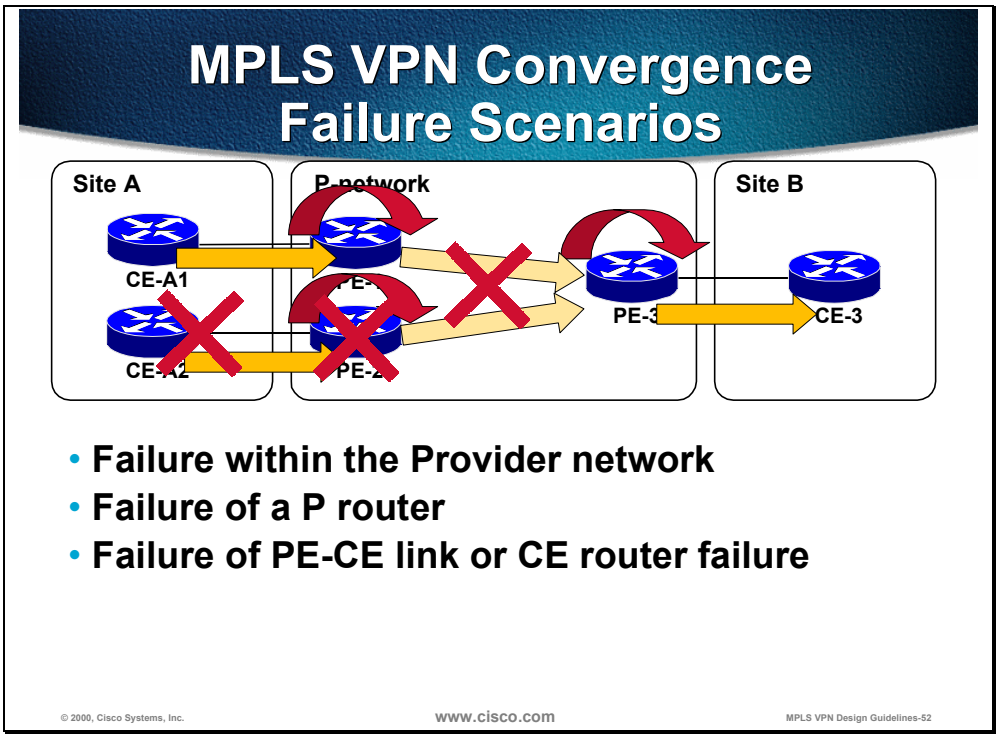

The following three types of convergence failures will be discussed:

- Failure somewhere in the Service Provider network,
- Failure of a PE router and
- Failure of a CE-PE link or the CE router.

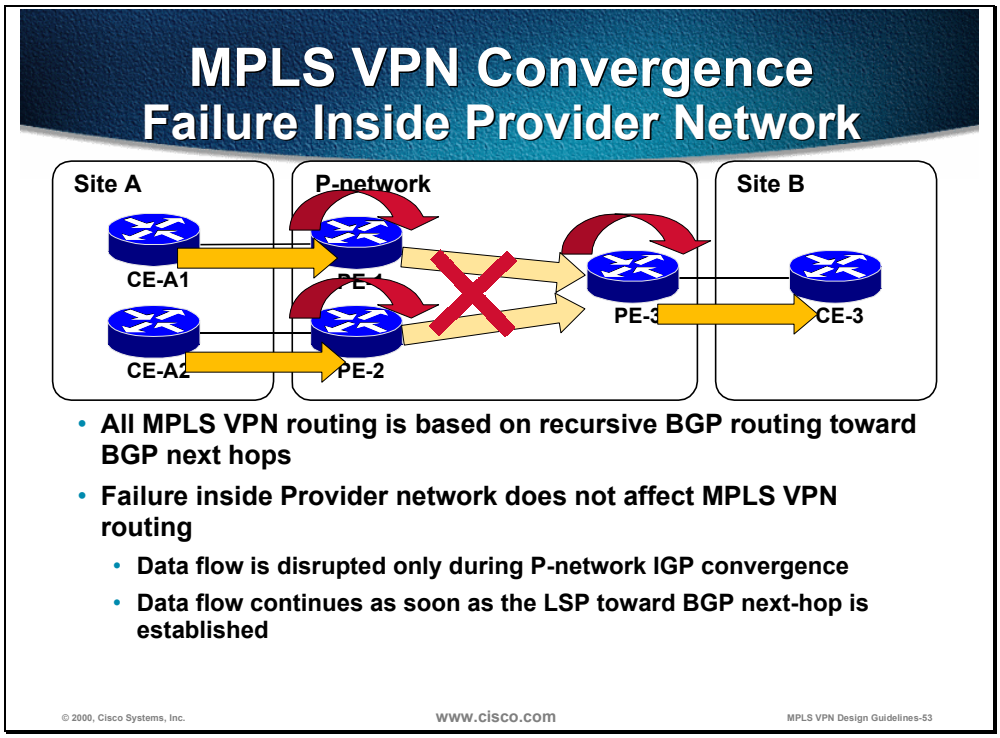

In the first failure scenario, a link goes down somewhere in the SP network. With proper design and implementation of routing protocols, BGP session between the PE-routers should not be lost and the convergence would only be influenced by the time it takes the core (P) routers to discover the failure, the time for IGP to converge and the time for LDP to converge (in certain cases LDP may already have a backup LSP ready).

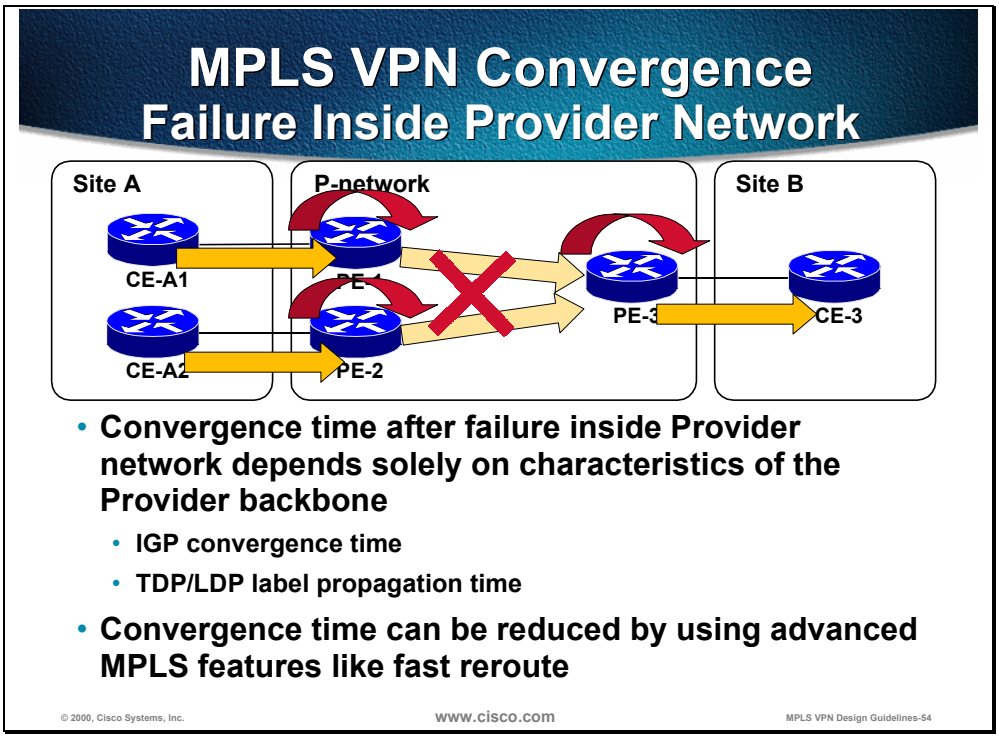

The convergence inside the Provider network is influenced by three factors:

- Time to discover a failure (depends on Layer 1, 2 or IGP),
- Time for IGP to converge, and
- Time for TDP or LDP to propagate the new label once a new path has been installed into the routing table. In some cases a second LSP may already be present due to liberal retention mode when using frame-mode MPLS.

Advanced MPLS traffic engineering features (like *fast reroute*) can be used in the Provider network to reduce the convergence time from a few seconds (which is the best time achievable with a fast and well-tuned IGP) to fewer than 50 milliseconds.

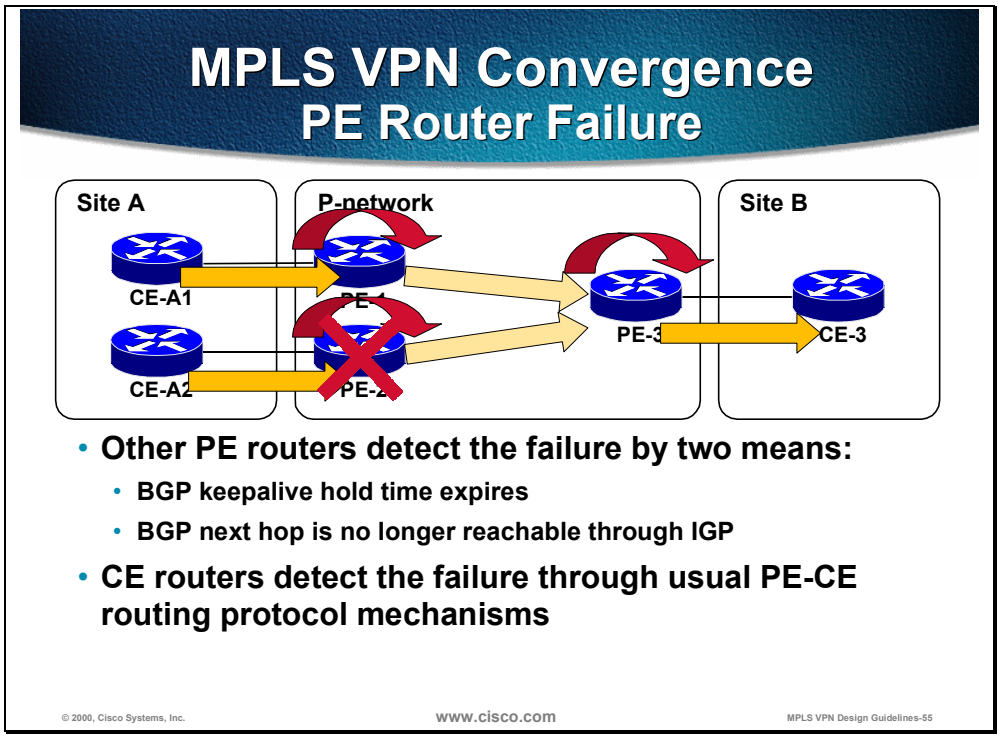

In the second scenario, a PE router fails and its BGP neighbors will require some time to realize their neighbor is lost (failure to receive three BGP keepalives will result in loss of BGP session and discarding of all networks announced by the failing router). A faster way of realizing a neighboring BGP router is down is by losing the path toward its loopback address. This can be achieved by not inserting any summary or a static route to null interface for the loopback address. The downside of this solution is that a prolonged failure somewhere in the P network might also cause the BGP neighborship to be lost even if a backup path exists.

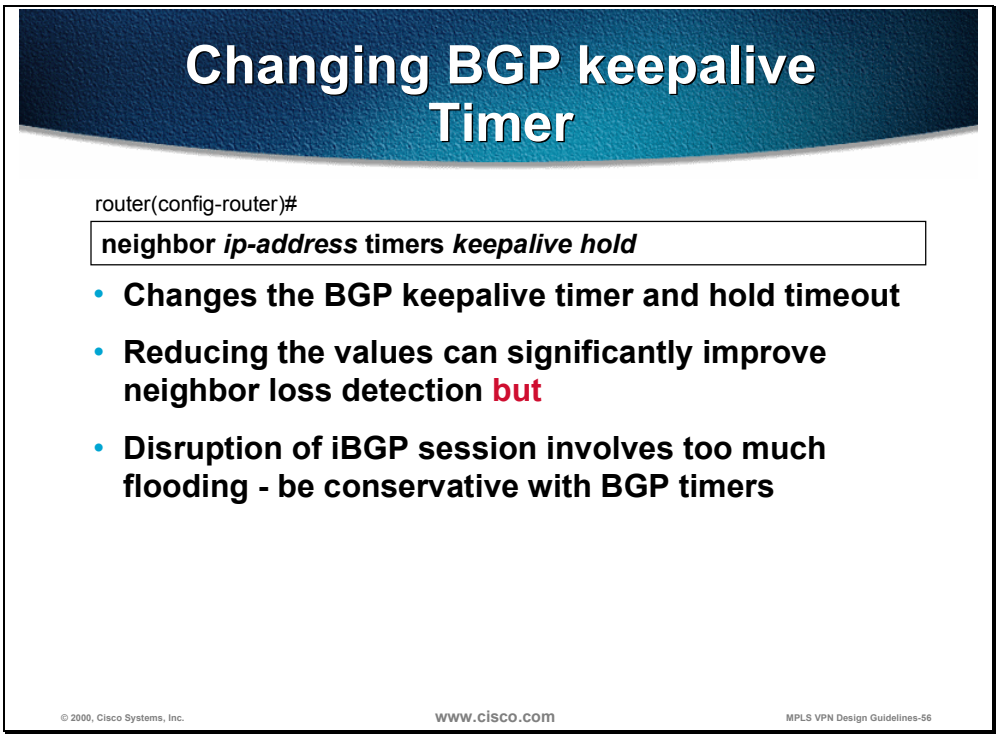

BGP implementation in Cisco routers allows tuning of many timers. Default values of these timers were designed for large amounts of routing information, which is typically the case in Service Provider networks. Changing (lowering) these timers improves performance but it also reduces stability of the network and requires more CPU power.

Changing neighbor keepalive timers normally does not have any drawbacks as long as the hold-time is three times the keepalive time. The recommended minimum values are one second for keepalive time and three seconds for hold time. An over-used CPU may, however, still lose a neighborship for failing to send or receive keepalives. It is, therefore, better to use more conservative values to prevent random BGP flaps.

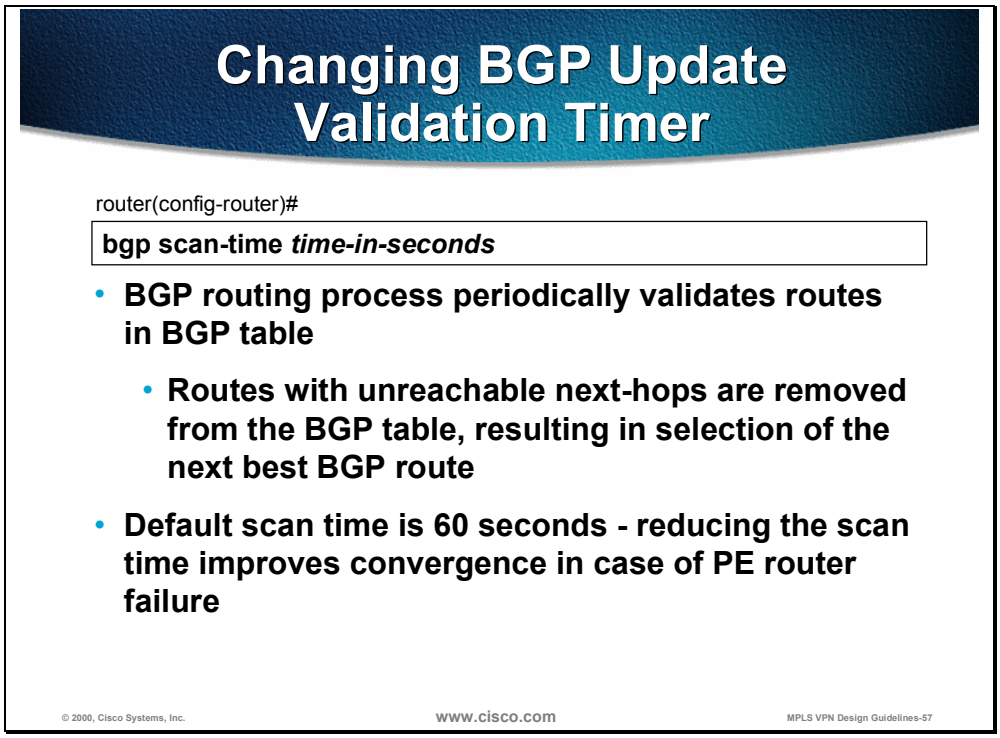

The BGP process periodically scans the routing table for changes. If a locally originated network is lost, BGP has to withdraw it from the BGP table. If the next-hop address is lost, BGP has to withdraw all networks using that next-hop address.

Reducing the timer setting may be acceptable when there is not much information in the BGP table. However it takes a long time for a router to scan the entire BGP table if it contains several hundred thousand networks.

Choosing the right timer setting is, therefore, influenced by the amount of routing information in the routing table and the strength of the CPU in the router.

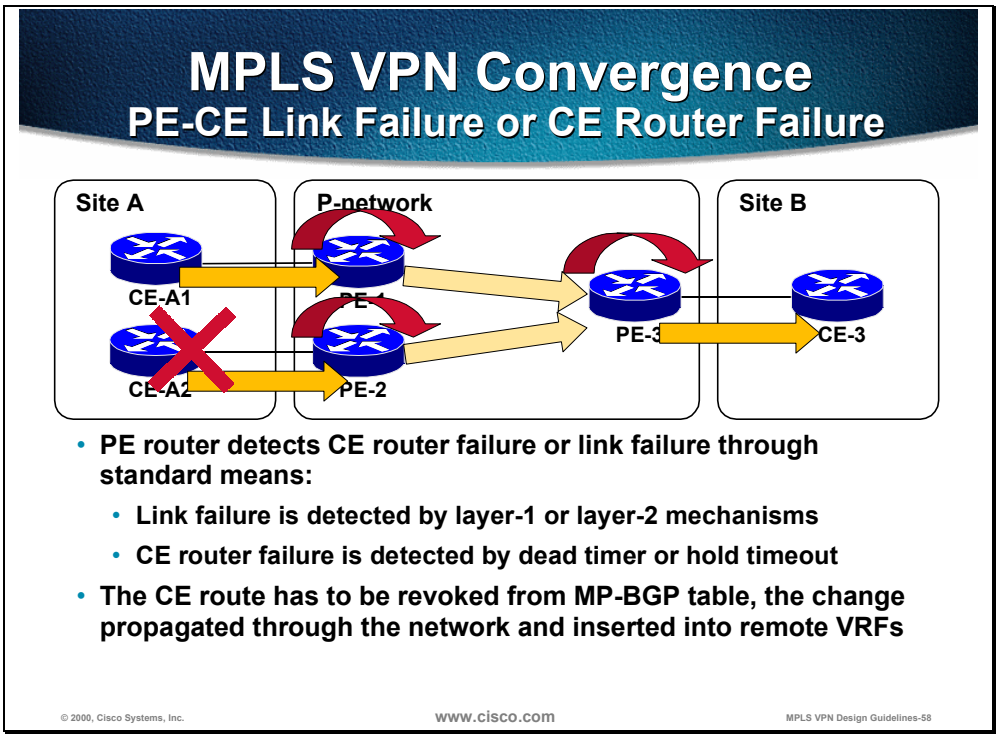

In the third scenario, a CE router fails or a CE-PE link goes down. Again the convergence depends on the time it takes to discover the failure, propagate the changes and recalculate the shortest path.

Typically, layer-1 or layer-2 signaling reports a failure. If not, then the IGP used between CE and PE routers should eventually recognize that a neighbor is no longer responding. After the IGP process removes the VPN route, the change has to be recognized by the BGP process (BGP scan timer), sent to other neighbors (BGP advertisement timer), imported (BGP import/export timer) and sent to other CE routers.

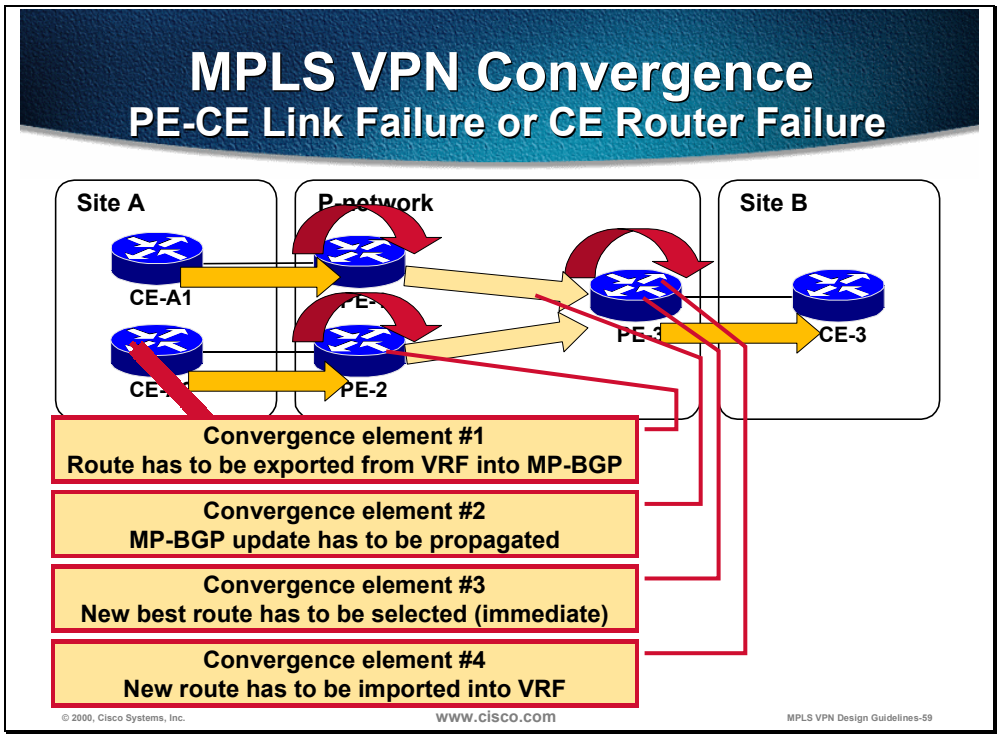

The figure above shows the necessary steps to allow information to be propagated across an MPLS/VPN network:

Step 1 Route (or lack of a route) has to be exported from a VRF into MP-BGP

Step 2 MP-BGP update has to be propagated

Step 3 New best route has to be selected on other PE routers

Step 4 New route has to be imported into the VRF and propagated to other CE-routers

All steps except #3 are periodic and their corresponding timer settings can be changed. See the next two pages for commands that change the BGP import/export timer and the advertisement timer.

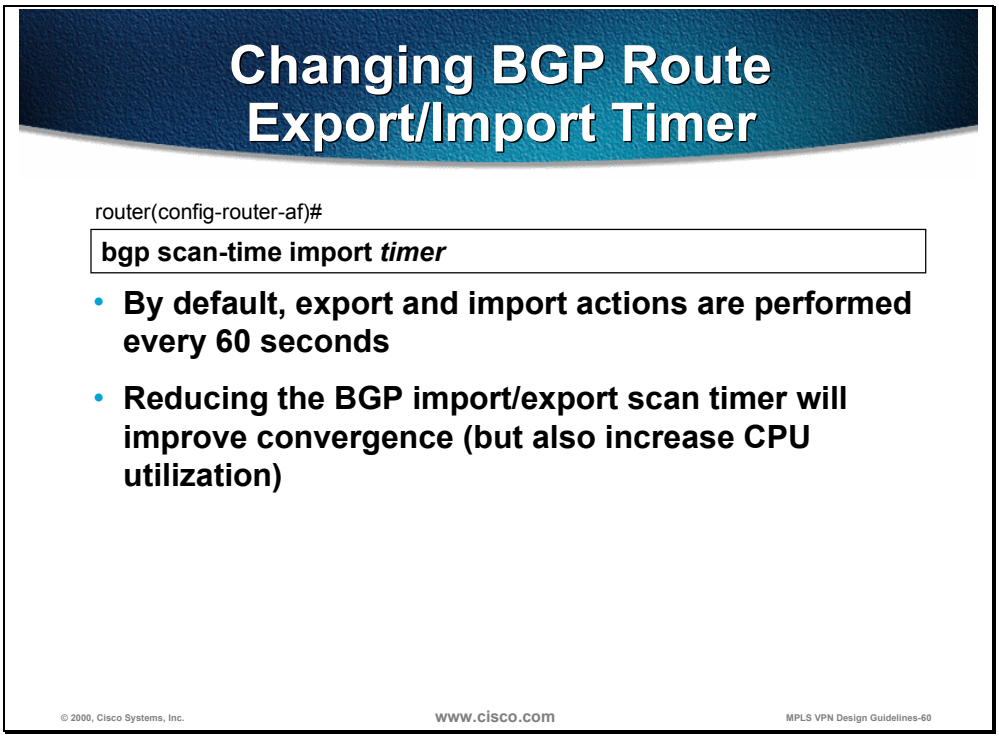

This command should be used with great care. It is very likely that in a large MPLS/VPN network there will be more routing information than is currently contained in the full Internet routing table. Lowering the timer setting may cause the router to be busy performing an import/export scan when it should already be starting a new scan.

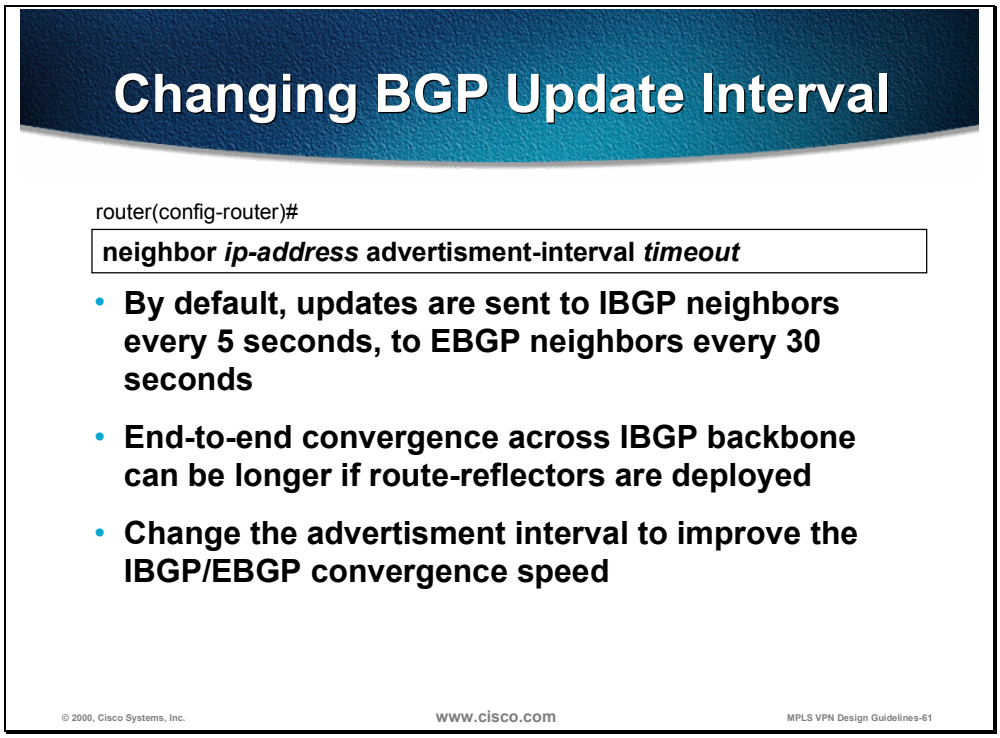

When a router receives a BGP update it will immediately forward the update and start the timer for the neighbor. If there is another update received, the router will first wait for the timer to expire (5 seconds for internal neighbors, 30 seconds for external neighbors) before forwarding the second update.

This command will change the default value. Setting the timer setting to zero may cause a flapping link to be seen on the other side of the MPLS/VPN with the same intensity – all routers have to recalculate the best path whenever a change occurs.

Summary

This section described the differences in convergence when using Overlay VPNs or MPLS/VPNs. The responsibility for fine-tuning this convergence falls mainly on the Service Provider.

To improve convergence in an MPLS/VPN network, the following factors to determine whether there are any opportunities for fine tuning:

- Time to realize a failure
- Time to propagate a change in IGP
- Time to redistribute between protocols
- Time to propagate a change in BGP
- Time to import/export between MP-BGP and a VRF
- Time to recalculate a new path in IGP or BGP

Review Questions

Answer the following questions:

- What are the major elements of end-to-end convergence in traditional overlay VPN networks?
- Which part of the end-to-end MPLS/VPN solution performs the most complex routing?
- What are the three common failure scenarios in MPLS/VPN solution?
- How is the MPLS/VPN routing influenced by a failure in a provider network?
- What influences the overall convergence after a failure in a provider network?
- How can a PE router detect the failure of another PE router?
- How can a CE router detect the failure of an adjacent PE router?
- Which parameters influence the MPLS/VPN convergence after PE router failure?
- How can a PE router detect the PE-CE link failure?
- Which convergence steps need to be taken after PE-CE link failure?
- Which parameters influence the MPLS/VPN convergence after PE-CE link failure?

Chapter Summary

After completing this chapter, you should be able to perform the following tasks:

- Select a proper addressing scheme for the MPLS/VPN backbone.
- Select the optimal Interior Gateway Protocol.
- Develop comprehensive Route Distinguisher and Route Target Allocation Schemes.
- Design BGP in the MP-BGP backbone.
- Optimize overall network convergence.

Answers to Review Questions

Backbone and PE-CE Link Addressing Scheme

- What are the drawbacks of using unnumbered links?

  Individual WAN interfaces are no longer reachable by **ping** or **telnet** if you use unnumbered links.

- Where should you use unnumbered links in the MPLS backbone?

  Unnumbered links are recommended in the ATM parts of the MPLS backbone.

- Where would you use unnumbered links between PE and CE routers?

  Using unnumbered links between PE and CE routers is highly discouraged. There are, however, applications like dial-up access that benefit from unnumbered links.

- Why would you use private address space in your IP backbone?

  IP backbones usually only use private address space if there is no public address space available.

- What are the drawbacks of using private address space in your IP backbone?

  Traceroute across a public IP backbone using private address space usually does not work.

- How would you hide the private address space from your customers?

  If you disable MPLS TTL propagation, the customers cannot see the P-routers. Using private address space between P-routers is then safe.

- What is the impact of using private backbone addresses on traceroute?

  ICMP replies received from private IP addresses would most likely be dropped by customer firewalls. IP address lookup through DNS would also fail.

- Why should you allocate PE loopback addresses from a separate address block?

  The PE loopback addresses should be allocated from a separate block to make sure they are not accidentally summarized in the backbone.

- Why should you use registered addresses for PE-CE links?

  Registered addresses should be used on PE-CE links to prevent potential overlap with the address space the customer is using.

- Why is the reuse of registered addresses between VRFs not advisable?

  You should not reuse addresses between VRFs, as a customer connected to a wrong interface might gain connectivity within the VPN of another customer.

- When can you reuse registered addresses in the same VPN between PE routers?

  You can reuse the same address range on several PE routers if you don't redistribute connected routes into MP-BGP.

Backbone IGP Selection and Design

- List three IGP selection criteria.

  Typical IGP selection criteria are convergence speed, stability and summarization support.

- What is the impact of higher convergence speed on network stability?

  Higher convergence speed always reduces network stability.

- How can you tune OSPF convergence?

  OSPF convergence can be fine-tuned by changing neighbor dead timeout and SPF timer.

- How can you tune IS-IS convergence?

  Many IS-IS parameters can be fine-tuned, from neighbor dead timeout to SPF timers, retransmission timers, LSP origination timeouts etc.

- What is the difference between OSPF and IS-IS route redistribution?

  Redistributed routes appear as separate LSA type-5 objects in OSPF, they appear as part of router LSP in IS-IS.

- Where can you summarize redistributed routes in OSPF?

  You cannot summarize redistributed OSPF routes.

- Where can you summarize redistributed routes in IS-IS?

  Routes redistributed into IS-IS can be summarized between level-1 and level-2 IS-IS areas.

- How do you avoid redistribution of connected interfaces when using OSPF?

  You include connected interfaces in the OSPF process and make them passive.

- Which routing protocols support MPLS Traffic Engineering?

  MPLS Traffic Engineering is supported by OSPF and IS-IS.

- Why is MPLS TE not supported by EIGRP?

  EIGRP cannot support MPLS TE because any router establishing MPLS TE tunnels require full knowledge of the backbone, which is only provided through link-state routing protocols.

- When can you use EIGRP as the IGP protocol in your MPLS/VPN backbone?

  You can use EIGRP as long as you don't plan to deploy MPLS Traffic Engineering.

- What is the impact of route summarization on MPLS/VPN?

Route summarization might break MPLS VPN connectivity if you summarize VPNv4 BGP next-hops (loopback addresses of PE routers).

- Why is IS-IS recommended for extremely large networks?

  Many large Service Providers use IS-IS, therefore there is more experience with running IS-IS in large networks.

Route Distinguisher and Route Target Allocation Scheme

- What is the function of the route distinguisher?

  Route distinguisher is used to make overlapping IPv4 addresses globally unique.

- Can you reuse the same route distinguisher on different PE routers?

  You can reuse the same route distinguisher as long as the VRFs on the PE routers have the same connectivity requirement.

- Is there any topology where every site requires a different value of route distinguisher?

  Hub-and-spoke topology requires a different value of route distinguisher for every site.

- What is the function of the route target?

  Route target controls the import of VPNv4 routes into VRFs.

- Do you have to make the route target equal to the route distinguisher?

  Route target can be different from route distinguisher.

End-to-End Convergence Issues

- What are the major elements of end-to-end convergence in traditional overlay VPN networks?

  The major elements are:

  - Neighbor loss detection
  - Routing update propagation
  - Computation of the new topology (SPF run)

- Which part of the end-to-end MPLS/VPN solution performs the most complex routing?

  Service Provider PE-routers perform the most complex routing.

- What are the three common failure scenarios in MPLS/VPN solution?

  The common failure scenarios are:

  - Failure in the P-network
  - Failure of the PE-router
  - Failure of the PE-CE link (most common).

- How is the MPLS/VPN routing influenced by a failure in a provider network?

Failure in a provider network shall not influence MPLS VPN routing, as long as the IGP in the P-network converges fast enough.

- What influences the overall convergence after a failure in a provider network?

  The overall convergence is affected only by the convergence speed of the IGP used in the P-network.

- How can a PE router detect the failure of another PE router?

  A PE-router can detect neighbor loss through BGP hold timer timeout or through loss of BGP next-hop.

- How can a CE router detect the failure of an adjacent PE router?

  CE router uses traditional routing protocol mechanisms (for example, dead timeout in OSPF or invalid timer in RIP).

- Which parameters influence the MPLS/VPN convergence after PE router failure?

  BGP neighbor timers and BGP scan-time affect MPLS VPN convergence after a PE-router failure.

- How can a PE router detect the PE-CE link failure?

  PE router could detect the PE-CE link failure through layer-1 or layer-2 signaling (for example, carrier loss or DLCI failure signaled by LMI). It can also detect PE-CE link failure with traditional routing protocol mechanisms (for example, dead timeout in OSPF or invalid timer in RIP).

- Which convergence steps need to be taken after PE-CE link failure?

  The following steps are taken:

  - VRF route is removed from the VRF routing table
  - VRF route is removed from the VPNv4 BGP table
  - Withdrawal of VPNv4 route is propagated to other PE-routers
  - Other PE-routers select a new best BGP route
  - The newly selected BGP route is imported into the VRFs on other PE-routers.

- Which parameters influence the MPLS/VPN convergence after PE-CE link failure?

  MPLS VPN convergence after PE-CE link failure is affected by BGP update interval and BGP import scan timer.

8

Large-Scale MPLS VPN Deployment

Overview

This chapter describes scalability issues encountered in large-scale MPLS VPN networks and presents a number of solutions that allow these networks to scale while growing.

It includes the following topics:

- MP-BGP Scalability Mechanisms
- Partitioned Route Reflectors

Objectives

Upon completion of this chapter, you will be able to perform the following tasks:

- Understand the MP-BGP scaling issues in large-scale MPLS VPN backbones
- Describe the built-in scalability mechanisms
- Design and implement networks using partitioned BGP route reflectors

MP-BGP Scalability Mechanisms

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Understand MP-BGP scaling issues
- Describe the automatic filtering in MP-BGP
- Describe the functions of the BGP Route Refresh feature
- Describe the Outbound Route Filter feature and its benefits

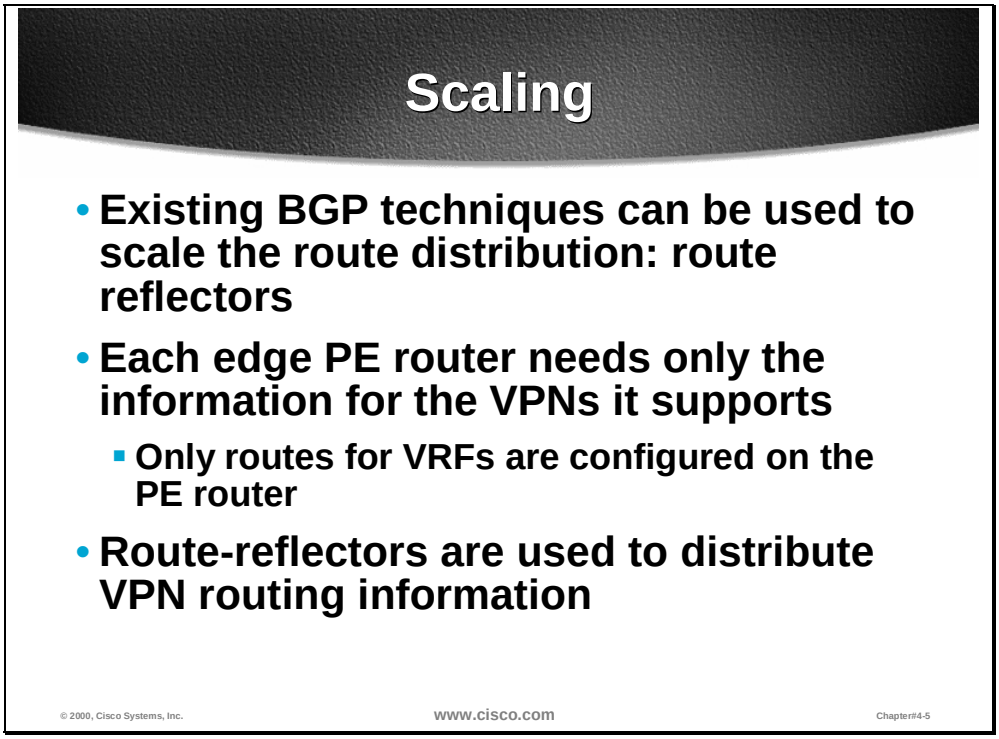

A network designer that wants to design a scalable MPLS VPN solution is always faced with a number of scalability issues, several of them being related to the MPLS VPN architecture:

- MPLS VPN uses internal BGP (IBGP) to propagate VPNv4 routes between PE routers. Default IBGP implementation requires a full-mesh of BGP sessions between PE routers—a design that is only appropriate for very small networks.
- As the number of MPLS VPN customers grows, the PE routers have to store more and more customer routes (in traditional overlay VPN implementations, the customer routes are not seen by the provider routers—this issue is therefore not present in overlay VPN implementations). In very large MPLS VPN networks, providing connectivity to large customers, the number of routes that need to be stored by the PE routers exceeds the current scaling capabilities of Cisco IOS BGP implementation as well as memory and CPU resources of the PE routers.

The IBGP full-mesh scalability roadblock is easily removed using traditional BGP scaling tools—**BGP route reflectors** and **BGP confederations** (both of them are described in the appropriate lessons of the BGP curriculum and their operations will not be discussed further in this section).

Note BGP route reflectors are a preferred scalability tool for MPLS VPN networks and their positioning will be covered extensively in the next section.

The memory and CPU requirements imposed on a PE router by a large number of customer routes can be easily reduced if the PE router only stores routes relevant to the VPN customers connected to it and ignores all the other VPNv4 routes. The incoming route filtering had to be configured manually with early MPLS VPN implementation. To reduce the configuration complexity, Cisco IOS releases 12.0(7) T and 12.1 provide automatic filtering of incoming Multi-protocol BGP (MP-BGP) updates.

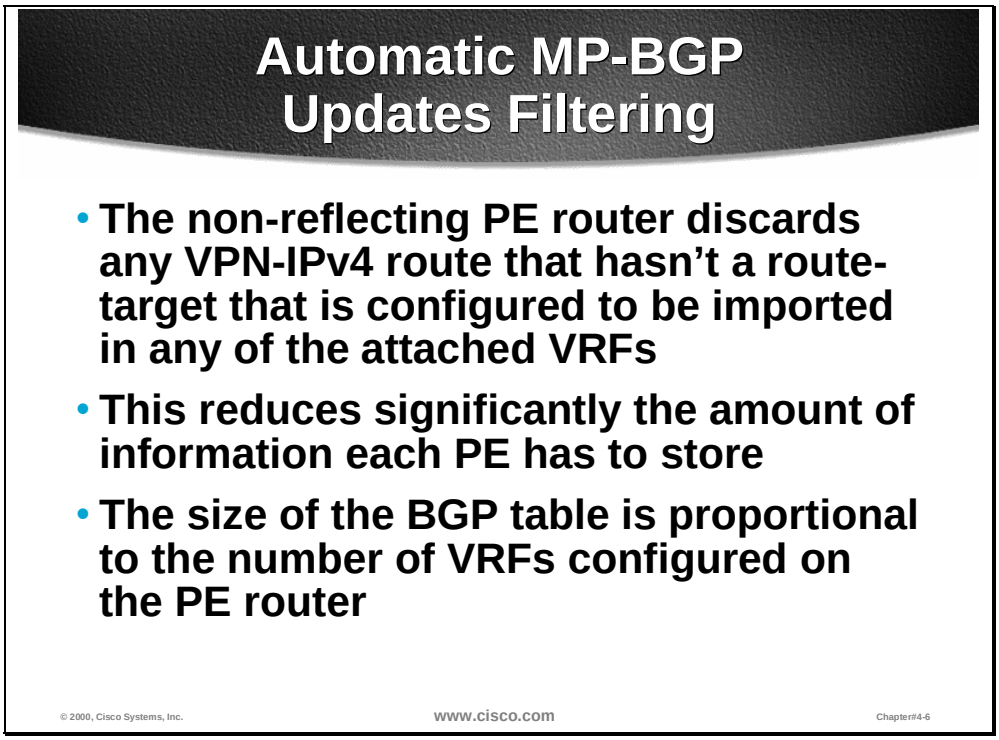

The automatic MP-BGP updates filtering uses a very simple algorithm—all VPNv4 routes received by a PE router that do not correspond to any VRF configured on the router are automatically ignored. This usually results in a significant reduction of VPNv4 BGP table on the PE router, as the size of the table becomes proportional to the number of VRFs configured on the PE router and not the overall size of the MPLS VPN network.

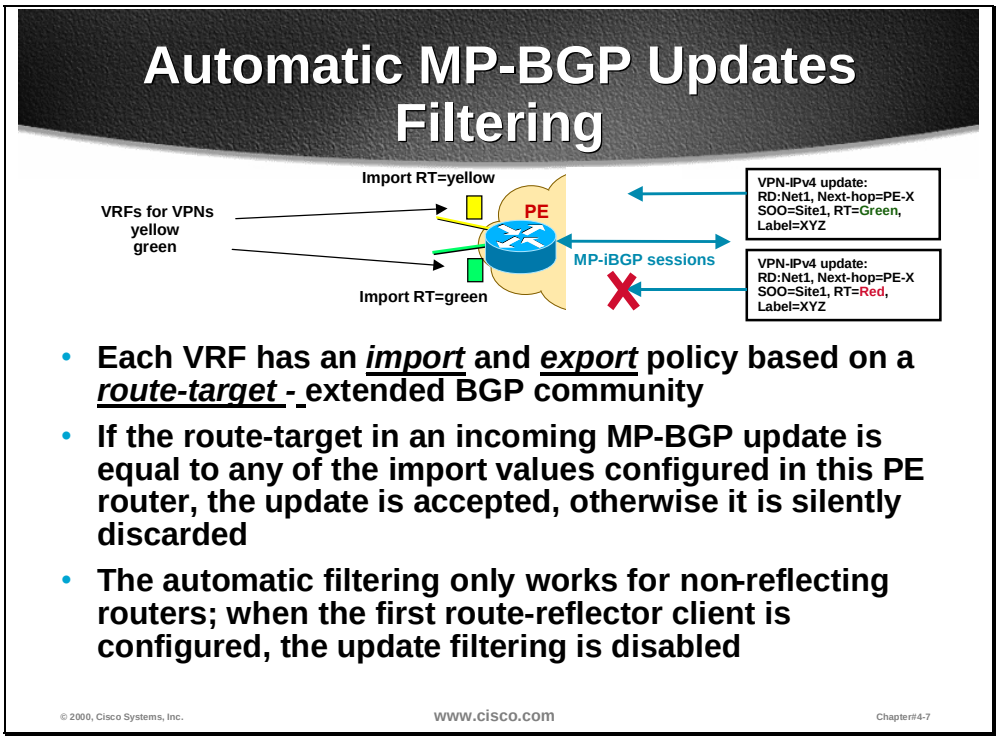

The filtering of incoming VPNv4 update is performed based on **import route-targets** configured in VRFs and the route targets attached to incoming VPNv4 routes. If the incoming VPNv4 route carries a route target that corresponds to an import route target of at least one VRF, the incoming route is potentially useful as it might get inserted into the VRF and is accepted by the PE router. Otherwise the incoming route is silently discarded (similar to other inbound BGP filtering mechanisms).

Note The incoming VPNv4 route that is accepted by automatic inbound filter might still be rejected by **import route-map** configured in the VRF, so the automatic filters are not perfect. Anyhow, taking import route-maps in consideration when filtering incoming VPNv4 updates would significantly increase the CPU load of the PE router.

The automatic inbound filters only work for PE routers that do not act as route reflectors. As there is no mechanism through which a route reflector might discover that one of its clients would need routes with a certain route target, the route reflectors do not filter inbound updates. The route reflectors therefore carry all VPNv4 routes in an MPLS VPN network.

Note A router starts acting as a BGP route-reflector the moment the first route-reflector is configured client with **neighbor route-reflector-client** configuration command. As soon as the first route-reflector client is configured, the automatic inbound filtering of VPNv4 routes is disabled.

The figure above shows an example of inbound filters. The PE router has two VRFs configured, one accepting routes tagged with route-target **green**, the other one accepting routes tagged with route-target **yellow**. When an incoming BGP

update carries a VPNv4 route with RT=**green**, the route is accepted. A VPNv4 route that only carries route target **red** is rejected, as **red** is not configured as an import route target of any VRF on this router.

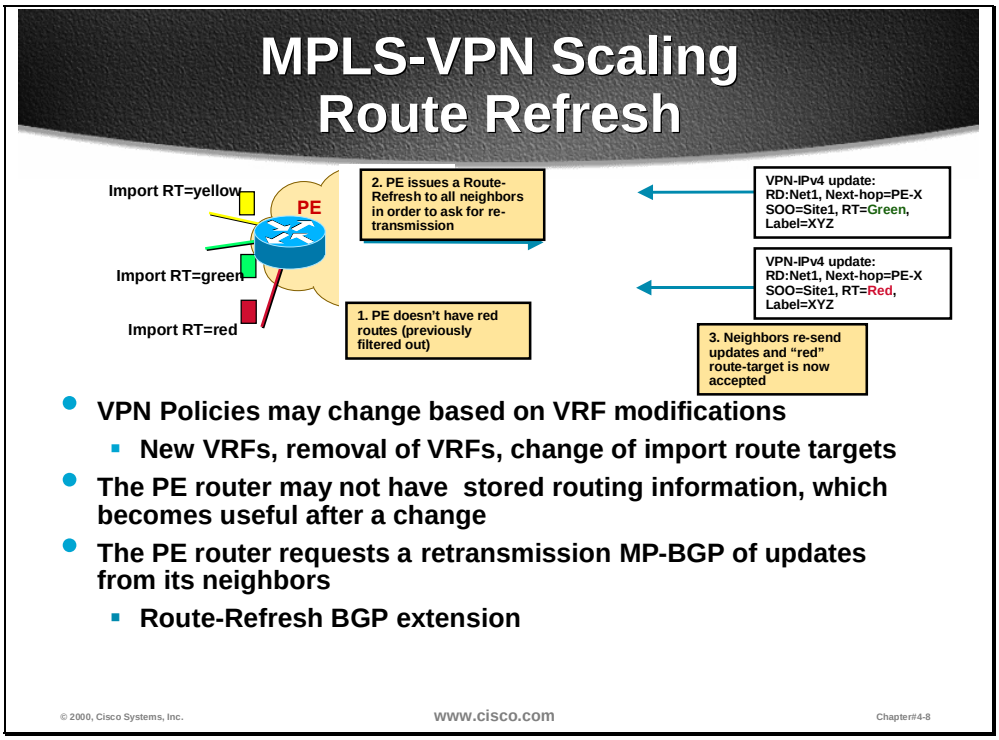

Automatic inbound route filters behave in exactly the same way as manually configured BGP inbound filters. Whenever the routing policy is changed (and the inbound filter is changed), the router might need routes that it has discarded previously. However, there is no mechanism that the router might use to request those routes from its BGP neighbors and the neighbors will never send those routes by themselves, as BGP has no periodic update mechanism.

Classical BGP implementation on Cisco IOS offers two ways to get the routes needed by a BGP router after a change in routing policy:

- The BGP session between the routers might be manually torn down and the neighbor will send all the routes after the session is reestablished.
- The BGP router might store an extra copy of routes sent by the neighbors.

Neither of these options is a viable option for large-scale MPLS VPN deployment because:

- Disruption of a BGP session results in a disruption of MPLS VPN service which is not acceptable for mission-critical customer traffic.
- Storing extra copies of BGP routes would defy the whole purpose of automatic inbound filters.

An extension to BGP, called **BGP route refresh**, was therefore introduced to BGP and subsequently implemented in Cisco IOS to allow a BGP router to **request** a resend of all BGP routes from its neighbor.

Note To optimize the amount of the BGP traffic exchanged between the PE routers, the **route-refresh** message specifies the address family where the refresh is needed. A PE router can thus request only a refresh of VPNv4 routes.

The figure above illustrates the **BGP route refresh** functionality:

- A PE router receives a VPNv4 route that does not contain any route target configured as VRF import route-target on this router. The update is ignored.
- A new VRF is configured on the PE router and the update that was previously ignored is now needed to gain connectivity for this new VRF. The PE router therefore sends a **route-refresh** message to its neighbors, requesting a resend of all their VPNv4 BGP routes.
- Routing update containing all VPNv4 routes is sent by the neighbor receiving **route-refresh** message. This update includes the routes that were previously discarded by inbound route filters.
- The modified inbound route filter accepts the VPNv4 route with **red** route-target and the new VRF is populated.

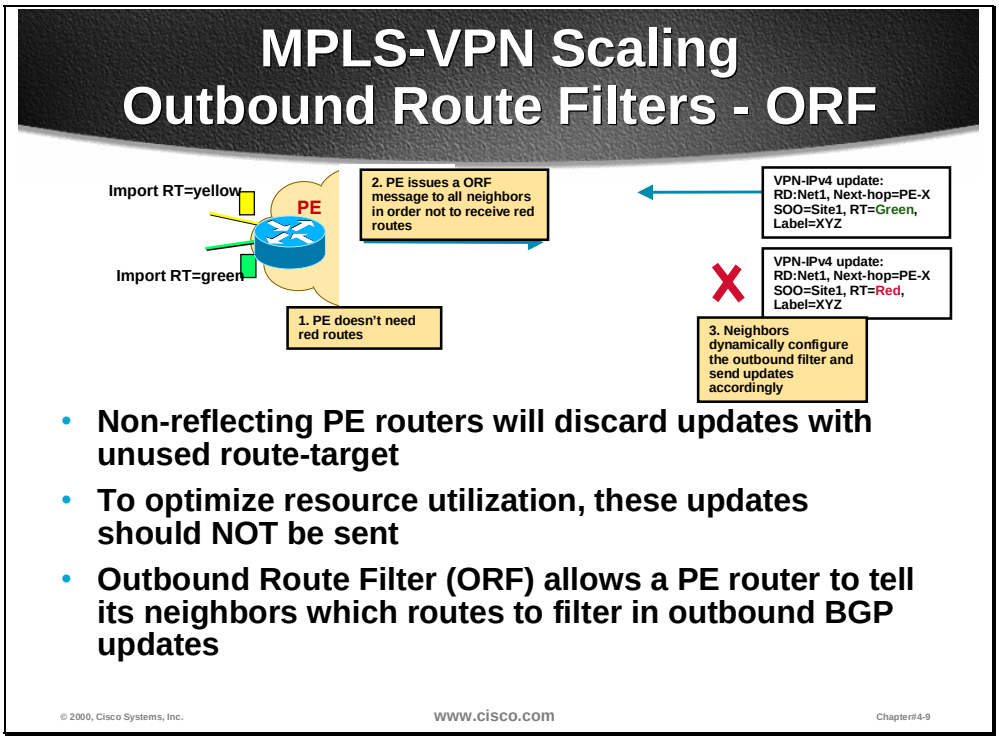

Automatic inbound filters on the PE routers are clearly suboptimal:

- The sending router spends its resources generating the BGP update.
- Network bandwidth is used to propagate the update.
- Receiving router spends its resources filtering the incoming update, only to discard the unnecessary route at the end.

The only way to reduce the overall resource usage would be to filter the BGP update at the sending router as it's being generated. The sending router, however, has no information on the inbound filter of the receiving router.

The **outbound route filter (ORF)** functionality introduced in BGP gives the receiving BGP router a way of downloading its inbound filter as an outbound filter of the sending router. Using ORF functionality, the receiving PE router can make sure that the sending PE router will discard all the routes that would be discarded by the receiving router, prior to sending the information to the receiving router.

Note The filtering capabilities of outbound route filters are severely limited when compared to the richness of BGP filters. The only two BGP filtering mechanisms currently supported by ORF are filters based on **prefix-lists** and automatic inbound filters based on MPLS VPN route targets.

The figure above illustrates the ORF functionality.

- The receiving PE router generates its automatic inbound filter permitting only VPNv4 routes with route-target **yellow** or **green** and downloads that filter as outbound filter to the sending PE router.

- The sending PE router will use this filter and discard the route carrying route-target **red** before it's sent to the receiving router.

Summary

Large-scale MPLS VPN deployments are usually faced with a number of scalability issues:

- The number of PE routers in the network is large and the corresponding MP-IBGP full-mesh does not scale.
- The amount of VPNv4 routing information in the network exceeds the scaling capabilities of BGP routers.

Scalable MP-IBGP design can be implemented using standard BGP scalability tools—BGP route-reflectors or BGP confederations.

The amount of VPNv4 routing information held by a PE router is reduced with automatic inbound filters. These filters discard all routes that are not relevant to the PE router (the routes that do not contain any route-targets configured as import route-targets on the PE router).

Configuration changes on the PE router might change the automatic inbound filter. As BGP routers don't send periodic routing information refreshments, a mechanism is needed to request missing information from other BGP routers – the **bgp route-refresh** functionality.

Outbound route filters are an additional optimization of automatic inbound filters. Through this function, a BGP router can download its inbound filter as an outbound filter of its neighbor, reducing its CPU utilization and the amount of BGP traffic in the network.

Review Questions

- Describe BGP scaling issues in a MPLS VPN network.
- Describe built-in MP-BGP scalability mechanisms.
- Why does the automatic filtering of inbound VPNv4 updates increase MPLS VPN scalability?
- What are the implications of automatic inbound filtering on BGP route-reflector design?
- Why do you need route-refresh functionality?
- When would a router send a route-refresh request to its neighbors?
- What is an outbound route filter (ORF)?
- Why are outbound route filters useful?

Partitioned Route Reflectors

Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe the partitioned route reflector design
- Design MPLS VPN networks using the partitioned route reflector design
- Implement partitioned route reflectors in a MPLS VPN network

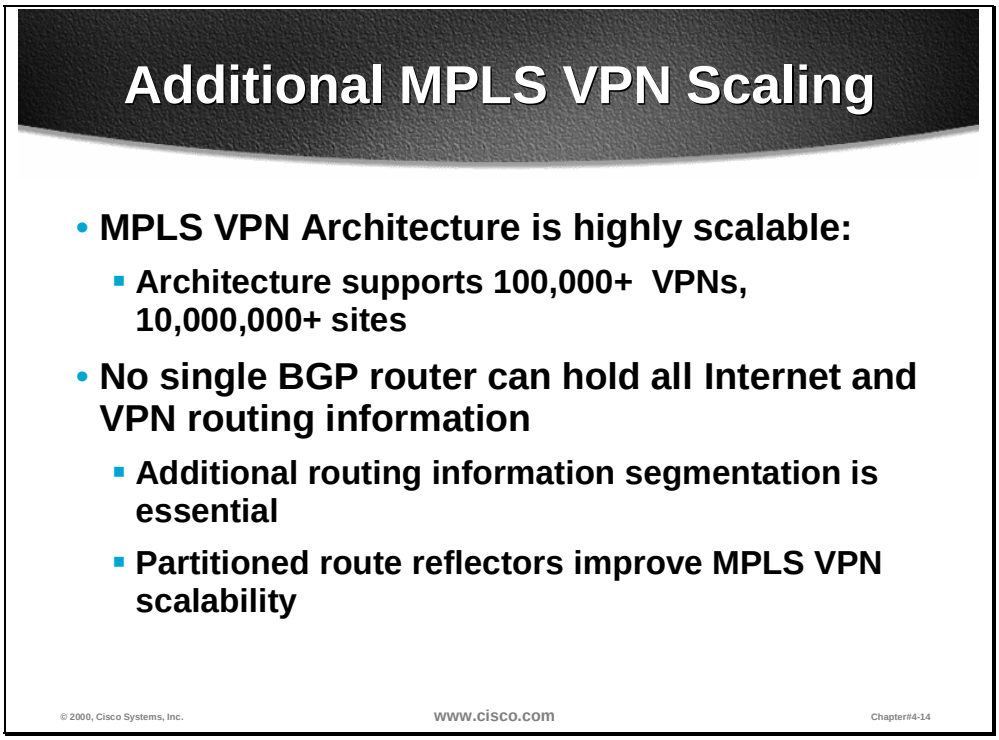

The MPLS VPN architecture is highly scalable and is being used in very large-scale networks supporting thousands of customers and potentially carrying millions of VPNv4 routes. MPLS VPN deployments with such a large overall number of VPNv4 routes defy any BGP implementation. The automatic inbound route filtering functionality provided by Cisco IOS is therefore no longer sufficient as the route-reflectors cannot store all VPNv4 routes any more.

Additional segmentation of routing information is necessary to allow MPLS VPN deployments in very large networks. The network design implementing the segmentation of VPNv4 information is called **partitioned route reflector design**. As the VPNv4 routing information is partitioned between a number of independent route reflectors, each of them stores only a portion of overall VPNv4 routing information.

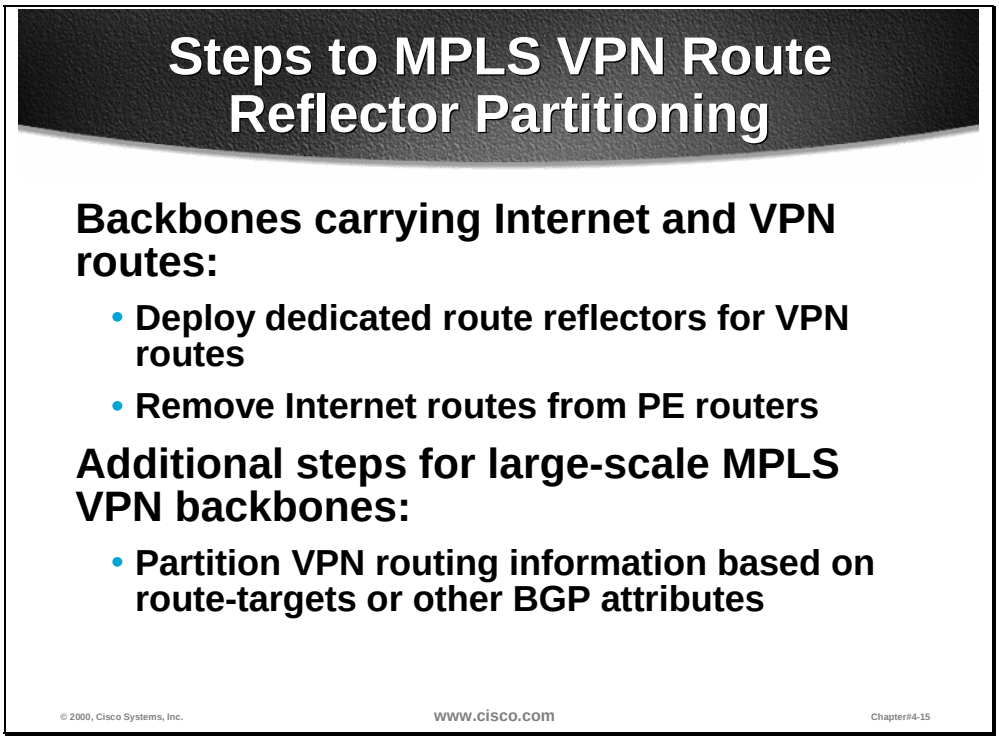

There are a number of intermediate steps that can improve MPLS VPN scalability even before the partitioned route-reflectors are introduced:

- Route-reflectors dedicated to reflecting VPNv4 routes can be introduced to reduce the number of routes carried by a route-reflector.
- Internet routes can be removed from the PE routers, resulting in further reduction of BGP table on the PE routers.

Partitioned route reflectors shall be deployed only when all these measures fail to address the needs of current or planned amount of VPNv4 routing information.

Partitioning of VPNv4 routing information is usually done based on route-targets, however, any BGP attribute (most often standard BGP communities) can be used for this purpose.

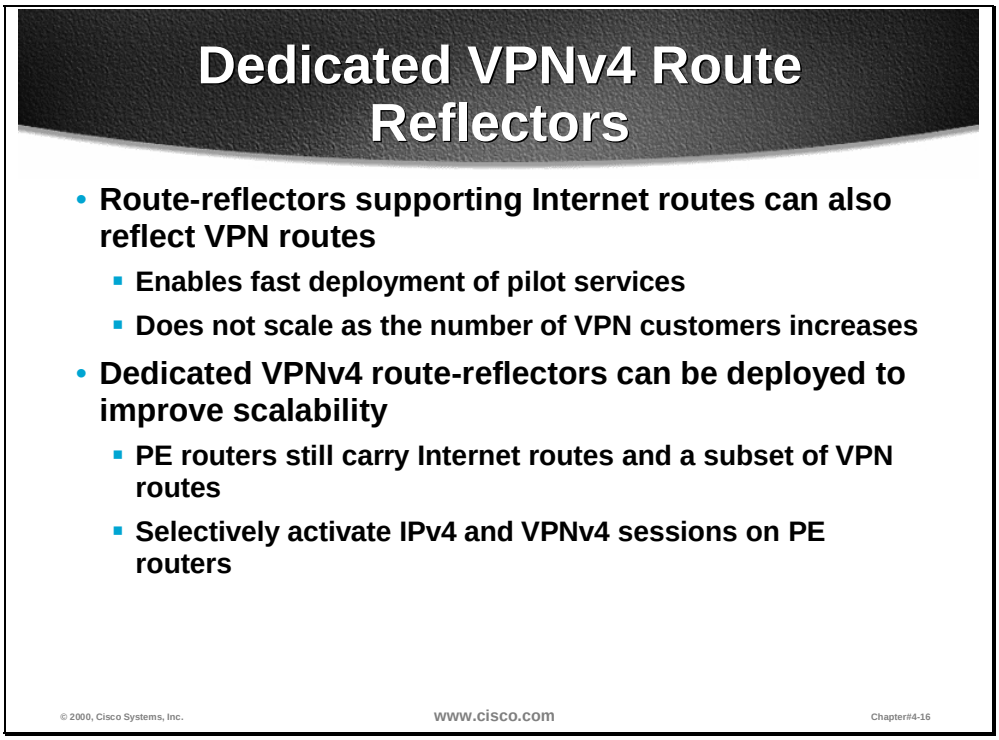

MPLS VPN pilots, as well as small-scale deployments, usually use existing BGP infrastructure to support the needs of MPLS VPN architecture. Existing routers are used as PE routers and existing BGP route reflectors are used to reflect the VPNv4 routes. This approach, while allowing fast deployment of new services, does not scale as the number of VPN customers start to increase.

The first step to increase the scalability of MPLS VPN network is deployment of dedicated VPNv4 route reflectors. This step reduces the load of IPv4 route reflectors, but does not reduce the load of the PE routers that still have to carry Internet routes and VPNv4 routes relevant for the VRFs configured on them.

The separation of IPv4 and VPNv4 routing information between two dedicated sets of route-reflectors is performed by selective activation of IPv4 and VPNv4 sessions on the PE routers and route reflectors.

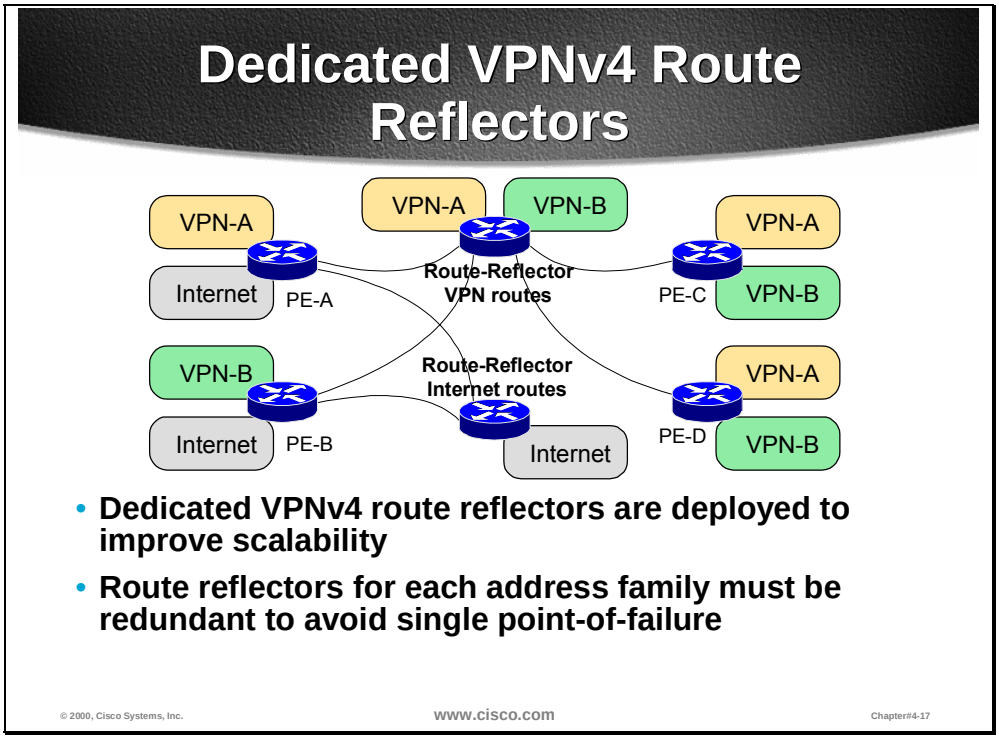

The diagram above displays an MPLS VPN network where the routing information has been split between a route-reflector carrying only VPNv4 routes and another route-reflector carrying IPv4 routes. The PE routers still carry full Internet routing (or partial Internet routing, based on the BGP design) as well as the VPNv4 routes relevant to them.

Note The example above **shall not** be used as a template for MPLS VPN network deployment. Route reflectors for each address family (VPNv4 and IPv4) shall be redundant to avoid single point-of-failure.

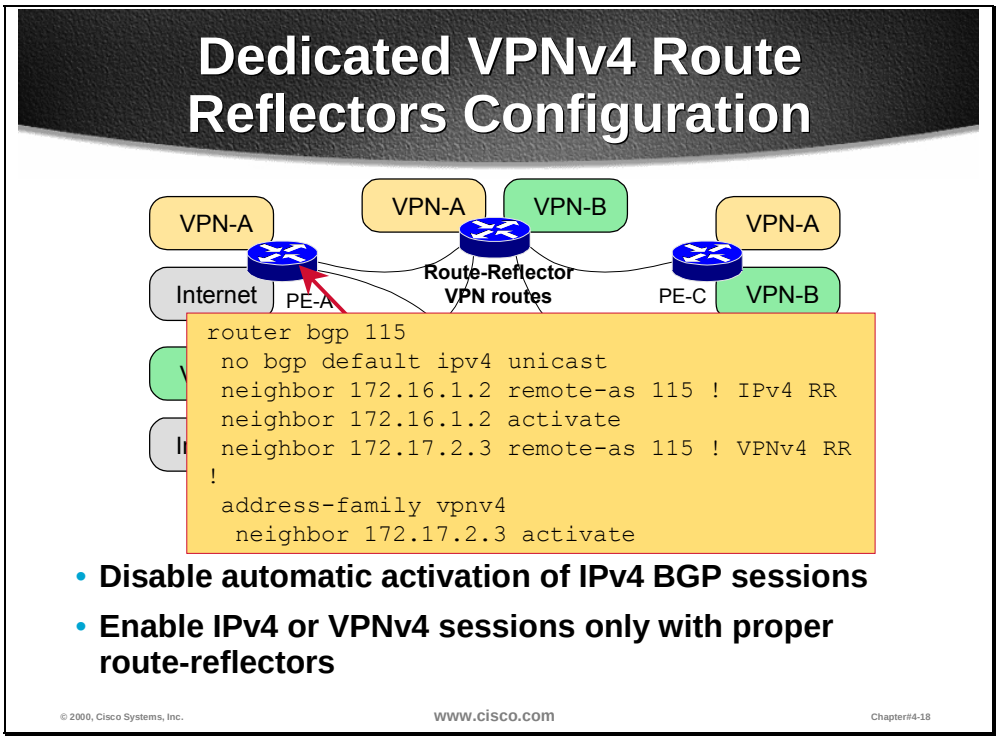

The example above displays the PE router configuration used to achieve separation of VPNv4 and IPv4 routes between two sets of route-reflectors. The automatic activation of IPv4 BGP sessions is disabled to make sure that the IPv4 routes are not sent to the route-reflectors carrying only VPNv4 routing information.

The route-reflectors are configured as BGP neighbors of the PE router. The IPv4 session is only activated toward the route-reflector carrying IPv4 routes (the BGP neighbor with the IP address 172.16.1.2) and the VPNv4 session is only activated toward the route-reflector carrying VPNv4 routes (the BGP neighbor with the IP address 172.17.2.3).

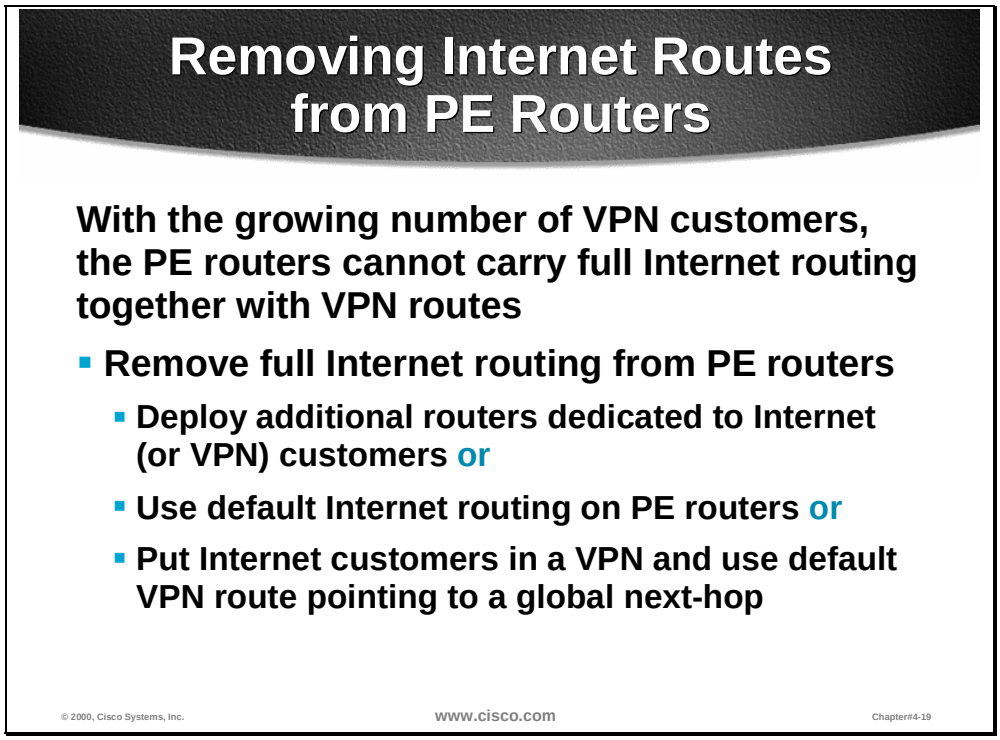

With the growing number of VPN customers, it will come to a point where the PE routers cannot carry full Internet routing together with the VPNv4 routes, even after the automatic inbound filters have reduced the number of VPNv4 routes carried by the PE router. When this point is reached, the next scalability measure is the removal of full Internet routing from the PE routers. This action might break the Internet routing and has to be preceded by a thorough network redesign and migration planning.

There are three ways to address this scalability step:

- Deploy additional routers, establishing routers that are dedicated to providing Internet services and another set of routers dedicated to providing MPLS VPN services. This approach requires a large number of changes in the transition period (including reconnecting a large number of customers to another router) and is therefore usually avoided as a migration step. There are, however, large Service Providers that have initially deployed MPLS VPN as a separate service and have always provided dedicated PE routers to address the scalability needs of MPLS VPN services.
- Use partial Internet routing in combination with the default route on the PE routers. This approach can only be applied to PE routers that are not in a transit path and can still get optimal routing (or close-to-optimal routing) when using a default route.

Note Please refer to the technical solutions in the BGP curriculum for further discussion on default route usage in networks supporting Internet services.

- Migrate your Internet customers into a VPN, using mechanisms explained in the **Internet Access from a VPN** chapter of this lesson.

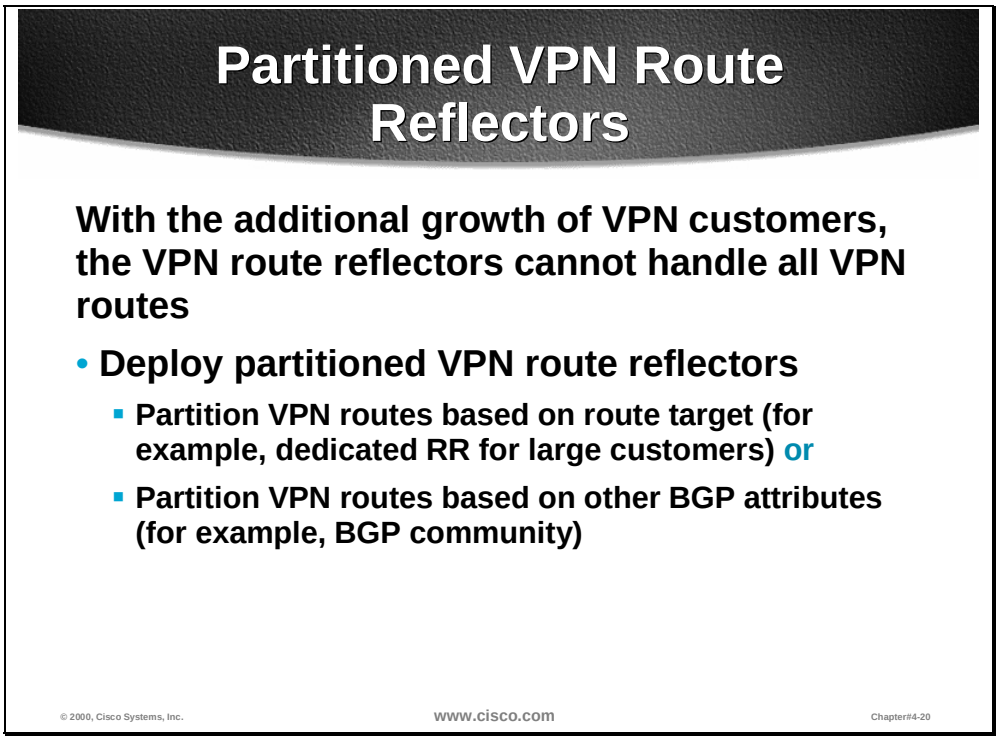

With additional growth of a MPLS VPN network, the route-reflectors carrying VPNv4 routes will not be able to handle the amount of VPNv4 routes that need to be propagated in the network. At this moment, the VPNv4 routing information has to be partitioned and additional route-reflectors deployed, where each set of route-reflectors will only carry a portion of overall VPNv4 routing information.

Partitioning of VPNv4 routes is usually done based on route-target, for example, a dedicated set of route-reflectors for a single very large MPLS VPN customer. However, it could be done on any other BGP attribute, for example, based on standard BGP community.

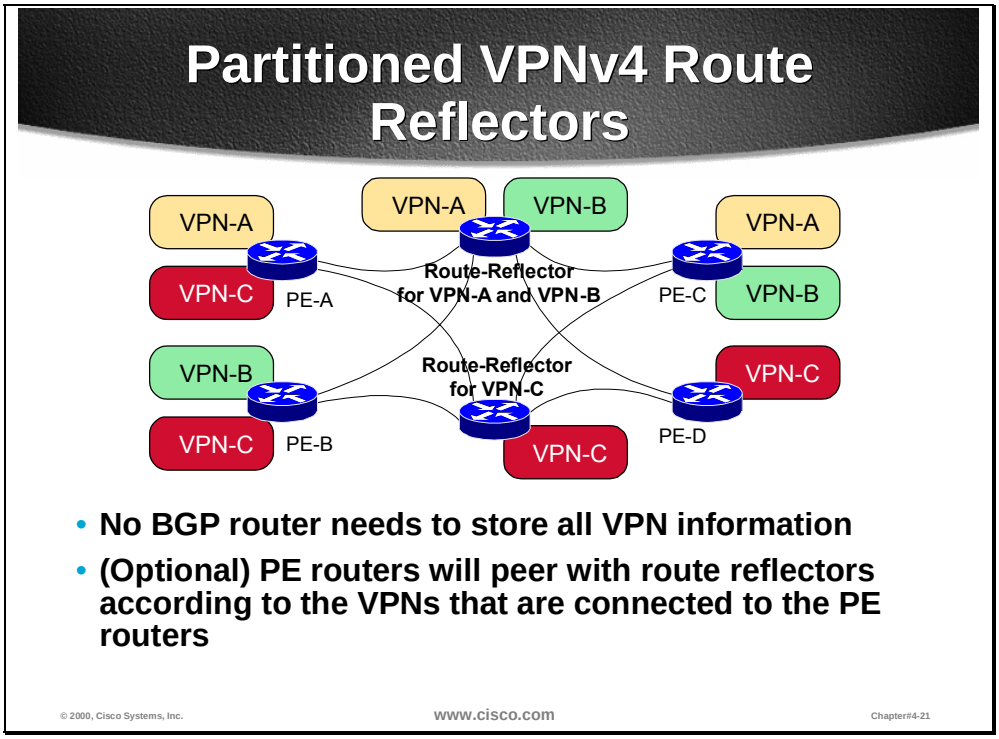

The diagram above demonstrates a partitioned VPNv4 route-reflector setup. The top route-reflector only accepts routes with route-target **green** and **yellow** and the bottom route-reflector only accepts routes with route-target **red**. In order to receive all the routing information required for proper operation, all PE routers need to have BGP sessions to all route reflectors.

Further reduction of resource utilization in the network can be made if the PE routers only peer with the route-reflectors that carry routing information relevant to the PE routers. This setup, although more optimal than the one presented above, introduces management and configuration complexity and is best avoided.

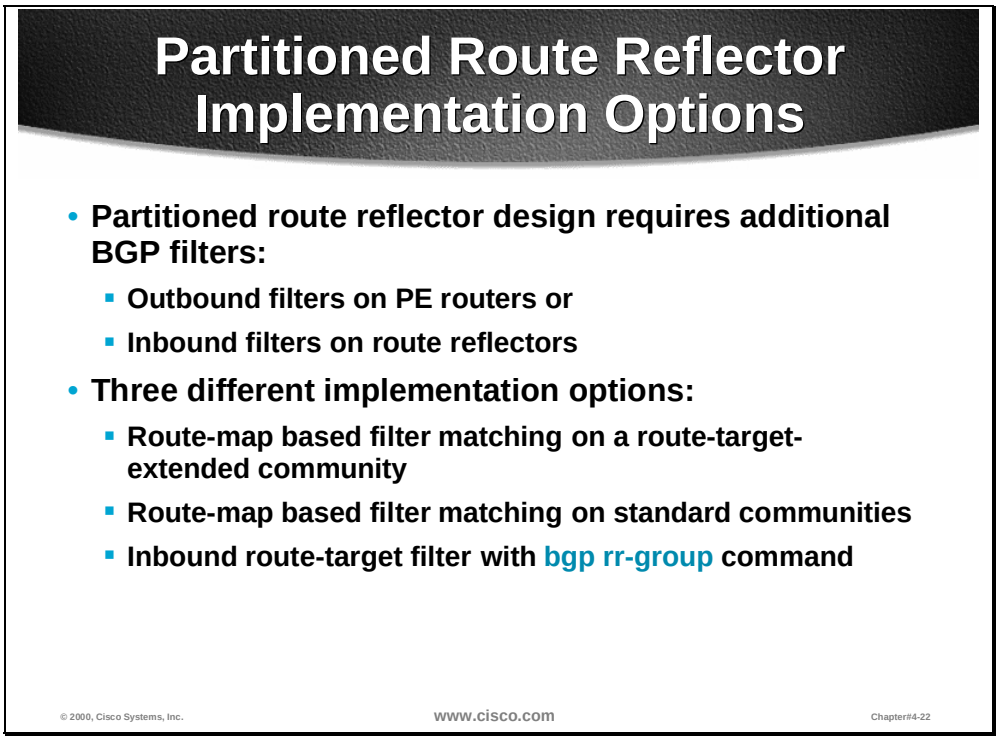

The partitioned route-reflectors can be achieved by configuring outbound filters on the PE routers or inbound filters on route reflectors. In both cases, the filtering can be performed with a route-map matching routes on any BGP attribute—usually on route-target or standard BGP community.

An additional filtering mechanism, configured with the **bgp rr-group** command, (an explanation follows) can be used to configure inbound route-target filter on the BGP route-reflector.

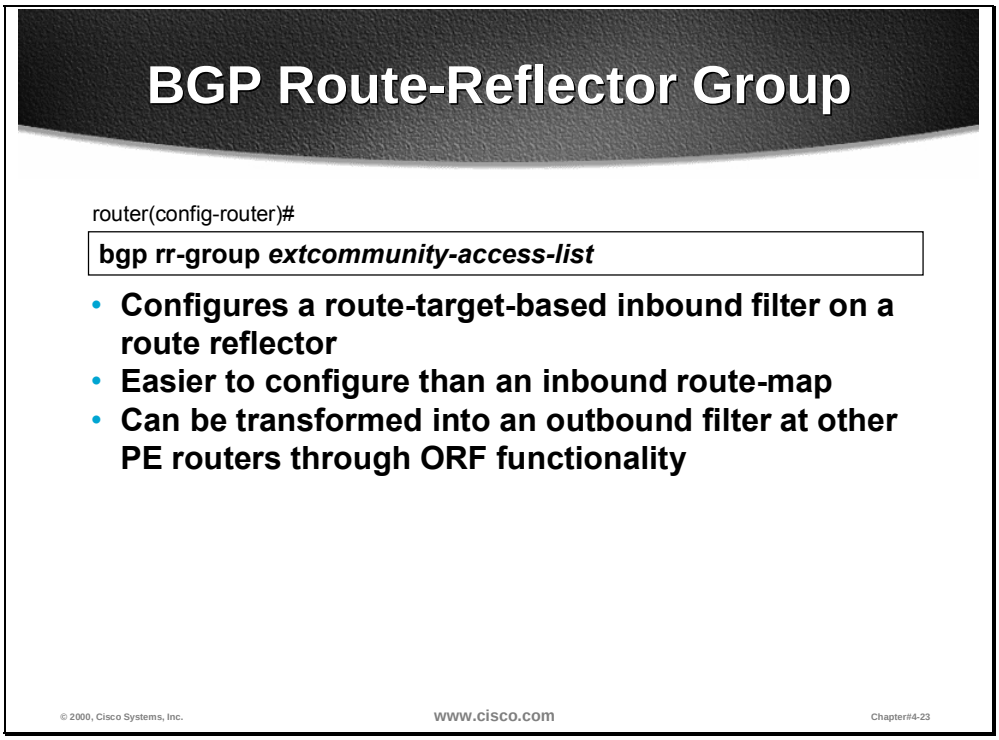

The **bgp rr-group** command is functionally equivalent to a route-map using the same extended community access-list to match routes. There are, however, a number of important differences between them:

- The **bgp rr-group** command is configured for the whole BGP process and applies to all BGP neighbors, introducing configuration consistency.
- The **bgp rr-group** command is easier to configure than a route-map
- The extended community access-list, configured with the **bgp rr-group** command, can be downloaded as an outbound filter to the PE routers. Whereas a route-map based input filter cannot be downloaded through the ORF functionality.

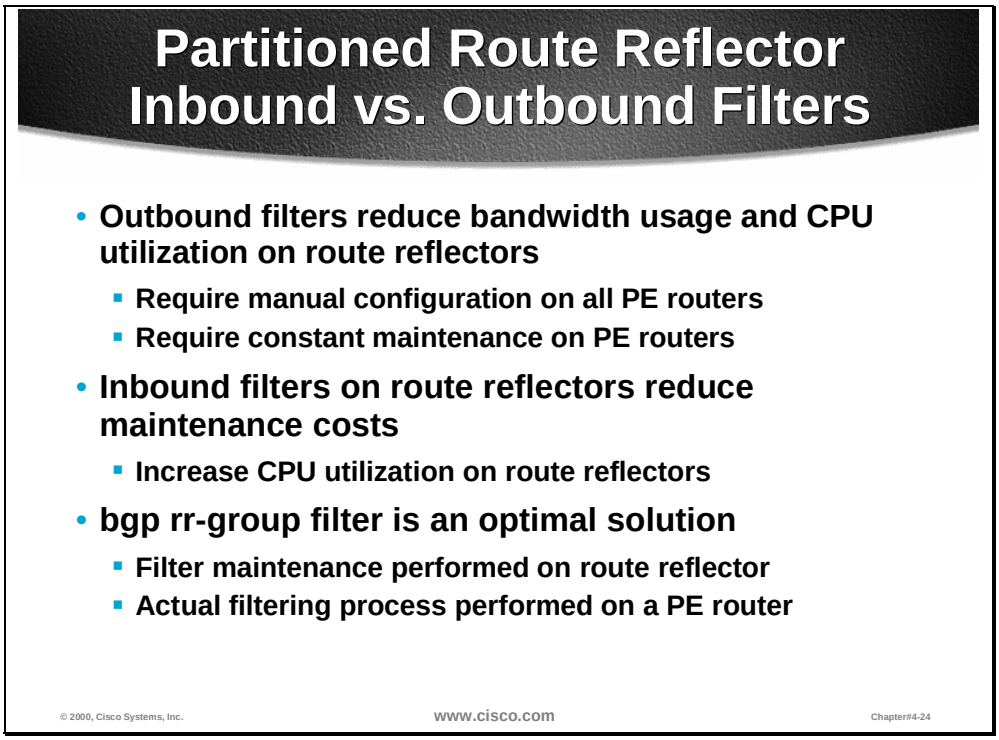

When deciding whether to use outbound route filters on the PE routers or inbound route filters on the route-reflectors to implement partitioned route reflector design, consider the following criteria:

- Outbound filters on PE routers reduce bandwidth utilization and CPU utilization of the route-reflectors (the CPU utilization of the route-reflectors might become an important point when the reflectors carry a large number of routes and serve a large number of clients). However, they require constant maintenance on **all** PE routers and are therefore discouraged from a maintenance and management perspective.
- Inbound filters on route-reflectors are optimal from a maintenance perspective, but increase the CPU utilization of the route reflectors.

The ideal solution (if it can be implemented) is the route-target based filter configured with **bgp rr-group** command, as the maintenance of the extended community access-list is performed on the route-reflector, but the actual filter is downloaded as an outbound filter to the PE routers through the ORF functionality.

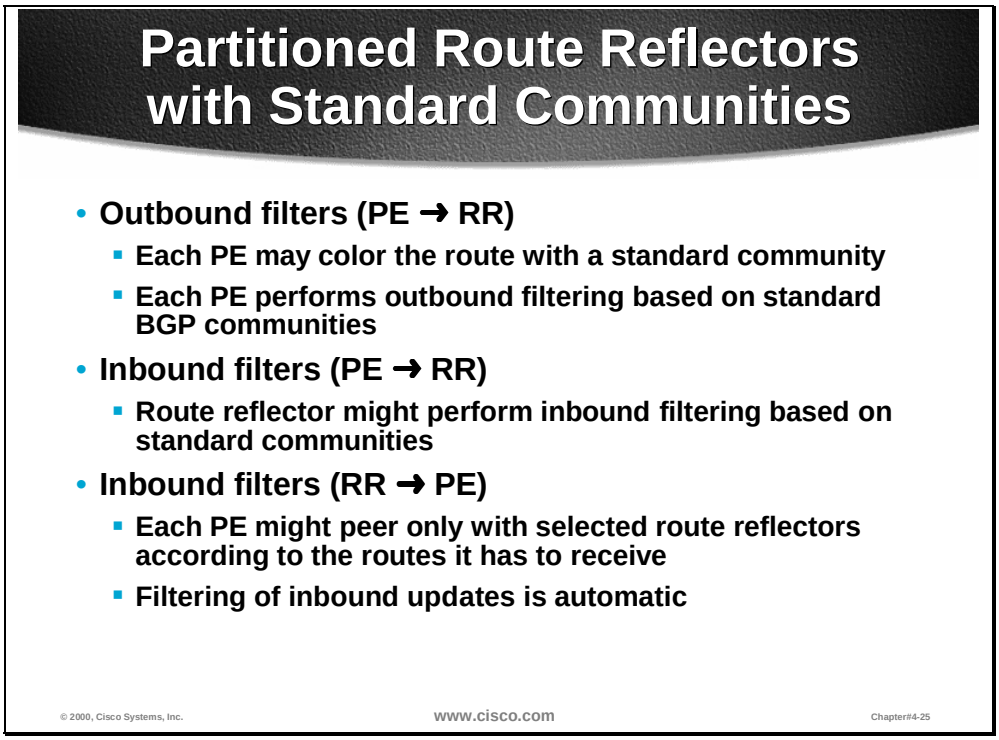

As an alternate solution, the VPNv4 routing information can be partitioned based on standard BGP communities. However, there are a number of different design and implementation methods:

- Filter outbound updates on the PE routers. As the PE routers have to attach standard BGP community to the VPNv4 route anyway, the filtering of outbound VPNv4 routing updates based on the standard BGP community does not represent an additional maintenance burden.
- Attach standard BGP communities to the VPNv4 routes on the PE routers, but perform the filtering on the route-reflectors. This design achieves a clean separation of the marking of customer routes from the partitioning of VPNv4 routing information.
- Configure inbound filters on the PE routers. This design will **not** reduce the amount of routing information on the route-reflectors, but only the number of VPNv4 routes on the PE routers, and is similar to automatic inbound filters based on route-targets. By going one step further, the PE router could peer only with the route-reflectors carrying the desired VPNv4 routes. In this case, there is no need for additional inbound filters.

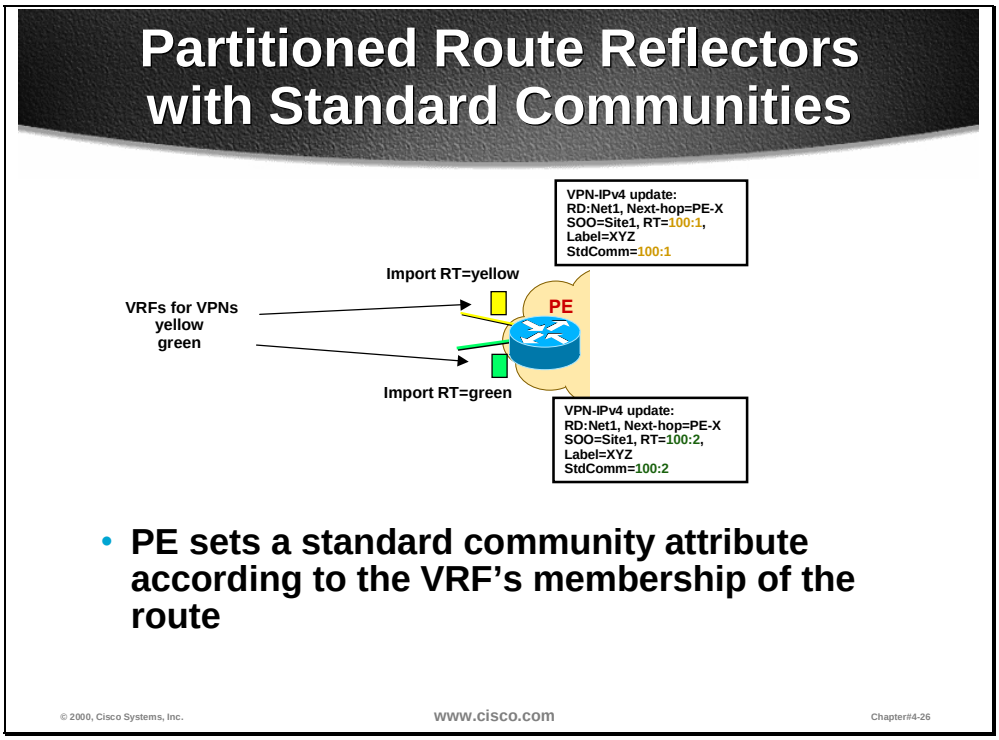

The example above illustrates the utilization of outbound filters on PE routers. As the first step, the PE router sets a standard community attribute to each VPNv4 route. The standard community attached to the route defines the partitioning of the VPNv4 routing information.

The VPNv4 routing information partitioning is usually done at a very low granularity and, therefore, all routes from a VRF would usually have the same community attached.

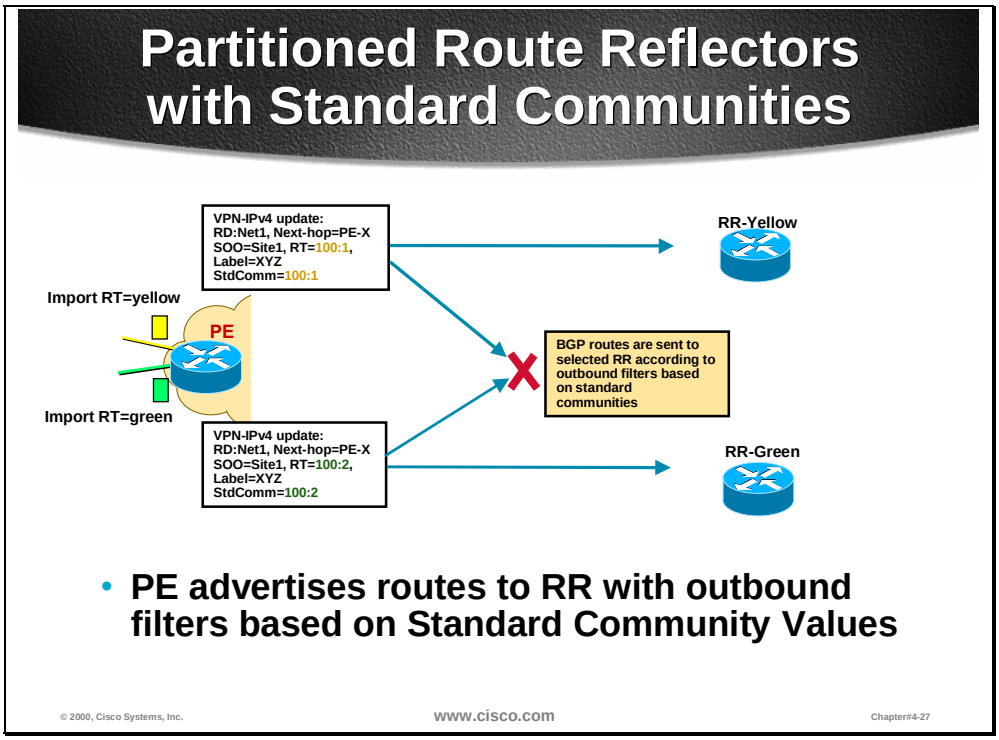

As the second step in this design, the PE router contains outbound route filters that filter VPNv4 routes based on standard BGP community before these routes are sent to the route-reflectors. For example, the route carrying yellow BGP community is not sent to the RR-Green.

Summary

Large MPLS VPN backbones might easily exceed the scaling limits of BGP route reflectors. Further reduction of BGP routing information on any single route reflector, through **partitioned route reflectors**, is therefore needed to facilitate additional growth of the MPLS VPN backbone.

Partitioning of BGP routing information can be performed based on the address-family (separate route-reflectors for IPv4 and VPNv4 routes). Additional partitioning of VPNv4 routing information can be performed based on route-targets attached to VPNv4 routes or any other BGP attribute (for example, standard BGP community). To partition VPNv4 routes based on route-targets, the **bgp rr-group** configuration command provides the optimal means of configuring the partitioning.

Review Questions

- What is the basic function of partitioned route reflectors?
- What are the benefits of partitioned route reflectors?
- Why are partitioned route reflectors needed in very large MPLS VPN backbones?
- How can you implement partitioned route reflectors?
- What are the benefits of using bgp rr-group functionality?
- Why would you choose implementation based on standard BGP communities?
- Why would you choose bgp rr-group implementation?

Chapter Summary

After completing this chapter, you should be able to perform the following tasks:

- Understand the MP-BGP scaling issues in large-scale MPLS VPN backbones
- Describe the built-in scalability mechanisms
- Design and implement networks using partitioned BGP route reflectors

9

MPLS VPN Migration Strategies

Overview

This chapter discusses potential migration strategies from existing IP backbones and existing VPN solutions towards MPLS VPN solutions.

It includes the following topics:

- Infrastructure migration
- Customer migration to MPLS VPN service

Objective

Upon completion of this chapter, you will be able to design the following migration strategies for an MPLS VPN deployment:

- Infrastructure migration strategy for existing IP backbones
- Phased migration strategy for pilot MPLS VPN service
- Migration strategy for customers using layer-2 overlay VPN solutions (Frame Relay or ATM)
- Migration strategy for customer running layer-3 overlay VPN solutions (GRE tunnels or IPSec)

Infrastructure Migration

Objective

Upon completion of this section, you will be able to develop various migration strategies away from existing backbones towards an infrastructure that supports MPLS VPN services.

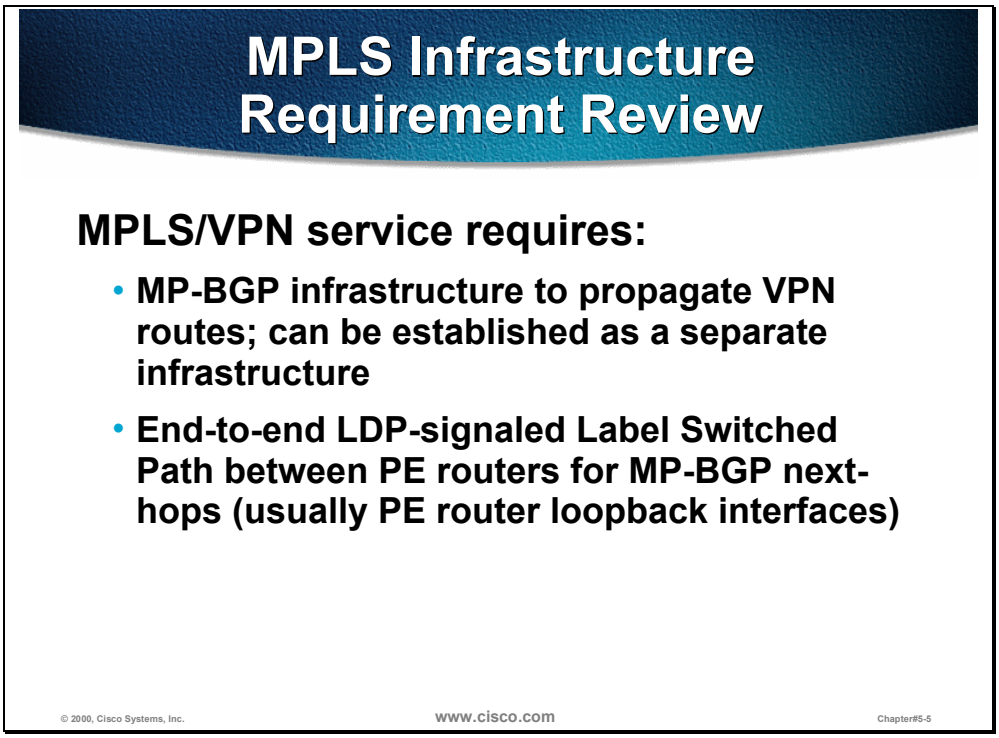

Two basic infrastructure requirements must be satisfied to establish MPLS VPN services in a Service Provider network:

- Multi-protocol BGP (MP-BGP) sessions must be run between Provider Edge (PE) routers. These sessions can be established as a separate infrastructure from the BGP sessions supporting Internet traffic to avoid any migration issues in the network core. Please refer to Chapter 4 of this lesson for more details.
- An End-to-end Label Switched Path (LSP) must be established between the PE routers signaled through Label Distribution Protocol (LDP) or Tag Distribution Protocol (TDP). A LSP must be established, at least, for all next hops of MP-BGP sessions (usually the loopback interfaces of the PE routers).

This section focuses on the migration steps needed to establish LSP between the PE routers.

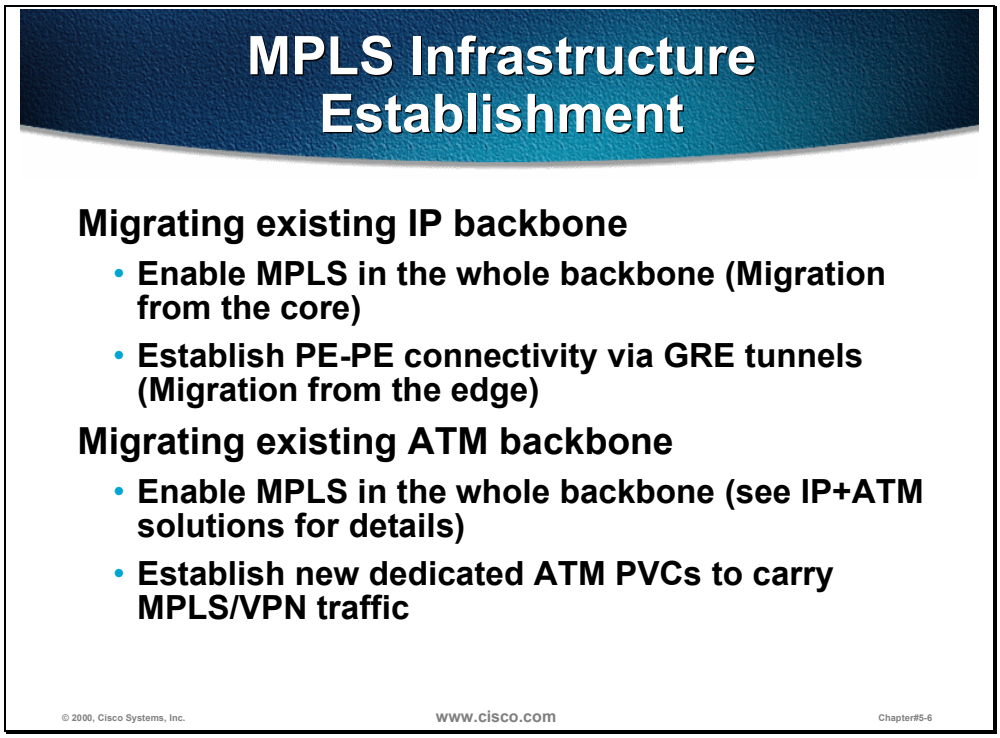

MPLS infrastructure in the network core can be established with a variety of migration strategies. The choice of strategy depends on the layer-2 structure of the existing network core: strategies for ATM-based cores differ from strategies for purely router-based network cores.

In a purely router-based network core, you can choose one of two migration strategies:

- MPLS is enabled in the whole network core (*Migration from core*)
- MPLS is enabled only in edge routers, resulting in disconnected islands of MPLS connectivity. These islands are connected via IP-over-IP tunnels using Generic Route Encapsulation (GRE) tunneling protocol (*Migration from edge*)

In an ATM-based network core, you can also choose one of two migration strategies:

- MPLS is enabled in the whole ATM network. (*Migration from core*)**.** Please refer to IP+ATM solution training for more details on this migration strategy.
- Additional Permanent Virtual Circuits (PVCs) are established directly between islands of MPLS connectivity (*Migration from edge*). Existing permanent virtual circuits can also be reused for this purpose.

Note Some service providers use single-protocol encapsulation (called AAL5MUX in Cisco IOS) on ATM virtual circuits in their core. This encapsulation type does not support concurrent IP and MPLS traffic and has to be changed to AAL5SNAP encapsulation prior to MPLS deployment.

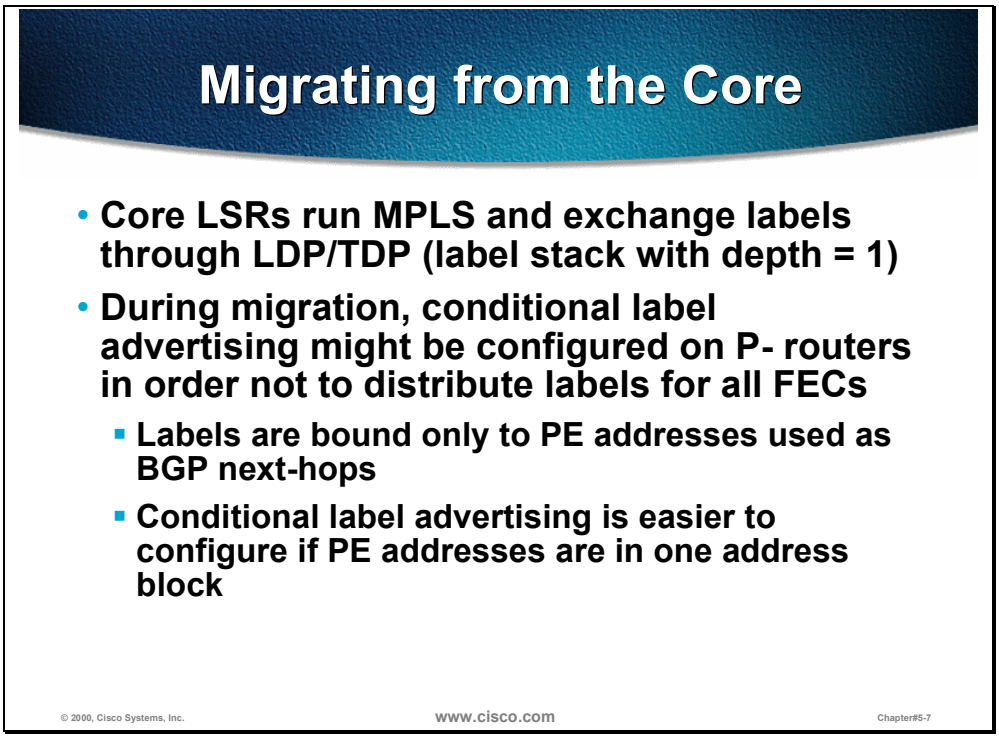

If you choose a *Migration from Core* strategy in your MPLS VPN deployment, you have to start LDP or TDP on all core routers and configure MPLS on all core interfaces. This operation might interfere with your existing IP traffic and you might decide to use conditional label advertising to prevent that.

With conditional label advertising, you can distribute labels only for selected destinations in your network (for example, only BGP next-hops of the PE routers). The IP traffic toward the other destinations will not be labeled, as the ingress routers would not receive labels for those destinations from their downstream neighbors.

Note Conditional label advertising for selected destinations is easier to achieve if these destinations are in one address block (and thus easily covered with an IP access list). It's therefore recommended that you assign loopback addresses of the PE routers from one address block.

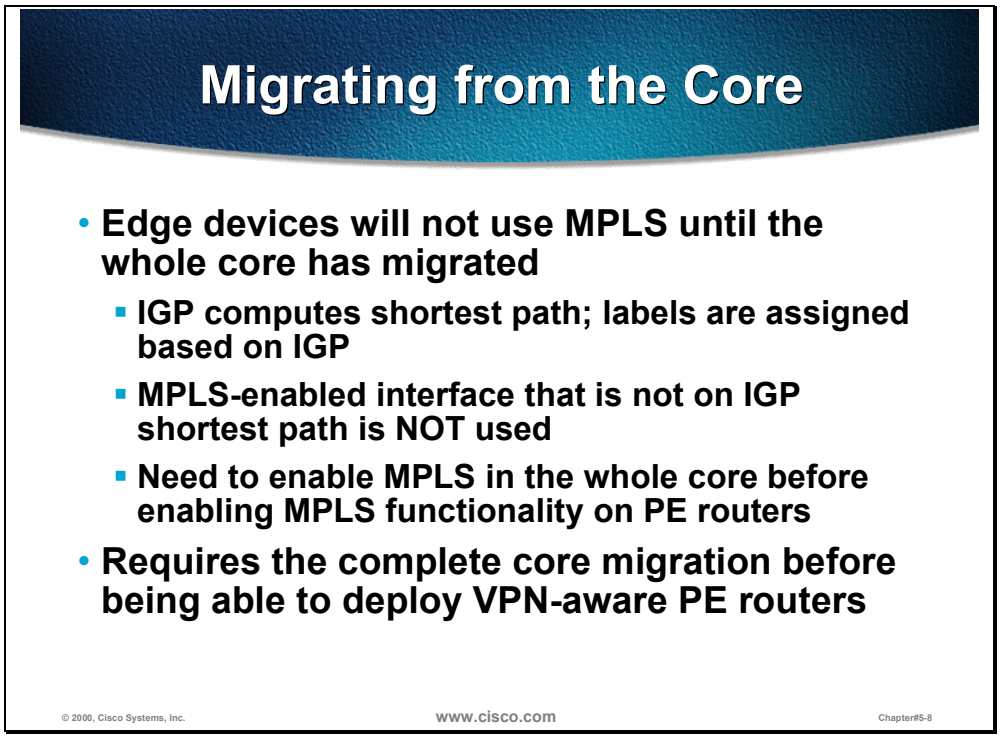

There are a number of caveats associated with *Migration from Core* strategy:

- LDP or TDP labels are assigned solely based on the contents of an IP routing table, which is driven by Interior Gateway Protocol (IGP) used in the network backbone.
- If the IP routing table directs traffic toward a PE router via an interface that is not MPLS-enabled, the label switched path toward that PE router is broken. There is no mechanism in TDP or LDP that allows MPLS traffic to avoid non-MPLS links if these links are in the IGP shortest path.

Note MPLS traffic could be redirected around non-MPLS-enabled parts of the network core, even if they are on IGP shortest path, by using MPLS Traffic Engineering. However, this solution is best avoided, as it unnecessarily increases the network complexity.

- MPLS-enabled interfaces that are not on IGP shortest path are not used for MPLS traffic forwarding.

In summary – when you use *Migration from Core* strategy, MPLS must be enabled on all core routers and on all interfaces in the IGP shortest path between the PE routers before you can start deploying MPLS VPN services.

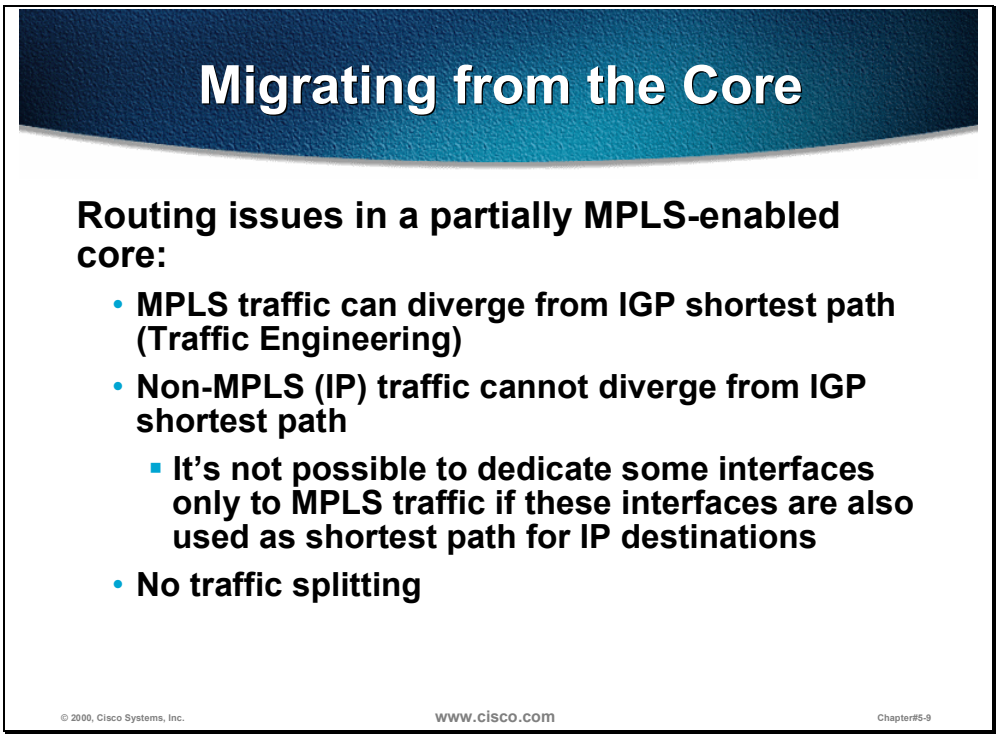

Some network designers would like to deploy MPLS Traffic Engineering in combination with the MPLS VPN services to optimize their backbone utilization. This goal is hard to achieve in backbones where the conditional label advertising has been implemented to minimize the impact of migration toward MPLS VPN because:

- While MPLS VPN traffic (or other labeled traffic) can diverge from the IGP shortest path by means of MPLS Traffic Engineering, the non-labeled traffic (pure IP traffic) cannot. It is therefore not possible to dedicate some interfaces to MPLS traffic (for example, additional links deployed to support MPLS VPN service) if these interfaces happen to be on IGP shortest path toward other IP destinations. As an intermediate step, IGP cost on these interfaces could be increased to discourage IGP from selecting them.
- As the non-labeled traffic is forwarded based only on IP routing tables, not on MPLS Traffic Engineering trunks established in the network core, it is hard to achieve traffic splitting between MPLS VPN and Internet traffic without deploying complex MPLS Traffic Engineering schemes for MPLS VPN traffic.

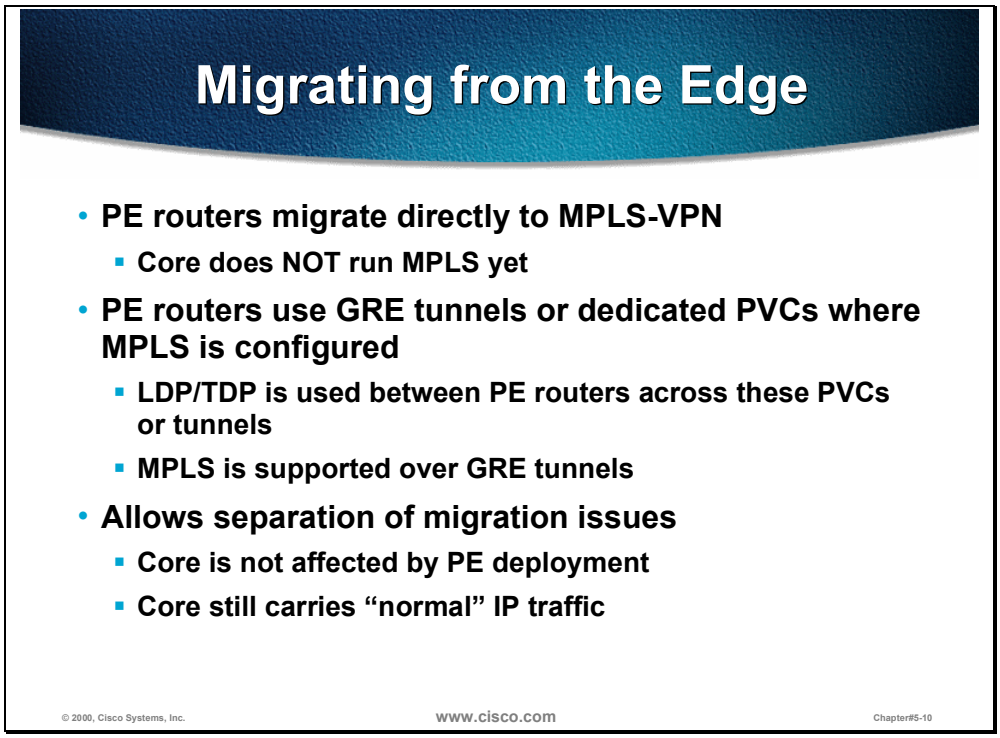

Migration from Edge strategy to deploying MPLS VPN services is easier and quicker to implement, as it does not involve reconfiguration of core devices in your network. The PE routers are MPLS-enabled and dedicated point-to-point links are used between the PE routers (or small islands of MPLS connectivity at the edges of the network) to enable MPLS transport across the network core. LDP or TDP is then run over these new point-to-point links to establish Label Switched Paths between PE routers.

The new point-to-point links needed to support MPLS connectivity across non-MPLS backbone can be implemented with ATM Virtual Circuits in ATM-based backbones or with IP-over-IP tunnels using Generic Route Encapsulation (GRE) technology.

The *Migration from Edge* strategy enables clear separation of migration issues, as the network core is not affected by MPLS VPN deployment and is still able to carry non-labeled IP traffic (for example, Internet traffic).

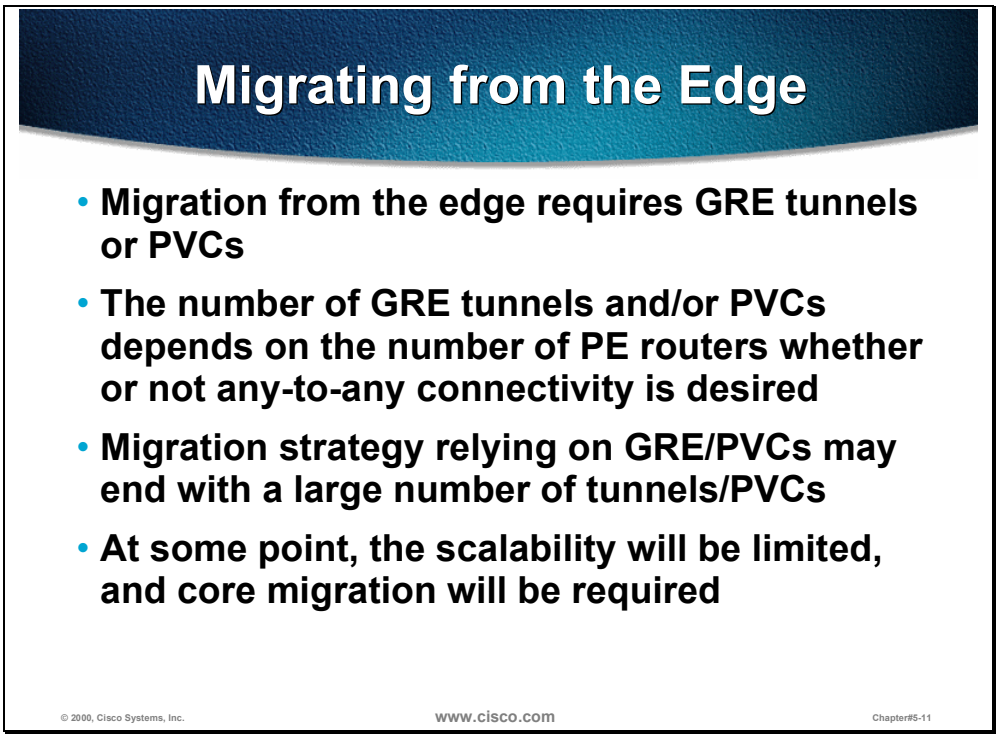

The *Migration from Edge* strategy, while easy to implement in a pilot network, suffers from severe scalability constraints.

The strategy requires point-to-point links between islands of MPLS connectivity. The number of these links depends on the number of PE routers, desired traffic pattern and potential requirement for optimal MPLS VPN traffic forwarding across the backbone. In most cases, the end result would be a full-mesh of GRE tunnels (or ATM virtual circuits), which is clearly not a scalable solution.

The scalability constraints of *Migration from Edge* strategy will eventually force anyone deploying this strategy to revert to *Migration from Core* strategy once the MPLS VPN service enters the production phase.

Note The Migration from Edge strategy also suffers from encapsulation overhead when implemented with the GRE tunnels. Every MPLS VPN packet propagated across the network core within a GRE tunnel incurs a 20-byte overhead of the IP and GRE header.

Summary

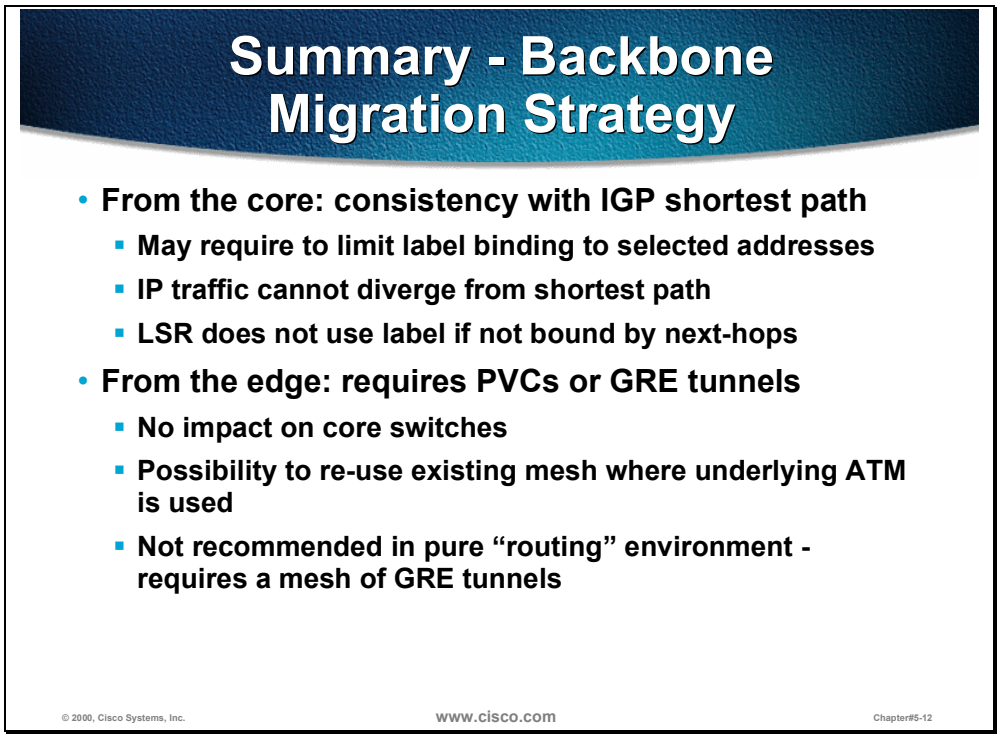

There are two basic migration strategies that can be used to establish MPLS connectivity as required by MPLS VPN service across network core:

- *Migration from Core* where end-to-end MPLS connectivity is established before the MPLS VPN service is deployed. The impact of this strategy on existing IP traffic can be minimized with deployment of conditional label advertising, but this technique prevents you from applying additional MPLS services (for example, MPLS Traffic Engineering) to your IP traffic.
- *Migration from Edge* where the small islands of MPLS connectivity on the network edge are connected via point-to-point links. This strategy has no impact on core switches and might be an optimal strategy in ATM environments where the full-mesh of ATM virtual circuits is already established between the edge routers. It should only be used for pilot projects in the router-based backbones, as it requires a mesh of GRE tunnels in order to enable MPLS transport across an IP backbone.

Review Questions

- How can you minimize the effect of core migration to MPLS for regular IP traffic?
- Can you allocate labels only to PE loopback addresses if you are using an ATM core?
- What are the benefits of edge-first migration toward MPLS infrastructure?
- What are the drawbacks of edge-first migration toward MPLS infrastructure?
- Which migration strategy is better suited for early MPLS VPN pilots?
- Which migration strategy is better for a large-scale MPLS VPN rollout?

Customer Migration to MPLS VPN service

Objective

Upon completion of this section, you will be able to develop migration strategies for the following customer types:

- Customers using layer-2 overlay VPN
- Customers using layer-3 overlay VPN
- Customers using IPSec-based VPN
- Customers using L2F-based VPN
- Customers using routing protocols that are not supported as PE-CE routing protocols

Generic Customer Migration Strategy

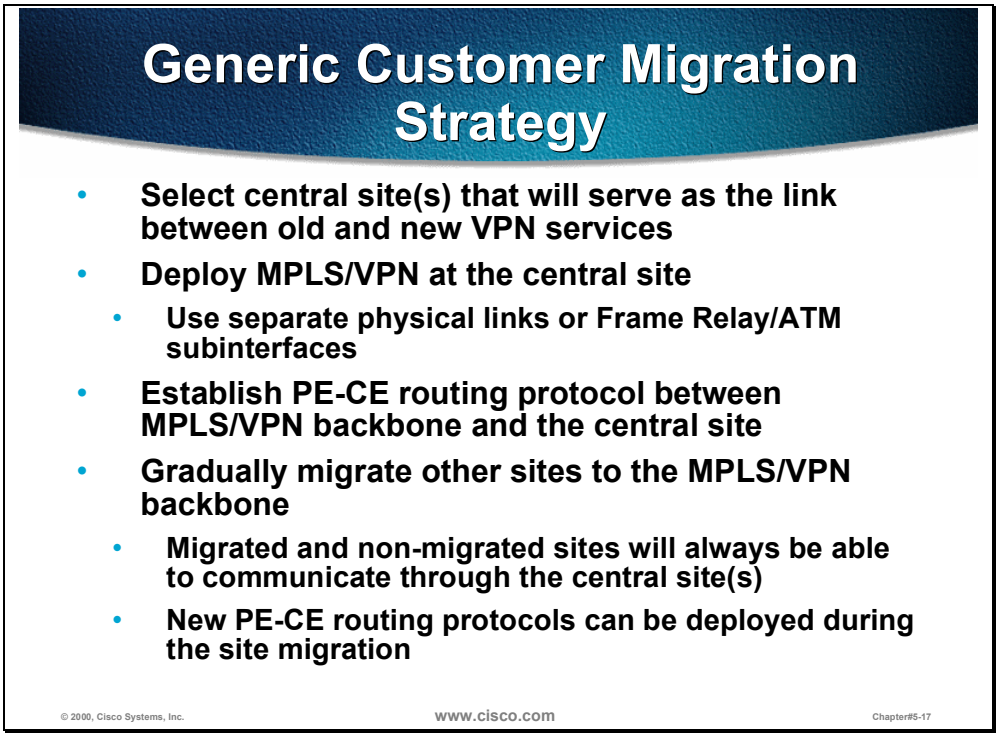

This section will discuss migration of existing VPN customers who might use a variety of overlay VPN technologies, ranging from layer-2 VPNs (Frame Relay, ATM) to IP-over-IP based overlay VPNs, toward an MPLS VPN service.

Whatever the current VPN technology, customer migration must be performed according to the following principles:

- The migration should have minimal impact on customer connectivity and traffic forwarding
- The migration should be performed gradually – it is impossible to migrate a large customer in one giant step
- Each step in the migration process should be easy to test and validate
- There should be an easy and quick fallback plan for each migration step – the customer connectivity should be easily and quickly reestablished if a particular migration step fails

The remainder of this section provides migration examples that conform to the principles outlined above. Each example is based on a common migration strategy involving four broad steps adapted to the particular overlay VPN technology the customer is using.

These are the four steps used in each of the examples:

Step 1 A site (or several sites) is selected to act as a transit site between the old and new VPN service. All traffic between sites using the old VPN technology and sites using the new VPN technology will flow through this transit site.

Note The transit site becomes a single point of failure during the migration process. It is therefore crucial that this site has redundant connectivity to the Service Provider network. Alternatively, you could deploy several transit sites to reach the desired level of redundancy.

Step 2 The transit sites are connected to the MPLS VPN backbone to enable forwarding of transit traffic between the old and the new VPN backbone. Parallel physical links could be used for this purpose or you could deploy additional Frame Relay or ATM Virtual Circuits (VC).

Step 3 The PE-CE routing protocol is established between the MPLS VPN edge routers and the customer routers. The customer network is now ready for site-by-site migration.

Step 4 Every customer site is migrated to the new backbone independently. The connectivity between the sites connected to the old backbone and the sites already migrated to the MPLS VPN backbone is provided through the transit sites. New routing protocols supported by Cisco's MPLS VPN implementation can be deployed on the customer sites during this migration process to replace non-supported routing protocols (for example, EIGRP or IS-IS).

After all sites have been migrated to the new MPLS VPN backbone, the connectivity between the transit sites and the old backbone can be removed.

Migration From Layer-2 Overlay VPN

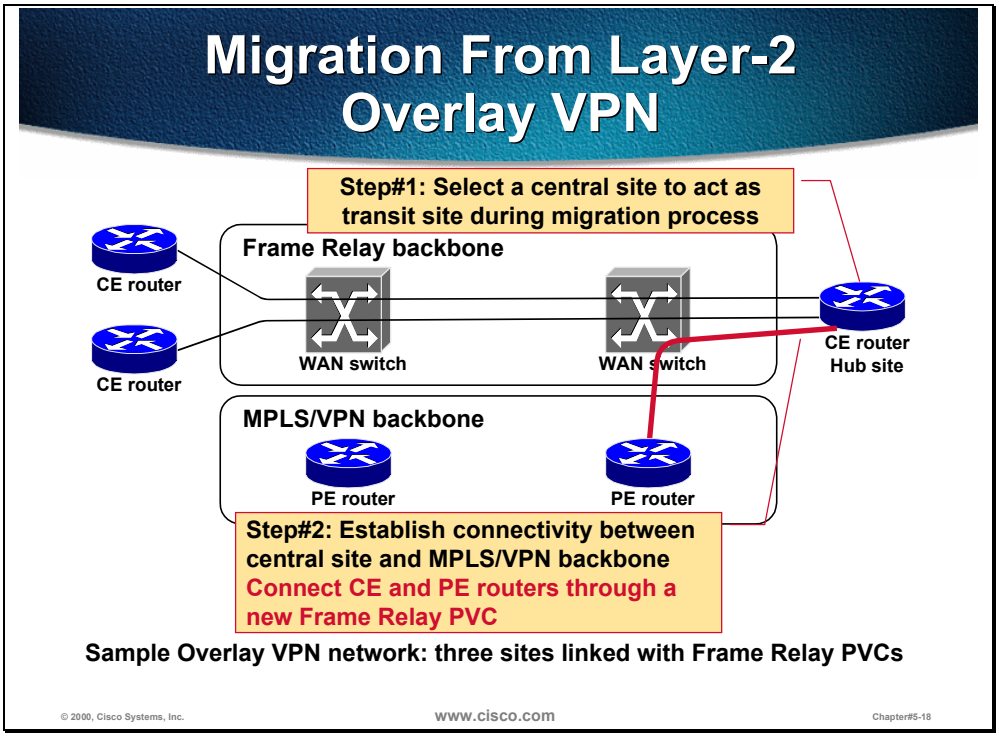

This first example will consider one of the easiest migration scenarios which is the migration of customers using layer-2 overlay Virtual Private Networks.

Note Throughout this migration scenario, we're assuming that the Service Provider has two backbones – Frame Relay backbone that provides existing VPN services and MPLS VPN backbone that is being deployed.

In the first preparatory steps, additional Frame Relay PVC is established between the CE routers at the transit sites and the closest PE router. The target PE-CE routing protocol is deployed on this new virtual circuit.

Note Route redistribution between the existing customer routing protocol and the new PE-CE routing protocol needs to be configured if the two routing protocols are not the same. The routing protocol migration will not be covered in individual migration examples, as it is covered in the last example of this section.

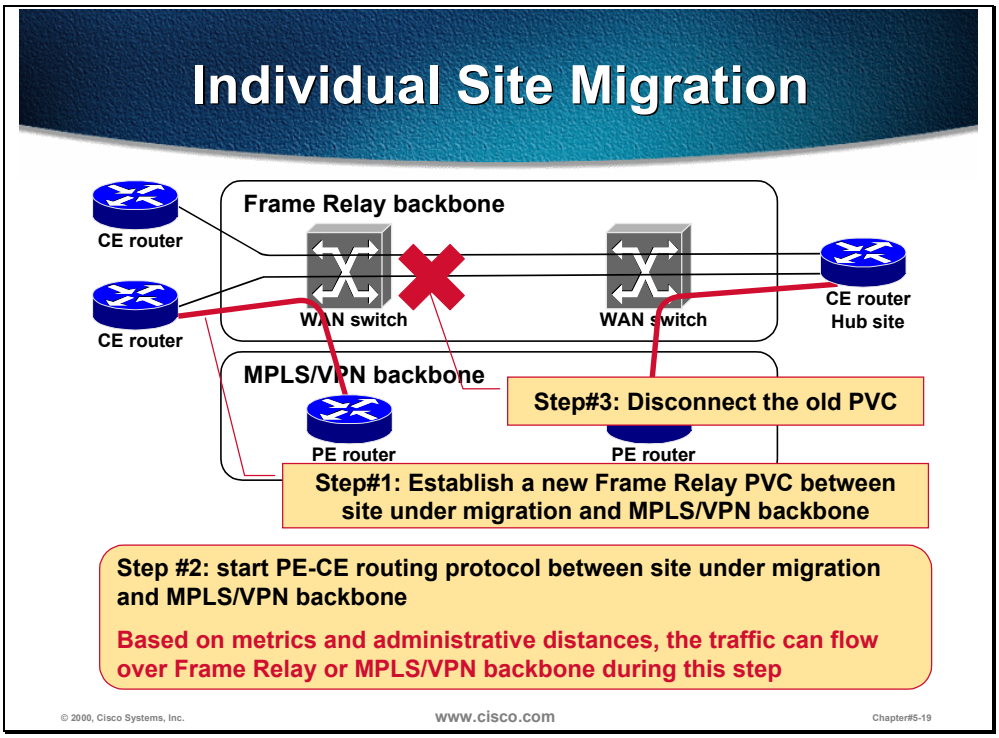

After the transit sites have been connected to the MPLS VPN backbone, individual site migration can start. Each customer site is migrated with three steps:

Step 1 New Frame Relay PVC is established between the site to be migrated and the nearest PE router.

Step 2 The target PE-CE routing protocol is deployed between the site to be migrated and the nearest PE router. Routing tables (or topology databases) on PE and CE routers can be examined at this point to verify proper route propagation across the MPLS VPN backbone.

Traffic between the transit sites and the site under migration can flow over the Frame Relay backbone or over the MPLS VPN backbone, based on the configured IGP metrics or on administrative distances of the deployed routing protocols. For example, if the customer uses OSPF as the routing protocol, but wishes to migrate to BGP as part of MPLS VPN migration, all the traffic will start to flow over the MPLS VPN backbone immediately as EBGP has lower administrative distance than OSPF.

Step 3 The Frame Relay PVC between CE routers is disconnected. Following the routing protocol convergence, the connectivity between the transit sites and the site under migration should be reestablished over the MPLS VPN backbone.

Note The fallback scenario for this step is very simple – re-enabling the PVC between the CE routers will reestablish overlay VPN connectivity.

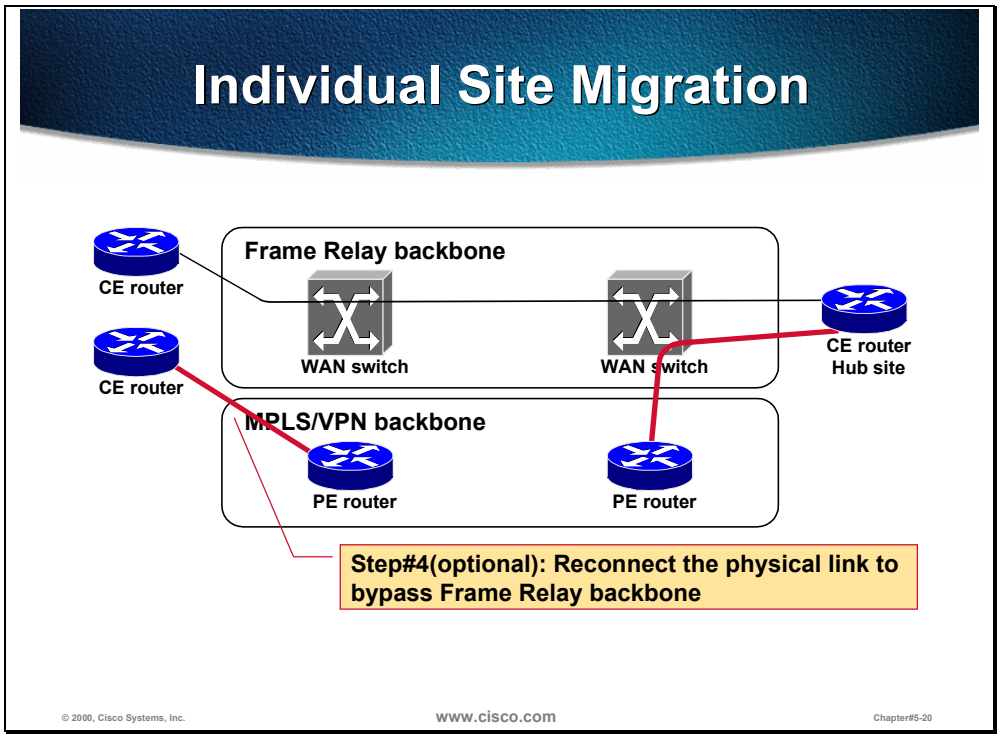

As the last migration step, the site migrated to the MPLS VPN backbone can be completely disconnected from the Frame Relay backbone and connected directly to the PE router. The decision whether to perform this migration step is based primarily on the access method the service provider is using for MPLS VPN service and the relative location of the Frame Relay switches and the PE routers. A few examples are listed below:

- If the service provider uses Frame Relay as the access backbone for MPLS VPN service, this step is not necessary.
- If the Frame Relay switches and PE routers are co-located and the link between the CE router and the Frame Relay switch is a physical link, the transition might be desired, but would require a physical intervention at the provider Point-of-Presence (POP).
- If the CE router is connected to the TDM access backbone, the switchover to PE router requires only reconfiguration of the TDM equipment. In some cases, the PE routers might be even closer to the customers than the Frame Relay switches.

Migration from GRE Tunnel-Based VPN

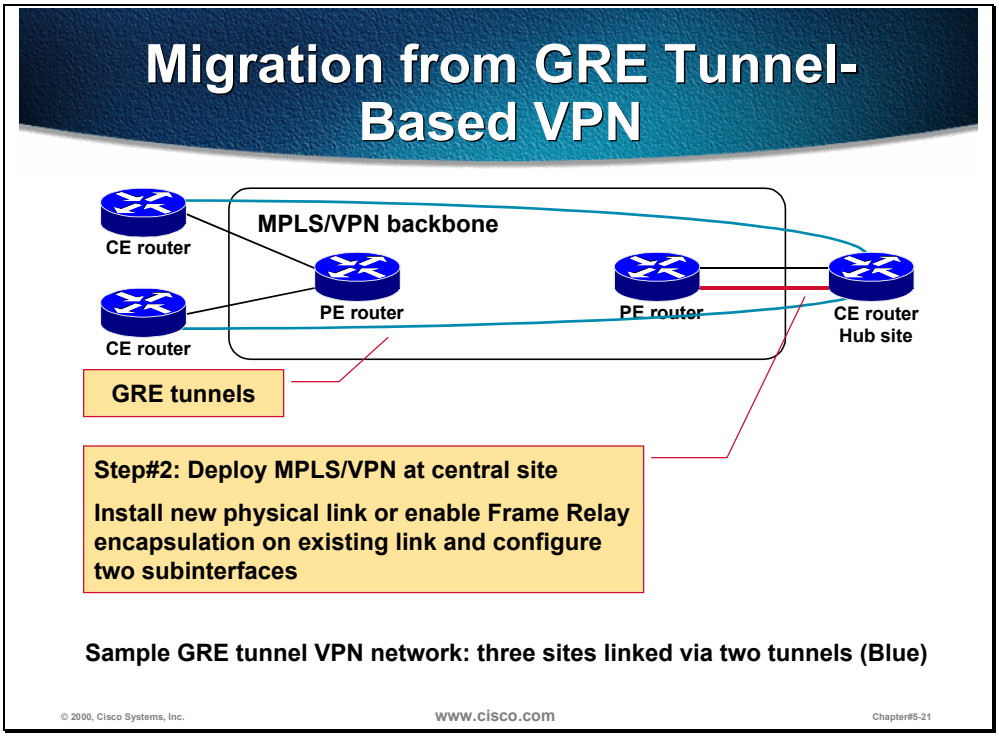

The second migration scenario describes migration of a customer using IP infrastructure of the Service Provider to establish a Virtual Private Network via IP-over-IP tunnels.

Note In most cases, the provider edge routers providing IP connectivity are the same routers that will provide MPLS VPN service in the future.

After the transit sites are selected, they need connectivity to the MPLS VPN backbone. This connectivity can be achieved by installing a new physical link between the PE and the CE router or by deploying Frame Relay encapsulation on existing link and configuring two subinterfaces. Please refer to Chapter 2 of this lesson for more information on combining Internet access and MPLS VPN service over the same physical link.

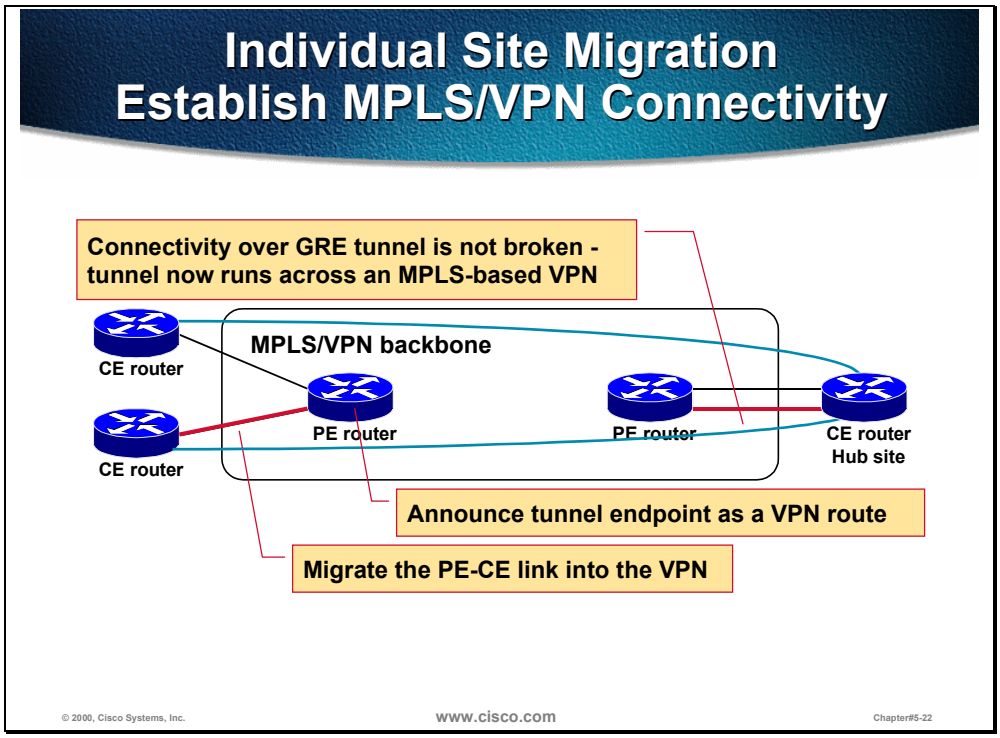

Assuming that the PE routers providing the existing IP connectivity are the same routers as the ones providing MPLS VPN service, the migration of an individual site into the MPLS VPN backbone is very easy – the interface on the PE router is put into the desired VRF and the IP address that was previously used on the interface is reconfigured. However, the GRE connectivity between the migrated site and the transit site is broken at this point.

To reestablish connectivity between the migrated site and the transit site, the tunnel endpoint of the migrated site is configured as a static route (or connected interface) in the VRF into which the site was migrated. The tunnel endpoint thus appears as being reachable through the MPLS VPN backbone by the transit site. As well, the GRE tunnel between the sites is reestablished, resulting in unhindered customer connectivity.

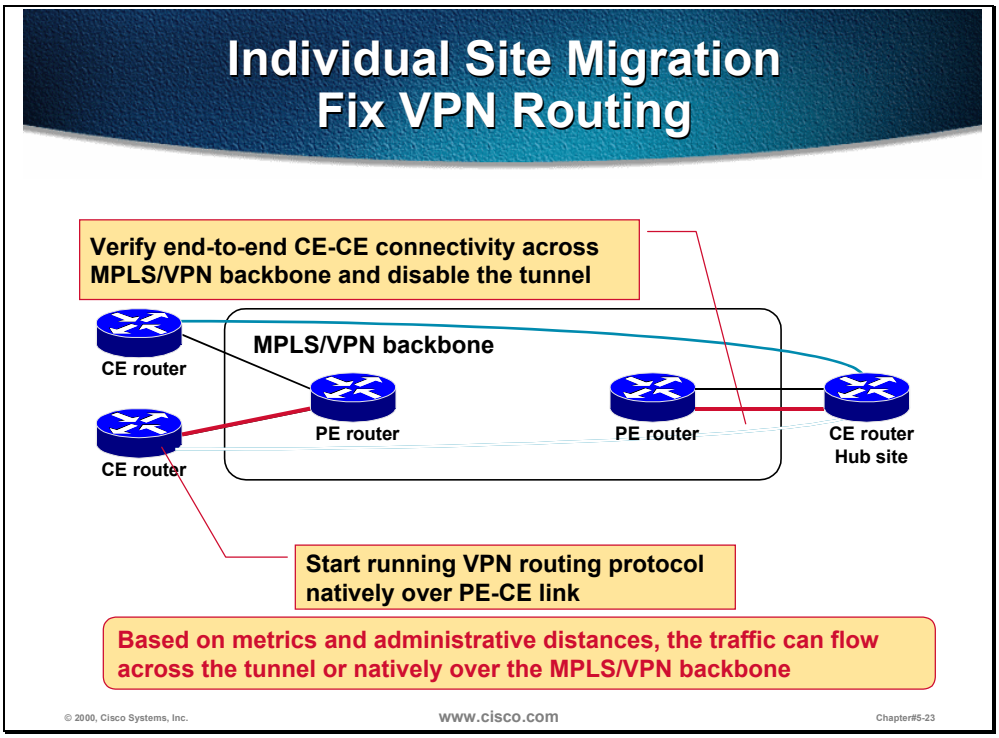

As the last migration step for customers using IP-over-IP tunnels, the routing between customer sites should be migrated from GRE tunnels to the native routing of the MPLS VPN backbone. This migration is performed in the same way as the migration for Frame Relay customer in the previous example:

Step 1 PE-CE routing protocol is started between the migrated site and the PE router to which it is connected. Routing tables are verified on the PE routers to make sure that the customer routes are propagated across MPLS VPN backbone.

Step 2 Tunnel interface is shutdown on the CE router. After the transit site CE router detects that its neighbor is no longer reachable over the GRE tunnel, the IP routing will re-converge based on the new information received from the MPLS VPN backbone. The connectivity between the migrated site and the transit site should be reestablished.

Step 3 Tunnel interfaces are removed from the CE routers on the migrated site and the transit site.

Migration from IPSec-Based VPN

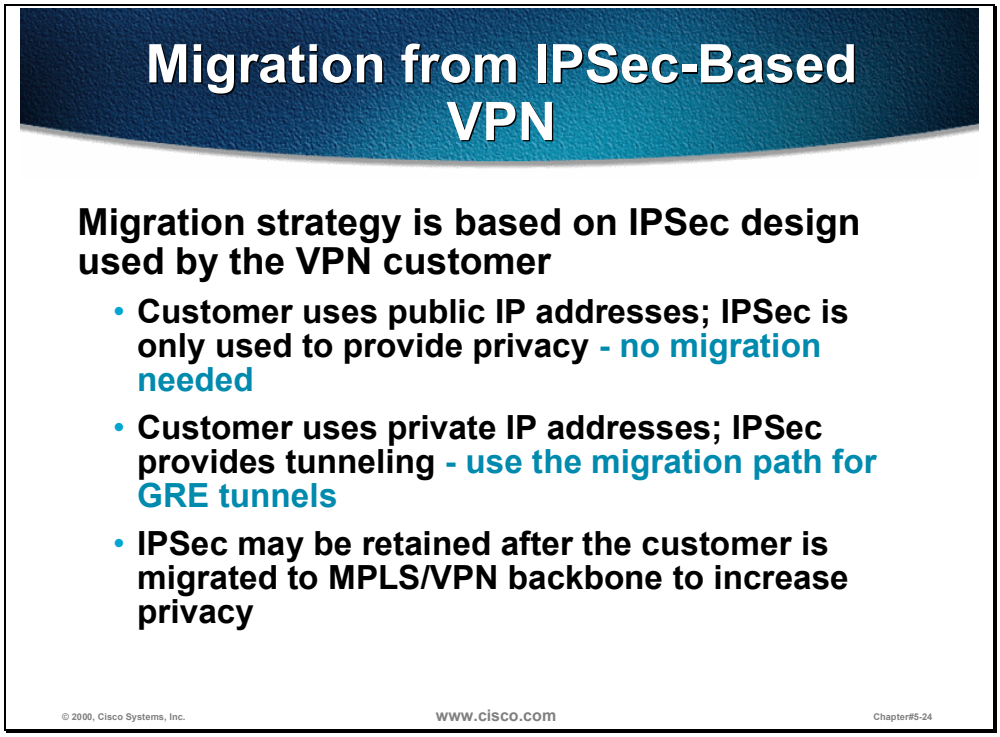

IP Security (IPSec) is used by many customers deploying Virtual Private Networks over the public IP infrastructure (for example, Internet) as the preferred VPN technology because it provides strong authentication and encryption.

Different migration paths toward MPLS VPN service can be used for customers using IPSec based on how they use the IPSec technology and their addressing structure:

- Customers using IPSec in transport mode have to use public IP addresses. These customers use IPSec only to ensure privacy over the public IP backbone. They can retain their IPSec setup between CE-routers even when they are migrated to MPLS VPN solution. Future releases of IOS will allow you to map an IPSec session through the MPLS cloud and terminate the IPSec session on the PE-router. When this is available, the customer will have the advantage of encryption and the SP will have the advantage of MPLS scalability.
- Customers that use IPSec in encapsulation mode are very similar to customers using IP-over-IP GRE tunnels (the only difference is in the technology they use for IP tunneling). These customers can be migrated to MPLS VPN backbone using the steps already outlined in the previous example.

After a customer using IPSec is migrated to MPLS VPN backbone, IPSec configuration might be retained to even further increase the privacy of customer data.

Migration from L2F-Based VPN

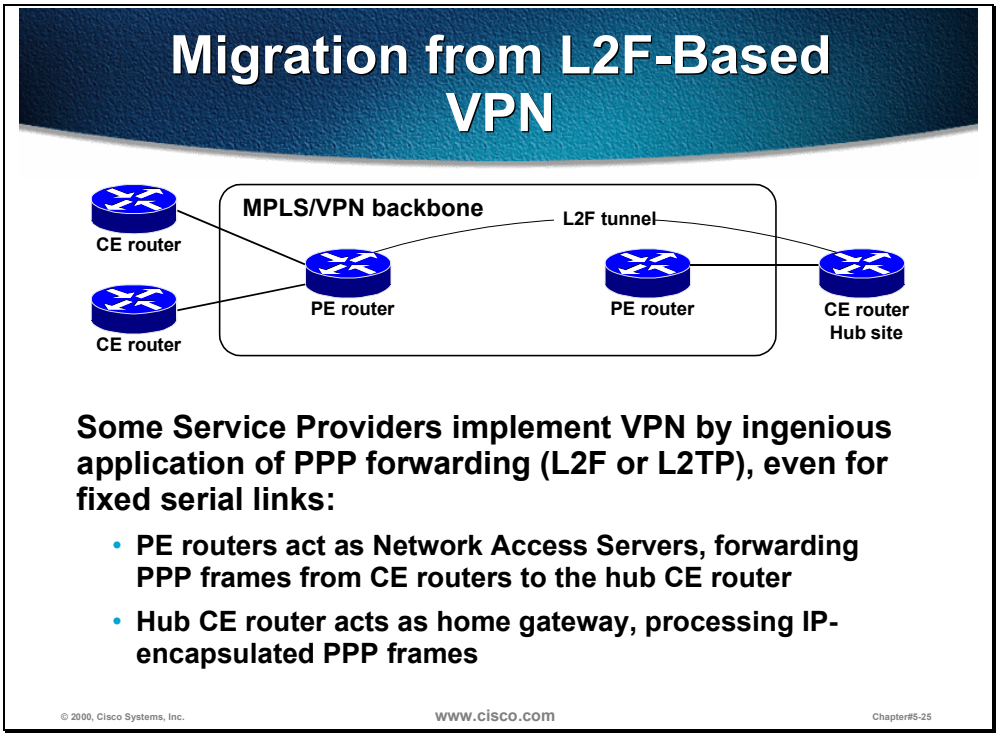

Faced with the complexities of IPSec or scalability issues of GRE tunnels, some Service Providers have started providing VPN service based on PPP forwarding technologies (Layer 2 Forwarding – L2F or Layer 2 Transport Protocol – L2TP). In these implementations, PE routers act as Network Access Servers (NAS), forwarding PPP frames received from point-to-point links between PE and CE to the central customer router, which acts as a home gateway.

Note When using this VPN implementation method, all the traffic from a particular site has to reach the customer home gateway first to be analyzed and forwarded to another site. The customer home gateway thus acts as a transit site between all customer sites.

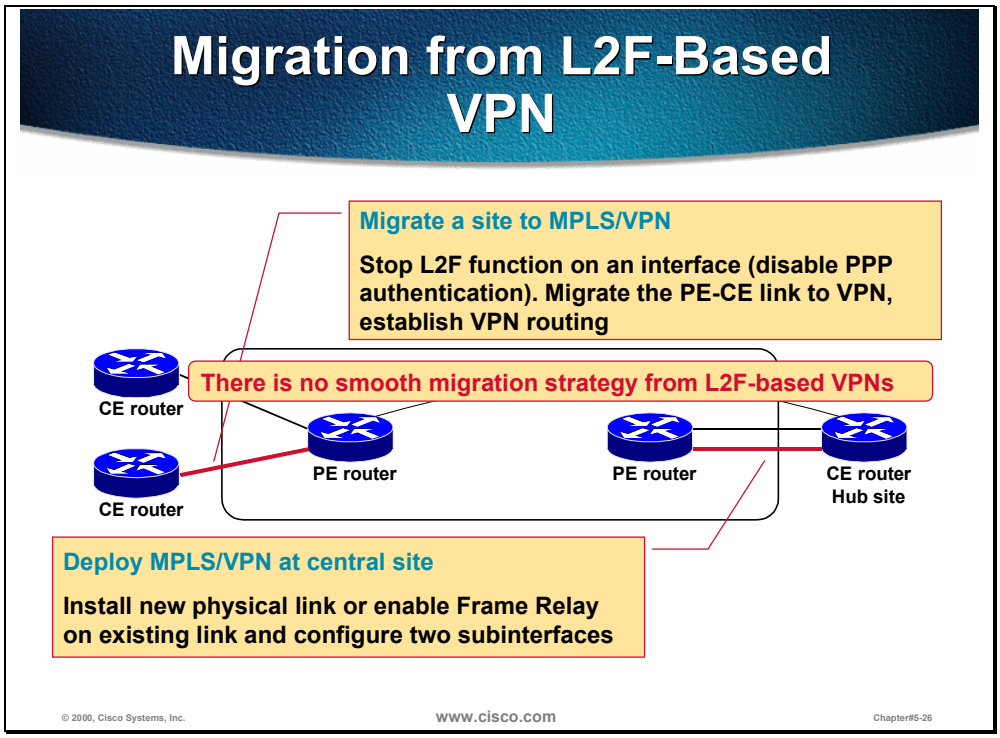

The migration of a customer using L2F or L2TP as the VPN-enabling technology toward MPLS VPN follows the same generic steps as the migration of a customer using GRE tunnels. This migration, however, is not as smooth as the migration of other customer types for these reasons:

- Frame Relay encapsulation cannot be configured on the link toward the migrated site as this would break the existing VPN connectivity.
- L2F or L2TP tunnel cannot be reestablished once the customer site has been migrated to the MPLS VPN backbone as the tunnel originates on the PE router (and thus in global address space), not on the CE router as IPSec or GRE tunnels do.

The migration of an individual site thus has to be performed in a single step. The VPN connectivity between the migrated site and the transit site will only be established after the routing protocol has been started between the PE router and the CE router of the migrated site and the routes have been propagated across MPLS VPN backbone.

Migration From Unsupported PE-CE Routing Protocol

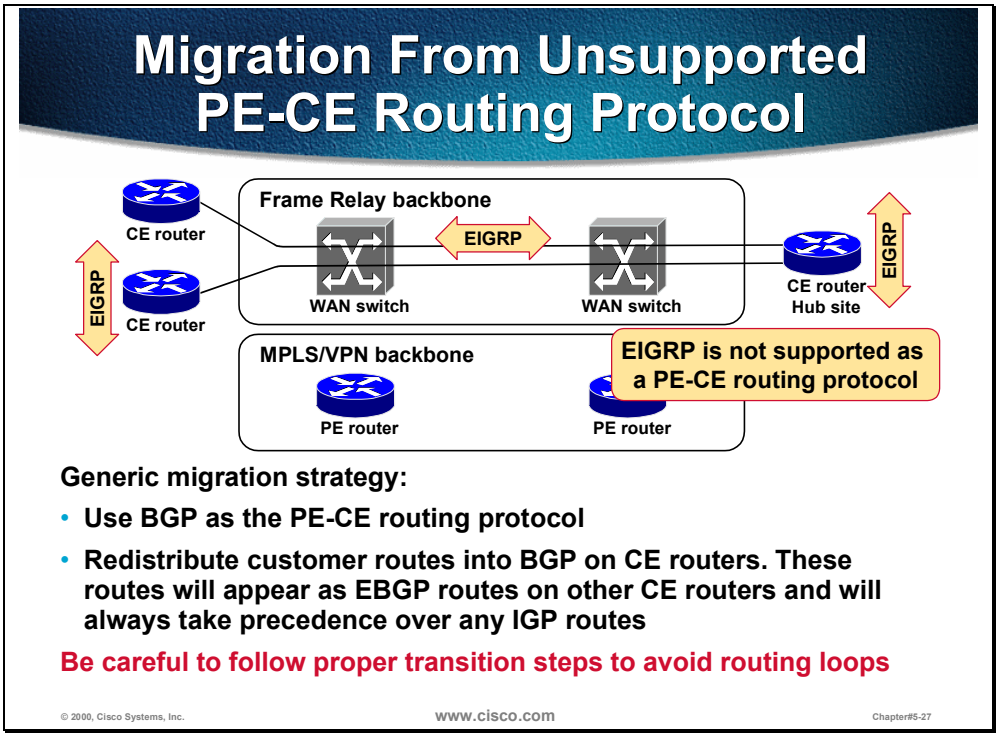

The previous migration scenarios have covered the replacement of virtual point-to-point links implemented with a variety of VPN technologies with the MPLS VPN service. None of them considered migration of a customer that is using a routing protocol not supported as a PE-CE routing protocol by Cisco IOS MPLS VPN implementation between the customer sites. IP routing protocols currently not supported as PE-CE routing protocols are IS-IS, EIGRP and RIP version 1.

Note OSPF is a supported PE-CE routing protocol. However, its CPU requirements, memory usage and the need to have an independent OSPF process for each Virtual Routing and Forwarding (VRF) table limit its usage to special cases.

The routing protocol not supported as a PE-CE routing protocol cannot be used between the CE-routers and PE-routers when the customer is migrating to an MPLS VPN backbone. The routing between CE-routers and PE-routers has to be migrated toward a supported protocol, while the customer can still retain the previous routing protocol within each individual site.

The routing protocol migration is independent of physical connectivity migration and can be performed in parallel. Furthermore, the same migration steps can be used regardless of the VPN technology currently used by the customer.

In the example discussed here, the customer is using EIGRP protocol everywhere in the existing VPN network:

- At the central site to distribute routes to other routers located at the central site.
- At all the other sites to exchange routes between routers located at these sites.
- Between the sites over the VPN backbone.

Whenever you need to change routing protocol during the migration toward an MPLS VPN backbone, BGP should be used as the target routing protocol between CE-routers and PE-routers as the administrative distance of EBGP is lower than the administrative distance of any other protocol, guaranteeing that the IP routing tables are always built from BGP information. Additional BGP attributes can also be used to prevent redistribution loops during the migration.

Note Migrating from one routing protocol to another usually involves complex two-way redistribution that can easily result in routing loops. Make sure that you closely follow the steps in this example to prevent them.

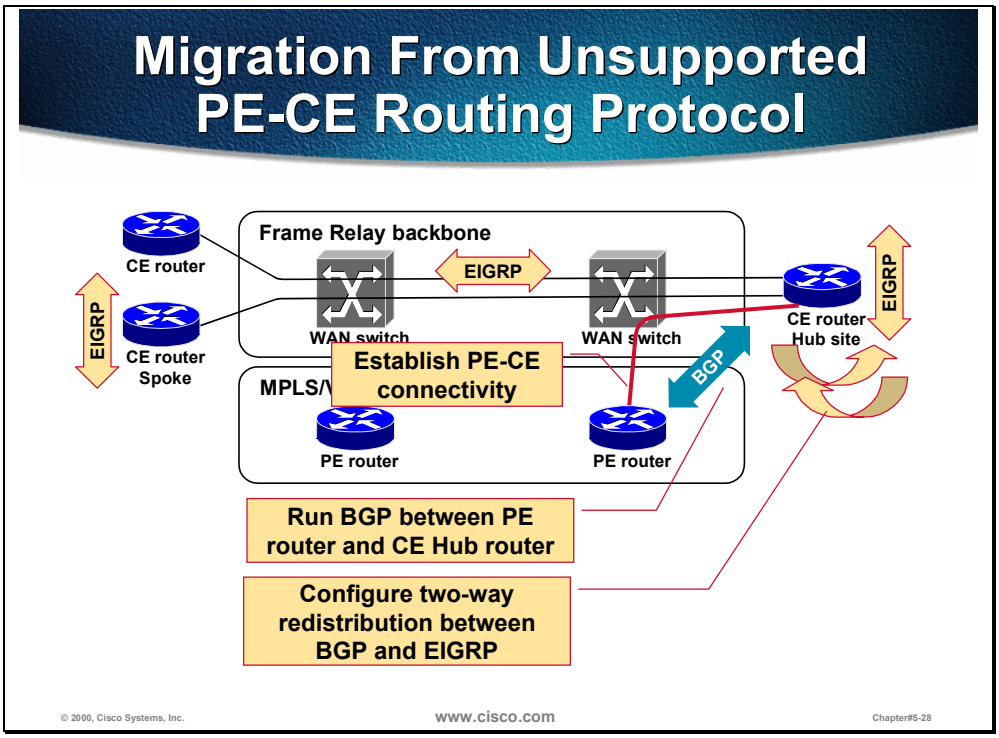

As the first migration step, BGP is established between the CE routers of the transit sites and the corresponding PE routers. Two-way redistribution between BGP and EIGRP is configured on the CE routers at transit sites. This ensures propagation of EIGRP routes (received from non-migrated sites) to MPLS VPN backbone as well as propagation of routes received from the MPLS VPN backbone via BGP to other EIGRP-speaking routers.

Note Originating default network from the CE router into the EIGRP is better than redistributing BGP routes into EIGRP. However, this approach might not be used in all networks, particularly when the customer already has a different default route, for example toward the Internet.

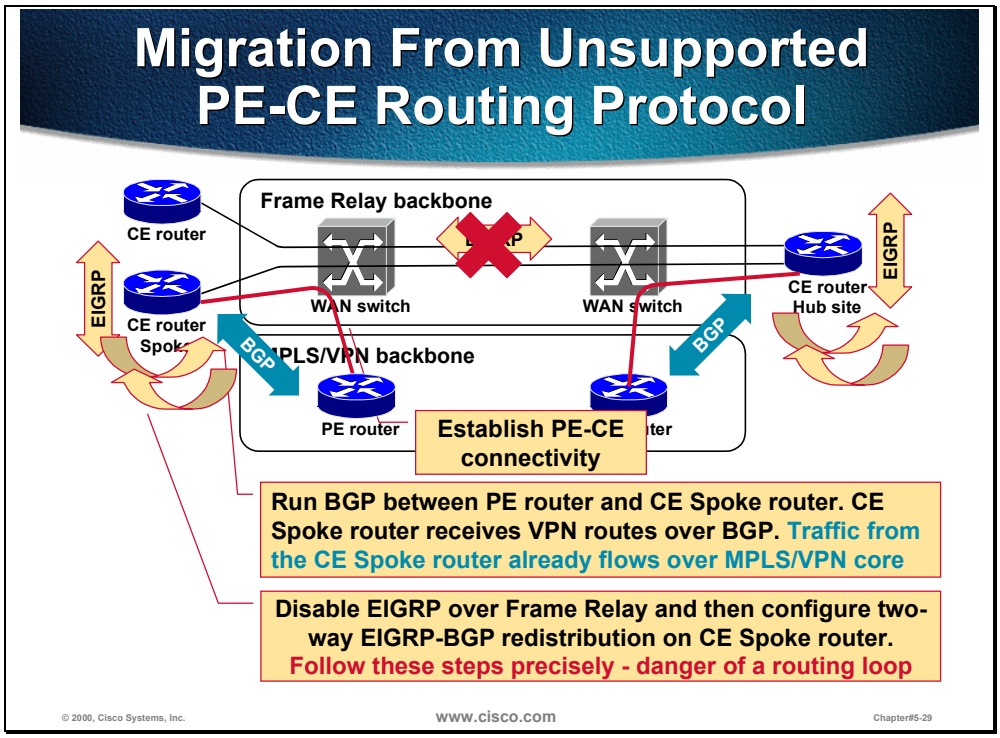

The next routing protocol migration steps are performed during migration of individual sites. In the first migration step, BGP is established as the PE-CE routing protocol. No route redistribution is configured.

At this moment, the CE-router of the migrated site receives all VPN routes as BGP routes sent by the PE-router. It also receives all VPN routes as EIGRP routes sent by the hub CE-router over the existing Frame Relay link. The EBGP routes received from the PE-router take precedence over EIGRP routes received over the old link due to lower administrative distance of EBGP routes. Traffic from the migrated site to all other sites starts flowing across the MPLS VPN backbone, resulting in asymmetrical routing.

When the exchange of BGP routes is verified, EIGRP has to be disabled on the link between the CE router under migration and the old VPN backbone. Only when the migrated CE router is isolated from the rest of the EIGRP network is it safe to configure two-way redistribution between EIGRP and BGP.

Summary

After you have deployed an MPLS VPN backbone in your network, you might want to migrate existing VPN customers to the new backbone to minimize your operational costs. The migration usually has to be performed in a gradual and non-disruptive way.

The steps necessary to migrate a customer to the new backbone vary based on the VPN technology used by the customer. However, the following principles apply to most of the migration strategies:

- Select a site (or a few sites) to act as a transit point between the old and the new VPN backbone.
- Connect the transit sites to the new VPN backbone.
- For every other site, establish connectivity with the new VPN backbone, test routing information exchange and then disable connectivity with the old backbone.

During the migration process, all the traffic between the migrated and non-migrated sites will flow over the transit sites. All migrated sites will, however, already enjoy the benefits of MPLS VPN – routing between them will be optimal.

Sometimes you have to change routing protocol the customer was using between sites during the migration to MPLS VPN as the customer might be using a routing protocol that is not supported as PE-CE routing protocol by the MPLS VPN implementation. In this case, it is strongly recommended that you use BGP as the routing protocol deployed between PE-routers and CE-routers and closely follow the steps outlined in the last example of this section to prevent routing loops.

Review Questions

- What are the steps in overlay VPN customer migration toward MPLS VPN?
- What are the necessary steps in layer-3 VPN customer migration toward MPLS VPN?
- Which protocol should you use as the PE-CE routing protocol when migrating customers are using EIGRP as their VPN routing protocol?

Chapter Summary

After completing this chapter, you should be able to design the following migration strategies for MPLS VPN deployment:

- Infrastructure migration strategy for existing IP backbones
- Phased migration strategy for pilot MPLS VPN service
- Migration strategy for customers using layer-2 overlay VPN solutions (Frame Relay or ATM)
- Migration strategy for customer running layer-3 overlay VPN solutions (GRE tunnels or IPSec)

A

Introduction to Laboratory Exercises

Overview

This chapter contains the information about your laboratory setup, details of the physical and logical connectivity in the laboratory and information on the addressing scheme, IGP routing and BGP routing pre-configured on routers.

It includes the following topics:

- Physical And Logical Connectivity
- IP Addressing Scheme
- Initial BGP Design
- Laboratory Exercises

The class will be divided into *workgroups*, each workgroup having its own Service Provider backbone and Customer routers. Each workgroup is further divided into two *subgroups.* Each subgroup will configure MPLS in part of the Service Provider backbone and implement MPLS VPN services for one of the customers.

Physical And Logical Connectivity

Routers in your workgroup are connected according to the setup in Figure 1. The light-gray routers in the figure are outside of your workgroup and are not configurable by you. They serve various functions, from injecting BGP routes into your Service Provider backbone to acting as Network Management stations.

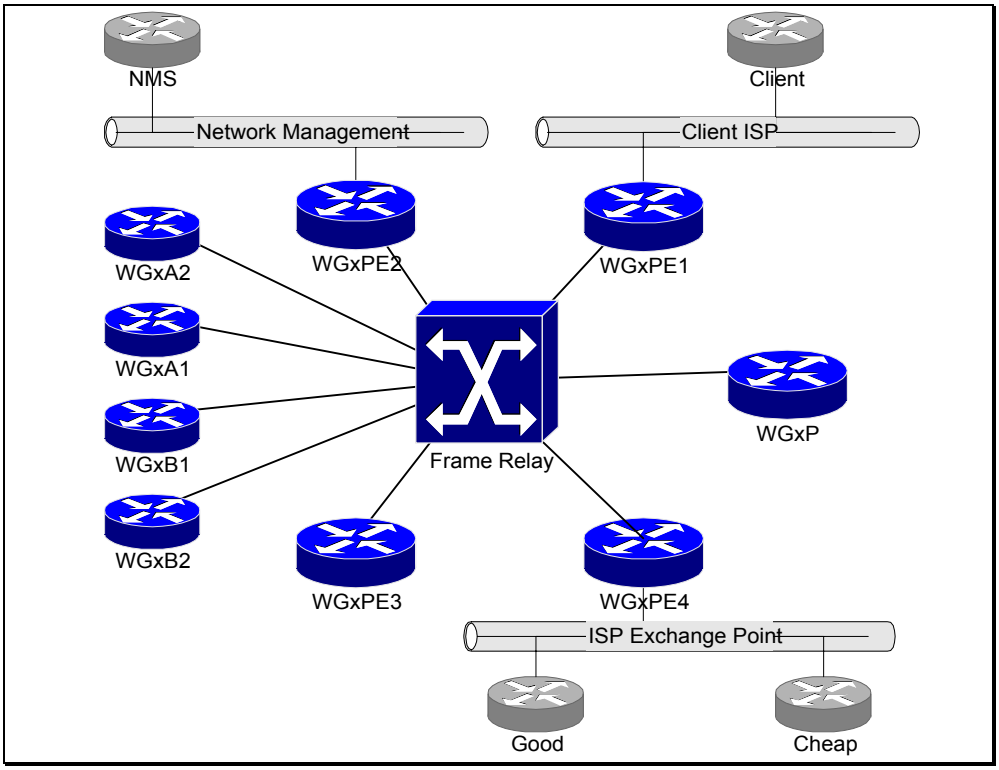

Figure 1: Physical connectivity

The first serial interface of your router is connected to the Frame Relay switch. The first Ethernet or Fast Ethernet interface of the router is connected to the LAN segment.

The routers in your workgroup have different roles as detailed in the following table:

Router name	Router role in the laboratory
WGxP	Provider router—a router in your Service Provider backbone with no customer connectivity
WGxPE1 … WGxPE4	Provider Edge routers—routers in your Service Provider backbone that connect to the Customer routers or to other Service Providers or Customers
WGxA1, WGxA2	Customer routers of customer A. The customer has two sites connected to different PE routers
WGxB1, WGxB2	Customer routers of customer B. The customer has two sites connected to different PE routers

Table 1: Roles of routers in your workgroup

The routers outside of your workgroup (the light-gray routers in Figure 1) have the following roles:

Router name	Router role in the laboratory
Good, Cheap	Upstream Service Provider routers. These two routers give your Service Provider upstream connectivity to the Internet.
Client	Your customer. This router represents an ISP customer with its own autonomous system, using your backbone for transit Internet services
NMS	This router acts as the Network Management station of your network.
WGxB1, WGxB2	Customer routers of customer B. The customer has two sites connected to different PE routers

Table 2: Roles of routers outside of your workgroup

The names of all routers in your workgroup follow this naming convention:

Router role	Router name
Provider router	WGxP, x being your workgroup number
Provider Edge router	WGxPE1 … WGxPE4.
Routers of Customer A (configured by subgroup A)	WGxA1, WGxA2
Routers of Customer B (configured by subgroup B)	WGxB1, WGxB2

Table 3: Names of routers in your workgroup

DLCIs are configured on the Frame Relay switch to give you the logical connectivity displayed in Figure 2. All the point-to-point connections in this figure are implemented with Frame Relay DLCIs, configured as point-to-point subinterfaces on the routers.

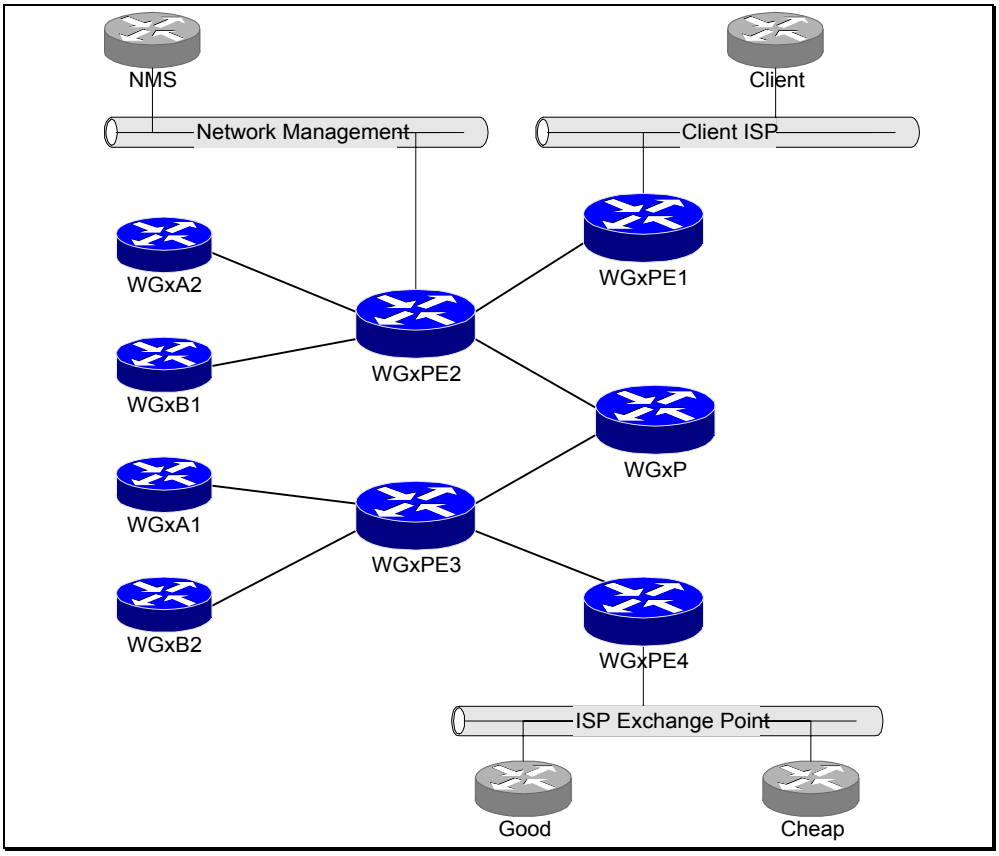

Figure 2: Initial logical connectivity of your workgroup

The DLCI values for all Frame Relay virtual circuits are shown in Table 4.

Source router	Destination router	DLCI
P	PE3	103
P	PE2	102
PE1	PE2	112
PE2	P	120
PE2	PE1	121
PE2	A2	212
PE2	B1	211
PE3	PE4	134
PE3	P	130
PE3	A1	231
PE3	B2	232
PE4	PE3	143

Table 4: Initial Core Frame Relay PVC parameters

IP Addressing Scheme

Figure 3 shows the IP addresses that have been preconfigured in the lab.

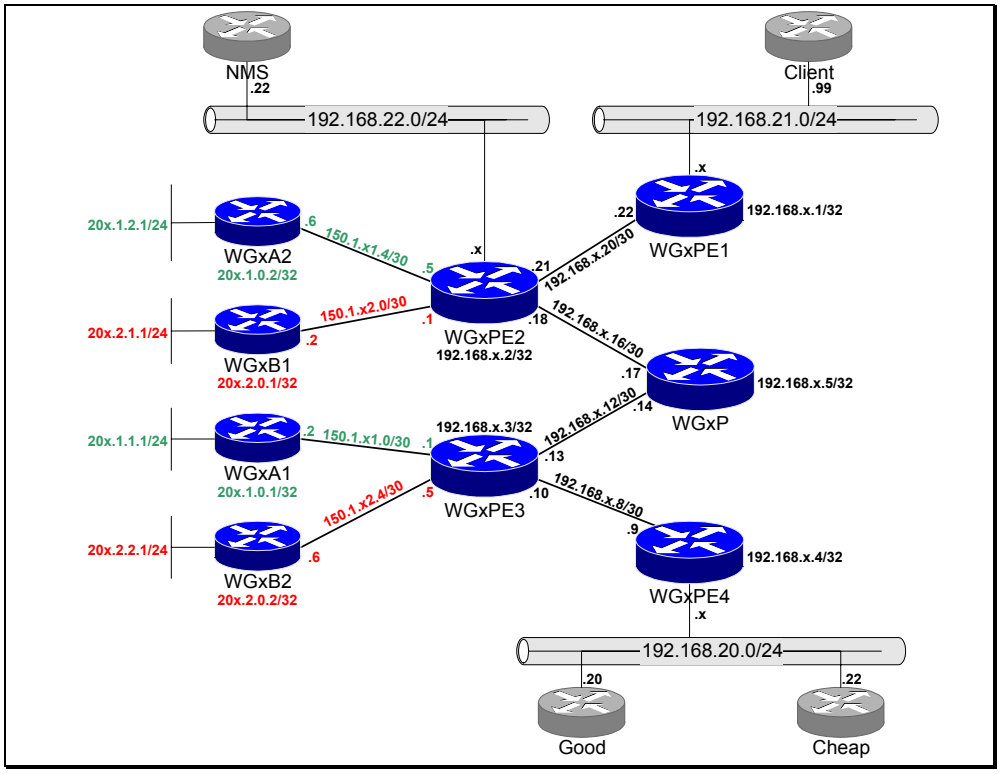

Figure 3: Addressing scheme

The addressing of Service Provider routers was performed using the following IP allocation scheme:

Parameter	Value
Router IP loopback addresses	192.168.x.1/32 … 192.168.x.5/32
Core WAN subnets	/28 subnets from 192.168.x.0/24, starting with 192.168.x.16/28.
IP address of WGxPE4 on ISP Exchange point subnet	192.168.20.x, subnet mask 255.255.255.0
IP address of WGxPE1 on Client ISP subnet	192.168.21.x, subnet mask 255.255.255.0
IP address of WGxPE2 on Network Management subnet	192.168.22.x, subnet mask 255.255.255.0

Table 5: Service Provider address space

The addressing of customer routers was performed using the following IP allocations scheme:

Customer	Address space
A (loopbacks)	20x.1.0.0/17
A (WAN links)	150.1.x1.0/25
B (loopbacks)	20x.2.0.0/17
B (WAN links)	150.1.x2.0/25

Table 6: Customer address space

Initial BGP Design

The routers have been preconfigured with IGP and BGP. The router configuration therefore includes the IS-IS and BGP configurations. You should, however, check the connectivity between customer routers inside your workgroup as well as connectivity between customer routers and external destinations before proceeding with the labs. Figure 4 shows the BGP design that has been implemented in the lab.

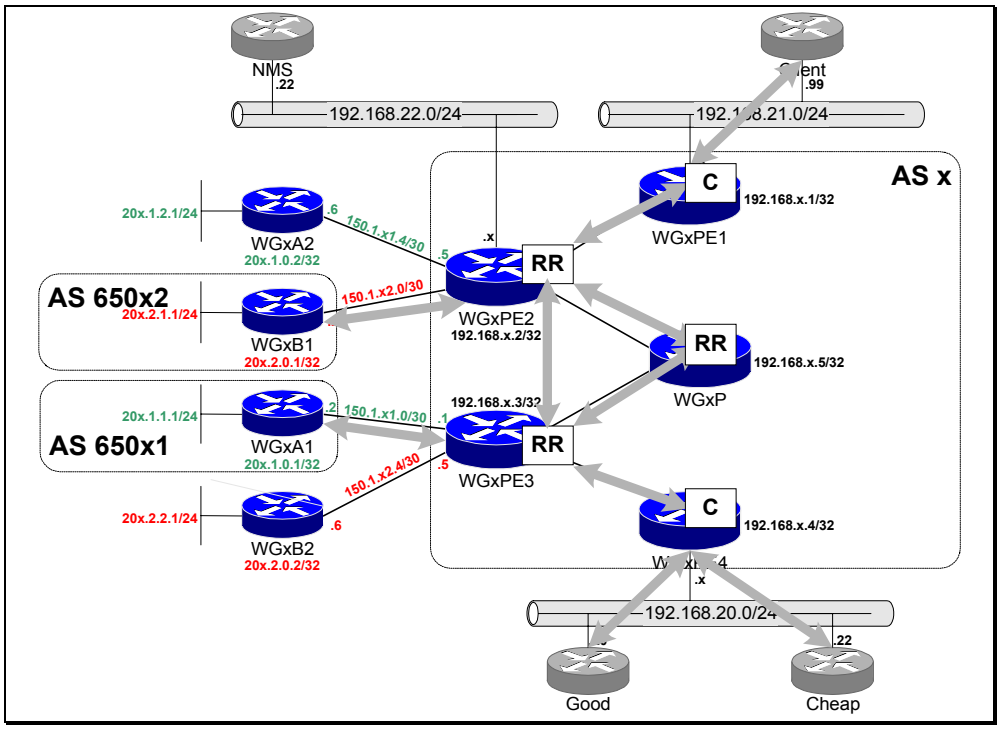

Figure 4: Initial BGP design

Notes Pages

Use the lab layout schemes in the following pages as your notepad during the lab exercises.

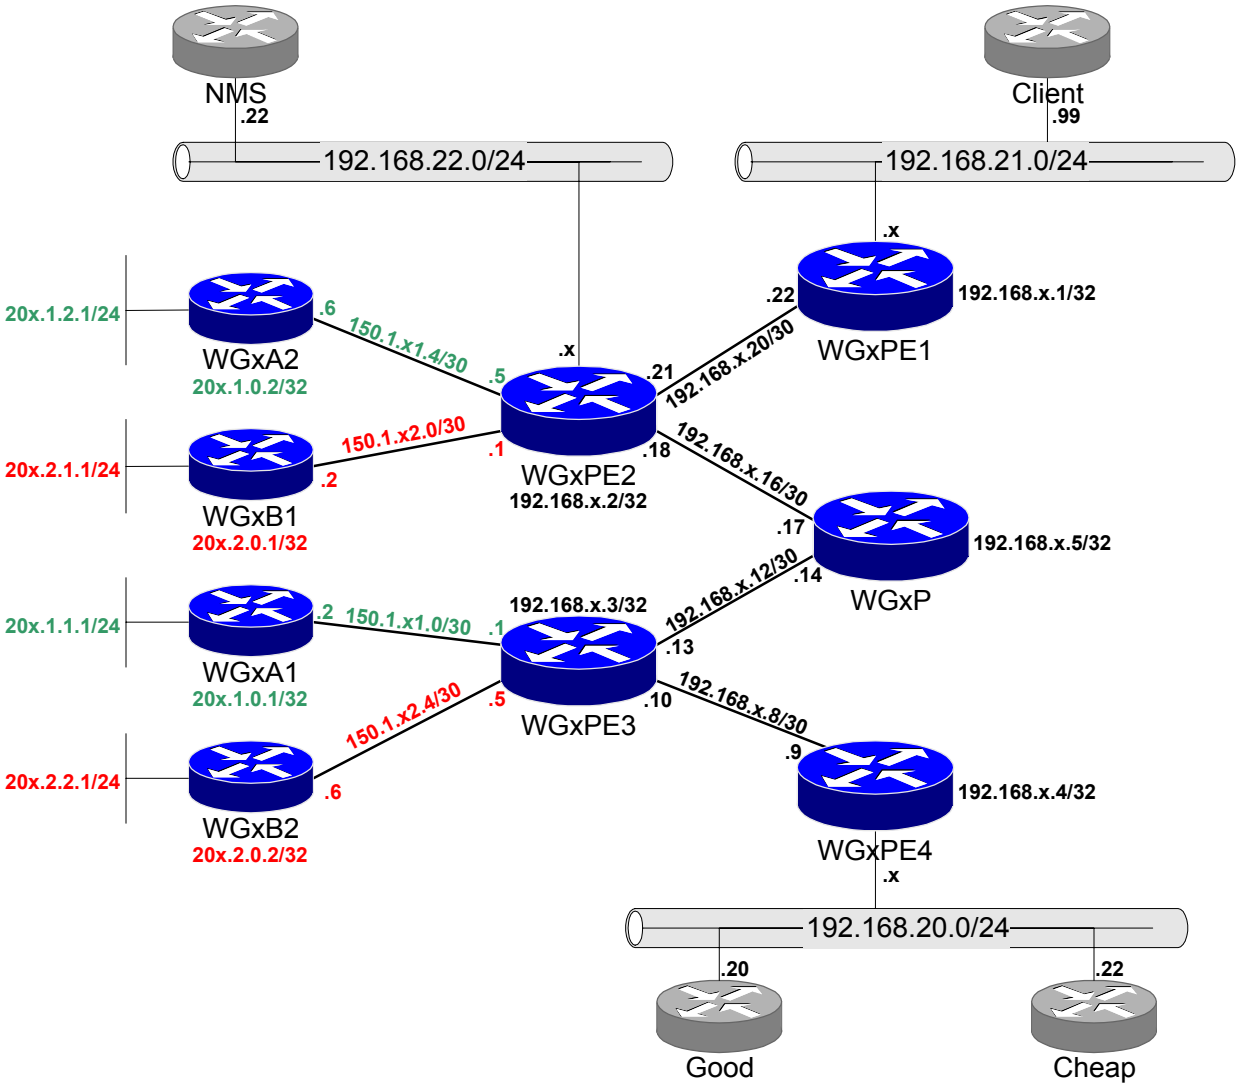

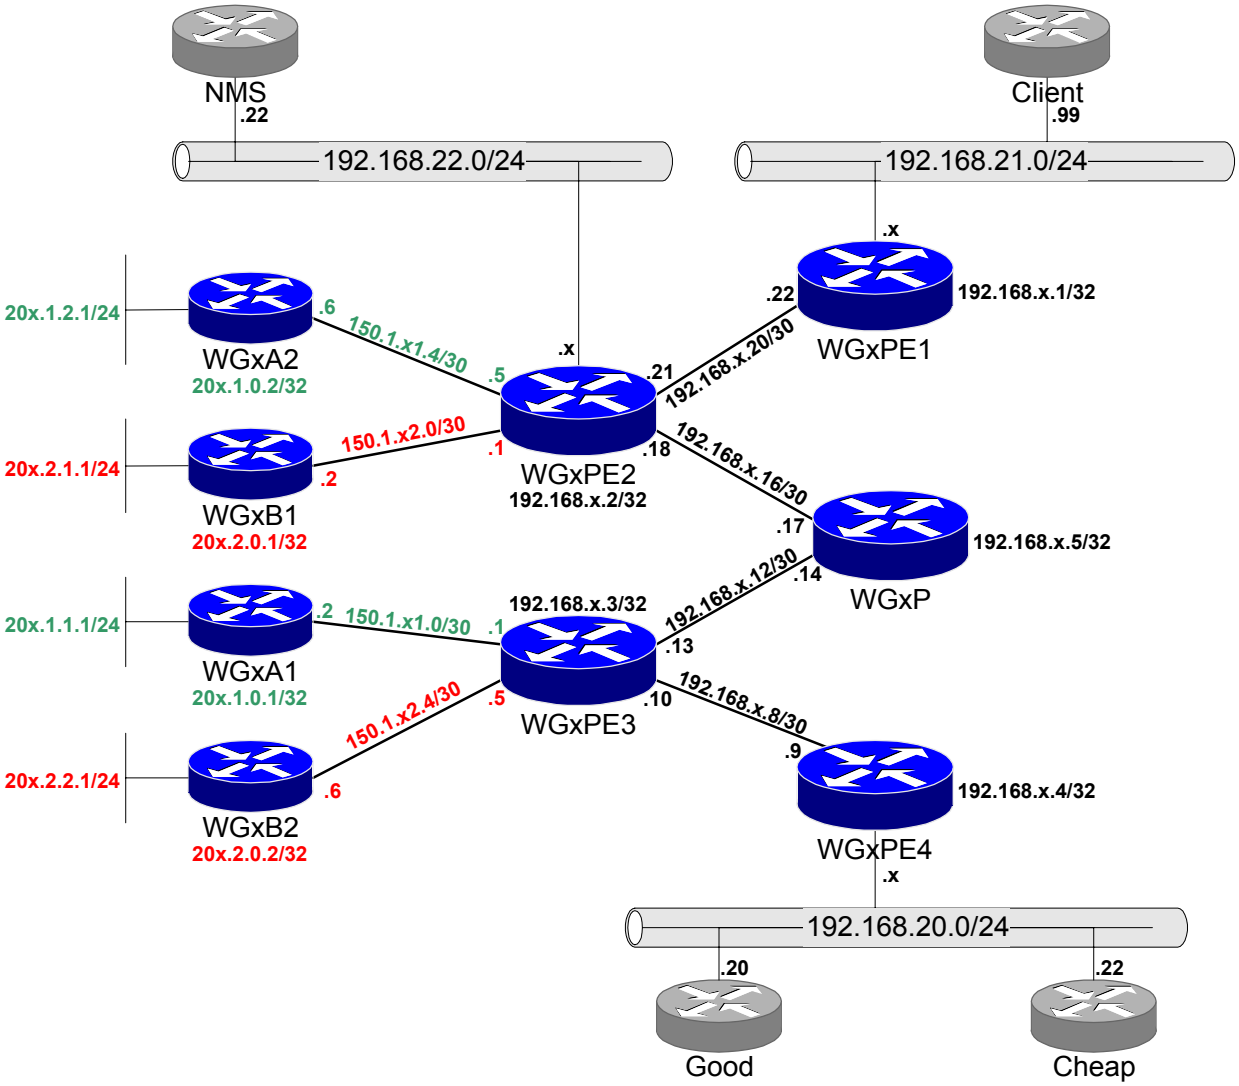
NMS
.22
Client
.99
192.168.22.0/24
192.168.21.0/24
.x
20x.1.2.1/24
.6
150.1.x1.4/30
.5
.x
.21
192.168.x.20/30
.22
192.168.x.1/32
WGxA2
20x.1.0.2/32
WGxPE1
20x.2.1.1/24
150.1.x2.0/30
.1
.2
.18
192.168.x.16/30
WGxB1
20x.2.0.1/32
WGxPE2
192.168.x.2/32
.17
192.168.x.5/32
WGxP
.14
192.168.x.12/30
20x.1.1.1/24
.2
150.1.x1.0/30
.1
192.168.x.3/32
.13
WGxA1
20x.1.0.1/32
WGxPE3
.10
192.168.x.8/30
150.1.x2.4/30
.5
.9
20x.2.2.1/24
.6
192.168.x.4/32
WGxB2
20x.2.0.2/32
WGxPE4
.x
192.168.20.0/24
.20
.22
Good
Cheap

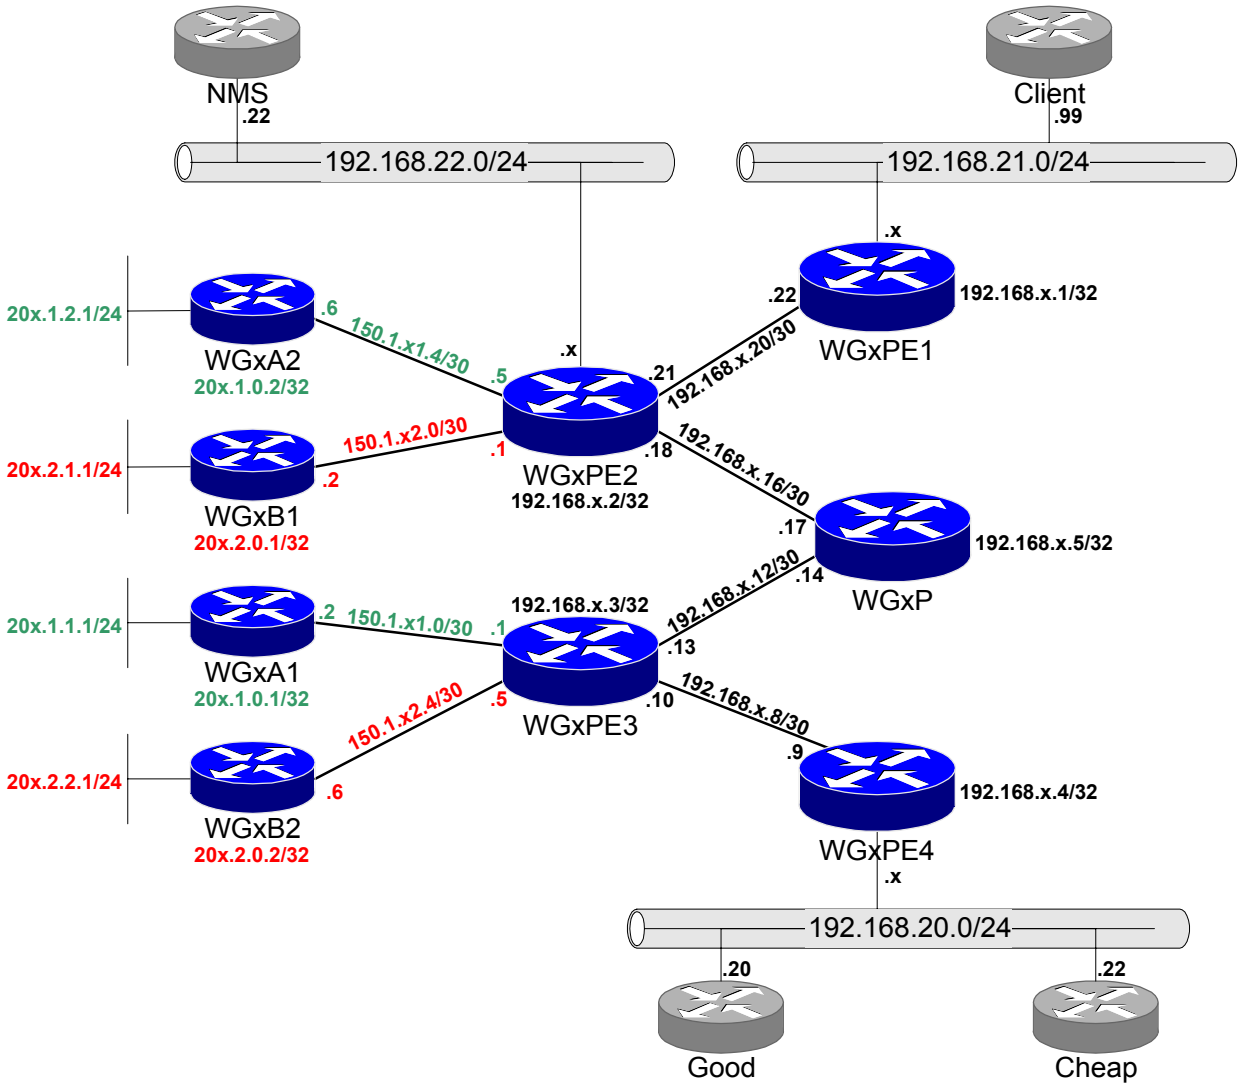

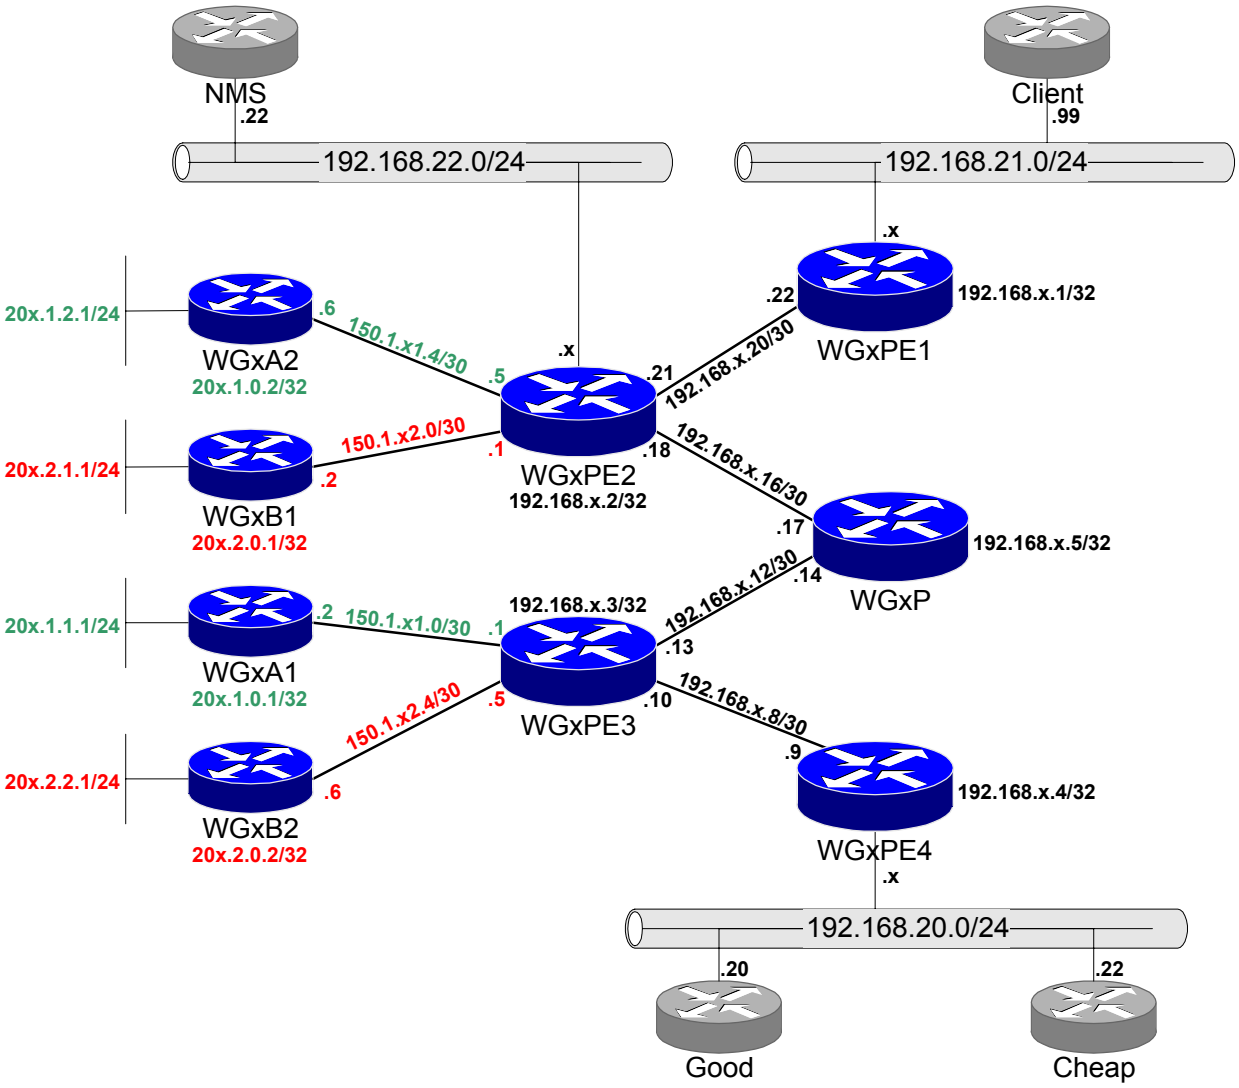
NMS
.22
Client
.99
192.168.22.0/24
192.168.21.0/24
.x
20x.1.2.1/24
.6
150.1.x1.4/30
.5
.x
.21
192.168.x.20/30
.22
192.168.x.1/32
WGxPE1
WGxA2
20x.1.0.2/32
20x.2.1.1/24
150.1.x2.0/30
.1
.2
.18
192.168.x.16/30
WGxPE2
192.168.x.2/32
WGxB1
20x.2.0.1/32
.17
192.168.x.5/32
.14
192.168.x.12/30
WGxP
20x.1.1.1/24
.2
150.1.x1.0/30
.1
192.168.x.3/32
.13
WGxA1
20x.1.0.1/32
150.1.x2.4/30
.5
.10
192.168.x.8/30
.9
20x.2.2.1/24
.6
192.168.x.4/32
WGxB2
20x.2.0.2/32
WGxPE3
WGxPE4
.x
192.168.20.0/24
.20
.22
Good
Cheap

B

Laboratory Exercises—Frame-Mode MPLS Configuration

Overview

This chapter contains exercises where you have to configure MPLS infrastructure in your core backbone. You will perform initial MPLS configuration, disable TTL propagation and configure conditional label advertising. These exercises support the core MPLS curriculum.

It includes the following exercises:

- Basic MPLS Setup
- Disabling TTL Propagation
- Conditional Label Advertising

Please read the Introduction to Laboratory Exercises chapter to become familiar with your laboratory before proceeding with these exercises.

Laboratory Exercise B-1: Basic MPLS Setup

Objectives

MPLS is enabled in Service Provider core networks to prepare the network core for MPLS VPN services or to gain additional benefits enabled by MPLS technology, for example, the ability to use traffic engineering.

In this laboratory exercise you will complete the following task:

- Configure basic label switching functionality in the ISP core network

Command list

Use the following commands to complete this exercise:

Command	Task
tag-switching ip	Enable MPLS on an interface
no router bgp	Disable BGP routing on core routers
no neighbor *ip-address*	Remove BGP neighbor after BGP is disabled on the neighbor router
show tag-switching tdp neighbors	Verify that the TDP neighbors are operational
show tag-switching tdp bindings	Verify label allocation and distribution between the neighbors

Table 7: Configuration and monitoring commands used to configure basic MPLS functionality

Task 1: Configure MPLS in your backbone

Subgroup A configures WGxPE1 and WGxPE2. Subgroup B configures WGxPE3 and WGxPE4. The WGxP router is configured by any one of them.

Step 1 Configure MPLS on all core interfaces. Do not configure MPLS on any interfaces toward customers or external backbones.

Task 2: Remove BGP from your P-routers

In traditional IP backbones, the Service Provider routers need to perform IP lookup at every hop. All P-routers therefore need full Internet routing. With the introduction of MPLS, the IP packets are labeled by the PE-routers and the P-routers no longer need full Internet routing. BGP can therefore be disabled on the P-routers.

Step 2 Remove BGP from all the core routers that do not need it any more

Verification:

After you have disabled the BGP on the core routers, perform the following tests:

- Display TDP neighbors on the core routers to verify proper TDP operation. You should get a printout similar to the one below:

```
WG1PE3#show tag-switching tdp neighbor
Peer TDP Ident: 192.168.1.5:0; Local TDP Ident 192.168.1.3:0
        TCP connection: 192.168.1.5.11003 - 192.168.1.3.711
        State: Oper; PIEs sent/rcvd: 1569/1562; ; Downstream
        Up time: 22:43:11
        TDP discovery sources:
          Serial0/0.1
        Addresses bound to peer TDP Ident:
          192.168.1.17    192.168.1.14    192.168.1.5
Peer TDP Ident: 192.168.1.4:0; Local TDP Ident 192.168.1.3:0
        TCP connection: 192.168.1.4.11006 - 192.168.1.3.711
        State: Oper; PIEs sent/rcvd: 1564/1582; ; Downstream
        Up time: 22:42:45
        TDP discovery sources:
          Serial0/0.2
        Addresses bound to peer TDP Ident:
          192.168.20.1    192.168.1.9     192.168.1.4
```

- Display TDP label bindings on your routers to verify that every IGP route has a local label and a label from all TDP neighbors. You should get a printout similar to the one below:

```
WG1PE3#show tag-switching tdp bindings
  tib entry: 192.168.1.1 255.255.255.255, rev 35
        local binding:  tag: 21
        remote binding: tsr: 192.168.1.5:0, tag: 20
        remote binding: tsr: 192.168.1.4:0, tag: 19
  tib entry: 192.168.1.2 255.255.255.255, rev 36
        local binding:  tag: 22
        remote binding: tsr: 192.168.1.5:0, tag: 22
        remote binding: tsr: 192.168.1.4:0, tag: 21
  tib entry: 192.168.1.3 255.255.255.255, rev 37
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.1.5:0, tag: 21
        remote binding: tsr: 192.168.1.4:0, tag: 20
```

- Perform **trace** from WGxA2 or WGxB1 toward 192.168.20.20. You should see all your core routers in the path. A sample **trace** printout is shown below:

```
WG1A2#trace 192.168.20.20

Type escape sequence to abort.
Tracing the route to 192.168.20.20

  1 150.1.11.5 44 msec 36 msec 32 msec                PE2 router
  2 192.168.1.17 164 msec 176 msec 168 msec           P router
  3 192.168.1.13 148 msec 156 msec 152 msec           PE3 router
  4 192.168.1.9 68 msec 76 msec 72 msec               PE4 router
  5 192.168.20.20 72 msec *  72 msec                  final destination
WG1A2#
```

- Perform **trace** from WGxA2 or WGxB1 toward the WGxP router. The **trace** should display WGxPE2 but fail at the WGxP router, similar to the printout below:

```
WG1A2#trace p

Type escape sequence to abort.
Tracing the route to P (192.168.1.5)

  1 150.1.11.5 44 msec 32 msec 32 msec                PE2 router
  2  *  *  *
  3  *  *  *
```

- Perform **trace** from WGxA1 or WGxB2 toward 192.168.21.99. Again, you should see all your core routers in the path

Review Questions

- Why can you trace across the WGxP router but not to the WGxP router?
- How does the WGxP router return the ICMP unreachable packet to the sender if it does not have the sender's IP address in its routing table?
- If you investigate the LIB on WGxPE2, you will discover that the TDP neighbors (WGxPE1 and WGxP) do not advertise labels for subnets toward WGxA2 and WGxB1. Why?

Laboratory Exercise B-2: Disabling TTL Propagation

Objective

In this laboratory exercise, you will complete the following task:

- Disable IP TTL propagation into MPLS labels to hide your core routers from the customers of your network

Command list

Use the following commands to complete this exercise:

Command	Task
no tag-switching ip propagate-ttl	Disable TTL propagation from IP packets to MPLS label header and vice versa

Table 8: Configuration commands used to disable TTL propagation

Task: Disable IP TTL Propagation

Step 1 Disable IP TTL propagation on all Service Provider routers that perform labeling of incoming IP packets.

Verification

- Perform **trace** from WGxA2 or WGxB1 toward 192.168.20.20. You should see only the ingress and egress core router in the path. A sample **trace** printout is shown below:

```
WG1A2#trace 192.168.20.20

Type escape sequence to abort.
Tracing the route to 192.168.20.20

  1 150.1.11.5 44 msec 36 msec 32 msec                PE2 router
  2 192.168.1.9 68 msec 76 msec 72 msec               PE4 router
  3 192.168.20.20 72 msec *  72 msec                  final destination
WG1A2#
```

- Perform **trace** from WGxA1 or WGxB2 toward 192.168.21.99. Again, you should only see PE3 and PE1 router in the **trace** printout.

Laboratory Exercise B-3: Conditional Label Advertising

Objective

In this laboratory exercise, you will complete the following task:

- Use the **conditional label advertising** feature of TDP to configure label switching for all addresses, except the WAN subnets in your core

Command list

Use the following commands to complete this exercise:

Command	Task
access-list …	Specify the access list that will match IP prefixes for which the labels will be advertised
tag-switching advertise for *acl*	Configure conditional label advertising for IP prefixes matched by the specified access list
show tag-switching tdp bindings	Verify label allocation and distribution between the neighbors

Table 9: Configuration and monitoring commands used to configure conditional label advertising

Task: Configure Conditional Label Advertising

Step 1 On all routers in your Service Provider backbone, configure the conditional label advertising. Your routers should not advertise labels for WAN subnets in your Service Provider core.

Verification

- Perform **trace** from WGxA2 or WGxB1 toward 192.168.20.20. You should only see the ingress and egress core router in the path (the rest of your core is invisible to customer trace). A sample **trace** printout is shown below:

```
WG1A2#trace 192.168.20.20

Type escape sequence to abort.
Tracing the route to 192.168.20.20

  1 150.1.11.5 44 msec 36 msec 32 msec                PE2 router
  2 192.168.1.9 68 msec 76 msec 72 msec               PE4 router
  3 192.168.20.20 72 msec *  72 msec                  final destination
WG1A2#
```


- Perform **trace** from WGxA2 or WGxB1 toward an IP interface assigned to a core WAN link on PE4. You should see all your core routers apart from WGxP in the trace, with a printout similar to the one below:

```
WG1A2#trace 192.168.1.9

Type escape sequence to abort.
Tracing the route to 192.168.1.9

  1 150.1.11.5 44 msec 32 msec 28 msec          PE2 router
  2  *  *  *                                    no response from P router
  3 192.168.1.13 236 msec 60 msec 56 msec       PE3 router
  4 192.168.1.9 72 msec *  72 msec              PE4 router – final destination
```

Review Questions

- Why does the WGxP router respond to trace toward external destination, but not to trace toward a WAN subnet?

C

Laboratory Exercises—MPLS VPN Implementation

Overview

This chapter contains exercises to enable you to configure your core MPLS VPN infrastructure and establish simple any-to-any VPN service for a customer. You will also test various PE-CE routing options, ranging from RIP and OSPF to running BGP between the PE and the CE router.

It includes the following exercises:

- Initial MPLS VPN Setup
- Running OSPF Between PE and CE Routers
- Running BGP Between the PE and CE Routers

These exercises rely on the **Frame Mode MPLS Configuration** exercises where you established MPLS connectivity in your backbone. If this is the first set of exercises you are performing, please refer to the Introduction to Laboratory Exercises chapter to familiarize yourself with the IP addressing and routing in your workgroup. Please also verify that MPLS has been enabled on all core interfaces in your backbone and that it has not been enabled on interfaces toward customer routers or other Service Providers.

Laboratory Exercise C-1: Initial MPLS VPN Setup

Objectives

In this laboratory exercise you will create a simple Virtual Private Network for your customer. To achieve this objective you will complete the following tasks:

- Establish MPLS VPN infrastructure in your backbone
- Migrate your customer from global routing to a simple VPN
- Change the routing protocol between the PE and CE routers to RIP

Background Information

The following diagram displays the parts of your MPLS VPN network that you will configure in this exercise:

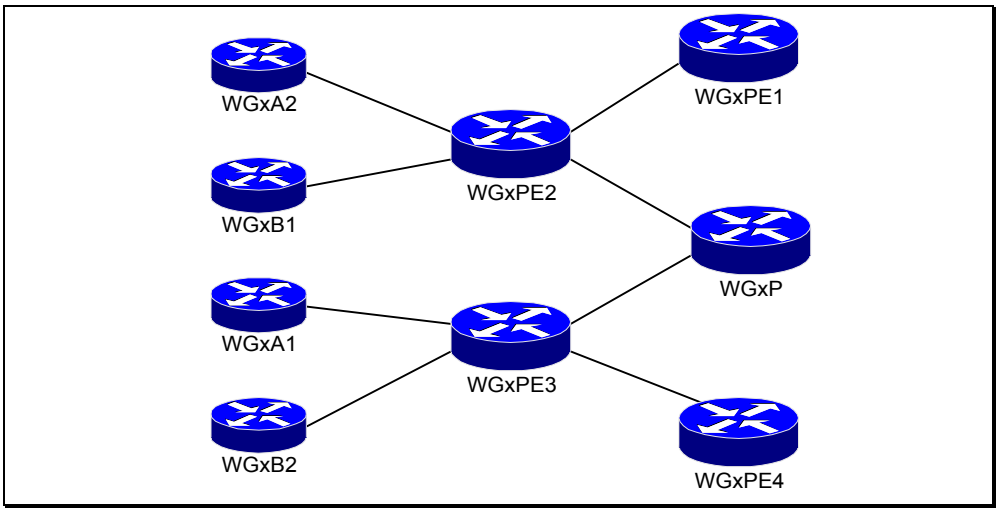

Figure 5: Parts of your network configured during this exercise

Command list

Use the following commands to complete this exercise:

Command	Task
router bgp *as-number*	Select BGP configuration
address-family vpnv4	Select VPNv4 address family configuration
neighbor *ip-address* activate	Activate exchange of routes from address family under configuration for specified neighbor
neighbor *ip-address* route-reflector-client	Configure a route-reflector client on a route-reflector
ip vrf *name*	Creates a virtual routing and forwarding table
rd *value*	Assigns a route-distinguisher to a VRF
route-target import|export *value*	Assigns a route target to a VRF
address-family ipv4 vrf *name*	Selects per-VRF instance of a routing protocol
ip vrf forwarding *name*	Assigns an interface to a VRF
redistribute bgp *as-number* metric transparent	Redistribute BGP routes into RIP with propagation of MED into RIP hop-count.
show ip vrf detail	Displays detailed VRF information
show ip bgp neighbor	Displays information on global BGP neighbors
show ip bgp vpnv4 vrf *name*	Displays VPNv4 routes associated with the specified VRF
show ip route vrf *name*	Displays IP routing table of specified VRF
telnet host /vrf *name*	Telnets to a CE router connected to the specified VRF
ping vrf *name host*	Pings a host reachable through the specified VRF

Table 10: Configuration and monitoring commands used to configure simple VPN with RIP routing

Task 1: Configure multi-protocol BGP

In this section of the exercise, you will configure multi-protocol BGP between PE routers. Subgroup A will configure multi-protocol BGP on WGxPE1 and WGxPE2; subgroup B will perform the same task on WGxPE3 and WGxPE4.

Complete the following steps:

Step 1 Activate VPNv4 BGP sessions between all PE routers in your Service Provider backbone.

Step 2 On the PE routers acting as route reflectors, configure the route-reflector clients under the VPNv4 address family.

Task 2: Configure Virtual Routing and Forwarding Tables

In this section and the following sections, you will establish simple Virtual Private Networks for Customer A and Customer B. Subgroup A will establish a VPN between the WGxA1 and WGxA2, subgroup B will establish a VPN between WGxB1 and WGxB2. Each workgroup is responsible for all PE-router configurations related to their customer. This division of work between workgroups applies to all further exercises.

Step 1 Design your VPN networks—decide on the route distinguisher and the route target numbering. Coordinate your number with the other subgroup.

Note The easiest numbering plan would use the same values for the route distinguisher and the route target. Use simple values, for example x:10 for customer A and x:20 for customer B.

Step 2 Create VRFs on the PE routers and migrate the PE-CE interfaces into the proper VRFs; use simple yet descriptive VRF names (i.e. **wg*x*a** and **wg*x*b**).

Step 3 Configure RIP in the VRF you have created.

Step 4 Configure redistribution of RIP into BGP within the **ipv4 vrf** address family.

Step 5 Configure redistribution of BGP into RIP within the **ipv4 vrf** address family. Configure RIP metric propagation through Multiprotocol BGP by using the **redistribute bgp metric transparent** command in the RIP process.

Step 6 Configure RIP on all the CE routers. Make sure you list all the networks (including loopbacks) in the RIP process.

Note Do not remove BGP from WGxA1 and WGxB1 as you will need it for later exercises. Disable BGP by using the **neighbor shutdown** command.

Additional Objective

There are two additional customer routers (WGxA3 and WGxB3) connected to your network as displayed in Figure 6. They are using RIP to propagate their networks. In this laboratory exercise you will complete the following task:

- Connect the router WGxA3 into the VPN of Customer A and WGxB3 into the VPN of Customer B.

Note WGxA3 and WGxB3 are controlled by your customers and are not configurable or accessible to you through a Telnet session.

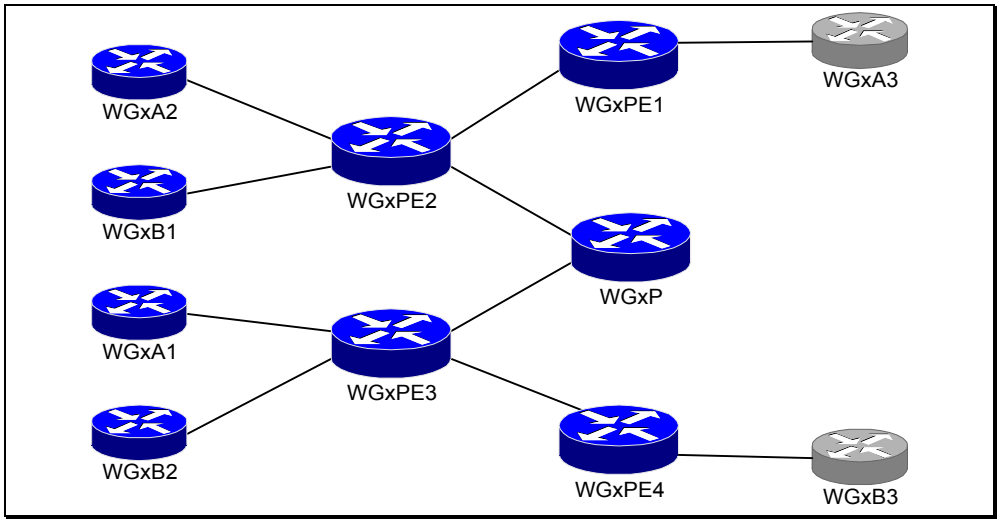

Figure 6: Additional customer routers

Task 3: Configuring Additional CE routers

Step 1 Configure your PE routers based on the parameters in Table 11.

CE router	Connected to PE router	DLCI on the PE router	WAN IP address on the PE router
WGxA3	WGxPE1	413	150.1.x1.129/30
WGxB3	WGxPE4	543	150.1.x2.129/30

Table 11: Connectivity parameters for customer RIP routers

Step 2 Configure customer VRF on WGxPE1/ WGxPE4.

Step 3 Configure RIP on WGxPE1/WGxPE4 to establish RIP routing between the PE routers and WGxA3/WGxB3. RIP is already configured on the CE routers.

Step 4 Follow the steps outlined in the **Configure Virtual Routing and Forwarding Table** task to complete the VPN routing configuration on WGxPE1/WGxPE4.

Verification

- Verify that you have the proper configuration of your Virtual Routing and Forwarding tables with **show ip vrf detail**. You should get a printout similar to the one below:

```
WG2PE4#sh ip vrf detail
VRF wg2b; default RD 2:20
  Interfaces:
    Serial1/0.4
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:2:20
  Import VPN route-target communities
    RT:2:20
  No import route-map
  No export route-map
```

- Check the routing protocols running in your VRF with the **show ip protocol vrf** command. When executed on WG2PE2 it will produce a printout similar to the one below:

```
WG2PE2#show ip proto vrf wg2b
Routing Protocol is "bgp 2"
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Route Reflector for address family IPv4 Unicast, 1 clients
  Route Reflector for address family VPNv4 Unicast, 1 clients
  Route Reflector for address family IPv4 Multicast, 1 clients
  IGP synchronization is disabled
  Automatic route summarization is disabled
  Redistributing: rip
  Routing for Networks:
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.2.3          200      01:29:55
  Distance: external 20 internal 200 local 200

Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 14 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Redistributing: bgp 2, rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Serial1/0.4           2     2
  Routing for Networks:
    150.1.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    150.1.22.2           120      00:00:00
  Distance: (default is 120)
```


- Verify the per-VRF routing table on the PE router with the **show ip route vrf** command. It will produce a printout similar to the one below:

```
WG2PE4#show ip route vrf wg2b
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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

B    202.2.2.0 255.255.255.0 [200/1] via 192.168.2.3, 01:28:39
B    202.2.1.0 255.255.255.0 [200/1] via 192.168.2.2, 01:28:29
R    202.2.134.0 255.255.255.0 [120/1] via 150.1.22.130, 00:00:16, Serial1/0.4
     202.2.127.0 255.255.255.255 is subnetted, 1 subnets
R       202.2.127.3 [120/1] via 150.1.22.130, 00:00:16, Serial1/0.4
     150.1.0.0 255.255.255.252 is subnetted, 3 subnets
C       150.1.22.128 is directly connected, Serial1/0.4
B       150.1.22.0 [200/0] via 192.168.2.2, 01:28:29
B       150.1.22.4 [200/0] via 192.168.2.3, 01:28:41
```

- Use the **show ip bgp vpnv4 vrf** command to display the BGP routing table associated with a VRF. The printout from WG2PE4 router is shown below:

```
WG2PE4#show ip bgp vpnv4 vrf wg2b
BGP table version is 24, local router ID is 192.168.2.4
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: 2:20 (default for vrf wg2b)
*>i150.1.22.0/30    192.168.2.2              0    100      0 ?
*>i150.1.22.4/30    192.168.2.3              0    100      0 ?
*> 150.1.22.128/30  0.0.0.0                  0         32768 ?
*>i202.2.1.0        192.168.2.2              1    100      0 ?
*>i202.2.2.0        192.168.2.3              1    100      0 ?
*> 202.2.127.3/32   150.1.22.130             1         32768 ?
*> 202.2.134.0      150.1.22.130             1         32768 ?
```

- On a CE router, use the **show ip route** command to verify that the router is receiving all VPN routes. On WG2B1, the printout is similar to the one below:

```
WG2B1#show ip rout
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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

R    202.2.2.0 255.255.255.0 [120/2] via 150.1.22.1, 00:00:07, Serial1/0.1
     202.2.0.0 255.255.255.255 is subnetted, 1 subnets
C       202.2.0.1 is directly connected, Loopback0
```

```
C    202.2.1.0 255.255.255.0 is directly connected, Loopback1
R    202.2.134.0 255.255.255.0 [120/2] via 150.1.22.1, 00:00:07, Serial1/0.1
     202.2.127.0 255.255.255.255 is subnetted, 1 subnets
R       202.2.127.3 [120/2] via 150.1.22.1, 00:00:07, Serial1/0.1
     150.1.0.0 255.255.255.252 is subnetted, 3 subnets
R       150.1.22.128 [120/1] via 150.1.22.1, 00:00:07, Serial1/0.1
C       150.1.22.0 is directly connected, Serial1/0.1
R       150.1.22.4 [120/1] via 150.1.22.1, 00:00:07, Serial1/0.1
```

- Use **ping** and **trace** on the CE routers to verify connectivity across the VPN.
- Use the **show ip route** command on the PE routers to verify that the customer routes are no longer in the global IP routing table.
- Use **ping** and **trace** on the PE routers to verify that you cannot reach your customer networks from global address space.

Laboratory Exercise C-2: Running OSPF Between PE and CE Routers

Objectives

Some customers insist on using OSPF as the routing protocol in their VPN, sometimes even combined with RIP or BGP at other sites.

In this laboratory exercise, you will deploy OSPF as the PE-CE routing protocol in your customer's VPN by completing the following tasks:

- Activate the OSPF as the routing protocol between PE and CE routers
- Transport OSPF routes between customer sites
- Configure connectivity with additional CE routers running OSPF

Visual Objective

Subgroup A configures OSPF between WGxA1 and WGxPE3 and between WGxA4 and WGxPE1. Subgroup B configures OSPF between WGxB1 and WGxPE2 and between WGxB4 and WGxPE4. WGxA4 and WGxB4 are additional customer routers, similar to the ones running RIP in the previous exercise.

Note WGxA4 and WGxB4 are controlled by your customers and are not configurable or accessible to you through a Telnet session.

The figure shows the part of your MPLS VPN network under configuration.

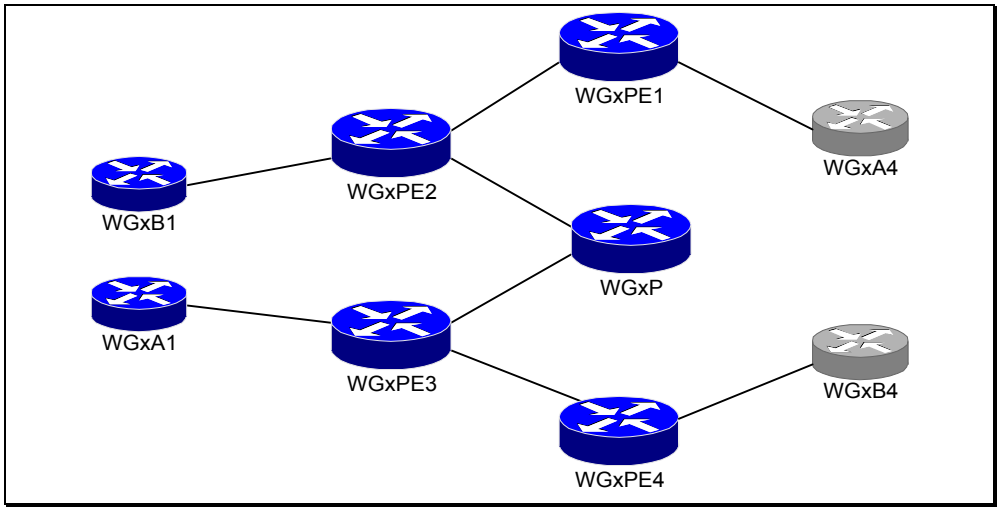

Figure 7: Configuring OSPF between PE-routers and CE-routers

Command list

Use the following commands to complete this exercise:

Command	Task
router bgp *as-number*	Select BGP configuration
ip vrf *name*	Creates a virtual routing and forwarding table
rd *value*	Assigns a route-distinguisher to a VRF
route-target import|export *value*	Assigns a route target to a VRF
address-family ipv4 vrf *name*	Selects per-VRF instance of a routing protocol
ip vrf forwarding *name*	Assigns an interface to a VRF
router ospf *process* vrf *name*	Starts an OSPF process within the specified VRF
redistribute bgp *as-number* subnets	Redistribute BGP routes (including subnet routes) into OSPF
default-information originate always	Generates a default route into OSPF
show ip vrf detail	Displays detailed VRF information
show ip ospf database	Displays OSPF database information
show ip bgp vpnv4 vrf *name*	Displays VPNv4 routes associated with the specified VRF
show ip route vrf *name*	Displays IP routing table of specified VRF
telnet host /vrf *name*	Telnets to a CE router connected to the specified VRF
ping vrf *name host*	Pings a host reachable through the specified VRF

Table 12: Configuration and monitoring commands used to configure simple VPN with OSPF routing

Task 1: Configure OSPF on CE routers

Step 1 Disable RIP and configure OSPF on WGxA1 and WGxB1 using the **router ospf** command. Configure OSPF areas in the CE router according to the table below:

Area	Interface(s)
Area 0	WAN interface toward PE-router Loopback 0
Area 1	Loopback 1

Table 13: OSPF areas configured in the CE-routers

Task 2: Configure OSPF on PE routers

Step 1 Configure OSPF in the VRFs on WGxPE3 and WGxPE2 using the **router ospf vrf** command. Use OSPF area 0 on the PE-CE link.

Step 2 Configure redistribution from OSPF to multi-protocol BGP using the **redistribute ospf** command inside the VRF address family configuration.

Step 3 Configure redistribution from multi-protocol BGP to OSPF using the **redistribute bgp subnets** command in the OSPF router configuration.

Verification

- Verify the OSPF adjacency on WGxA1 and WGxB1 or on the PE routers using the **show ip ospf neighbor** command.
- Check the OSPF topology database on WGxA1 and WGxB1. You should see router link states (resulting from OSPF connectivity between the PE and the CE router) and type-5 external link states (all other VPN routes originated in RIP or BGP), but no summaries. A sample printout from WG2A1 is shown:

```
WG2A1#show ip ospf data

            OSPF Router with ID (202.1.1.1) (Process ID 65021)
                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
150.1.21.1      150.1.21.1      1324        0x80000009 0x98E0   2
202.1.1.1       202.1.1.1       1721        0x8000000A 0xB684   4

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
150.1.21.4      150.1.21.1      1324        0x80000002 0x66F0   0
150.1.21.128    150.1.21.1      1324        0x80000002 0x8951   0
202.1.0.0       150.1.21.1      1324        0x80000002 0xE157   0
202.1.2.0       150.1.21.1      1324        0x80000002 0xCB6B   0
202.1.127.3     150.1.21.1      1324        0x80000002 0x496D   0
202.1.134.0     150.1.21.1      1324        0x80000002 0x1A98   0
```

- Verify connectivity across VPN by using **ping** and **trace** commands on the CE routers and **ping vrf** and **trace vrf** commands on the PE routers.

Task 3: Configure OSPF connectivity with additional CE routers

Step 1 Configure connectivity between the PE routers and WGxA4/WGxB4 using parameters in Table 14.

CE router	Connected to PE router	DLCI on the PE router	WAN IP address on the PE router
WGxA4	WGxPE1	414	150.1.x1.133/30
WGxB4	WGxPE4	544	150.1.x2.133/30

Table 14: Connectivity parameters for customer OSPF routers

Step 2 Configure OSPF between PE routers and WGxA4/WGxB4 on the PE router using the **router ospf vrf** command. Use OSPF area 0 on the PE-CE link. OSPF is already configured on WGxA4/WGxB4.

Step 3 Do not redistribute BGP routes into this instance of OSPF. Use the **default-information originate always** configuration command to insert the OSPF default route.

Step 4 Redistribute OSPF routes from WGxA4/WGxB4 in multi-protocol BGP using the **redistribute ospf** command in the VRF address family configuration.

Verification

- Verify you have the proper routing protocol configuration by using the **show ip protocol vrf** command on WGxPE1/WGxPE4.
- Verify OSPF connectivity between the CE routers and WGxPE1/WGxPE4 with the **show ip ospf neighbor** command
- Examine the OSPF topology database on WGxA1/WGxB1. The OSPF topology database should contain summary net link state objects—the OSPF routes received from the WGxA4/WGxB4 routers. The printout should be similar to the one produced on the WG2A1 router:

```
WG2A1#show ip ospf data

            OSPF Router with ID (202.1.1.1) (Process ID 65021)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
150.1.21.1      150.1.21.1      1324        0x80000009 0x98E0   2
202.1.1.1       202.1.1.1       1721        0x8000000A 0xB684   4

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
150.1.21.132    150.1.21.1      1324        0x80000002 0x449A
202.1.127.4     150.1.21.1      1324        0x80000002 0xBFED
202.1.135.0     150.1.21.1      1324        0x80000002 0x8F1A

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
150.1.21.4      150.1.21.1      1324        0x80000002 0x66F0   0
150.1.21.128    150.1.21.1      1324        0x80000002 0x8951   0
202.1.0.0       150.1.21.1      1324        0x80000002 0xE157   0
202.1.2.0       150.1.21.1      1324        0x80000002 0xCB6B   0
202.1.127.3     150.1.21.1      1324        0x80000002 0x496D   0
202.1.134.0     150.1.21.1      1324        0x80000002 0x1A98   0
```

- Examine the BGP routes generated from the redistributed OSPF routes on the PE routers. You should see additional OSPF-related route targets:

```
WG2PE1#show ip bgp vpnv4 vrf a 202.1.135.0
BGP routing table entry for 2:10:202.1.135.0/24, version 50
Paths: (1 available, best #1, table a)
  Advertised to non peer-group peers:
  192.168.2.2
  Local
    150.1.21.134 from 0.0.0.0 (192.168.2.1)
      Origin incomplete, metric 782, localpref 100, weight 32768, valid, sourced
, best
      Extended Community: RT:2:10 OSPF RT:0:2:0
```

- Verify proper VPN operation by performing **trace** between WG*x*A1, WG*x*A2 and WG*x*A4 (or WG*x*B1, WG*x*B2 and WG*x*B4).

Laboratory Exercise C-3: Running BGP Between the PE and CE Routers

Objectives

In this laboratory exercise, you will establish a backup link between a PE-router and the central router of your customer (WGxA1 or WGxB1). You will use BGP as the PE-CE routing protocol between the WGxA1/WGxB1 and the WGxPE2/WGxPE3. To reach this objective, you will complete the following tasks:

- Configure a backup link between WGxA1/WGxB1 and the MPLS VPN backbone, converting the A1 and B1 sites into multi-homed sites
- Configure BGP as the routing protocol between WGxA1/WGxB1 and the PE routers
- Using local preference and MED on WGxA1/WGxB1, select the primary and backup links
- Verify the switchover from primary to backup link following the primary link failure and the reactivation of primary link after the physical connectivity is reestablished

Background Information

An additional Frame Relay DLCI between WGxA1 – WGxPE2 and WGxB1 – WGxPE3 is configured on the Frame Relay switch as shown in Figure 8.

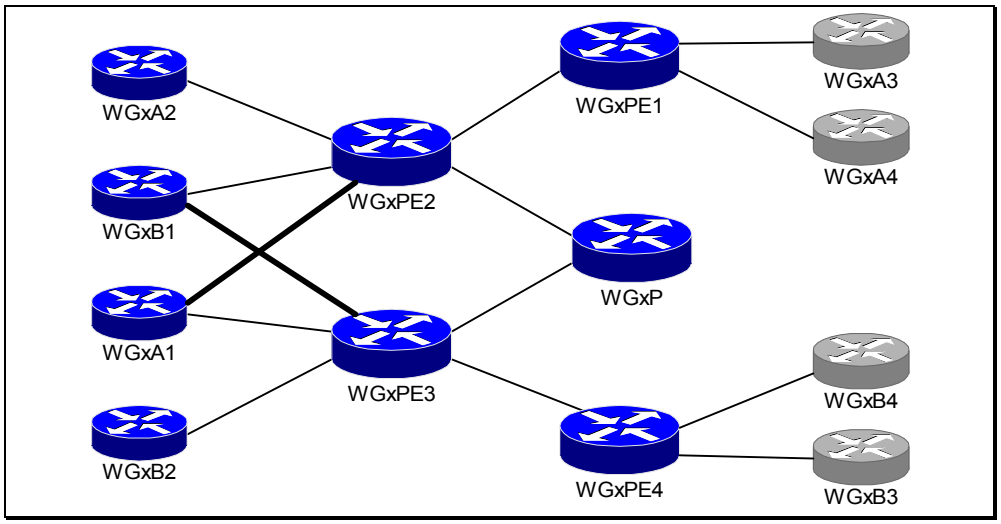

Figure 8: Multi-homed sites WGxA1 and WGxB1

Command list

Use the following commands to complete this exercise:

Command	Task
router bgp *as-number*	Select BGP configuration
ip vrf *name*	Creates a virtual routing and forwarding table
rd *value*	Assigns a route-distinguisher to a VRF
route-target import|export *value*	Assigns a route target to a VRF
address-family ipv4 vrf *name*	Selects per-VRF instance of a routing protocol
ip vrf forwarding *name*	Assigns an interface to a VRF
no neighbor *ip-address* shutdown	Enables a BGP neighbor previously disabled with the **neighbor shutdown** command
route-map *name* permit *seq*	Creates an entry in a route-map
set metric *value*	Sets BGP MED attribute in a route-map
neighbor *ip-address* route-map *name* in|out	Applies a route-map to BGP updates received from or sent to the specified neighbor
show ip bgp vpnv4 vrf *name*	Displays VPNv4 routes associated with the specified VRF
show ip route vrf *name*	Displays IP routing table of specified VRF
telnet host /vrf *name*	Telnets to a CE router connected to the specified VRF
ping vrf *name host*	Pings a host reachable through the specified VRF

Table 15: Configuration and monitoring commands used to configure BGP as the routing protocol between the PE-routers and the CE-routers

Task 1: Configure Additional PE-CE link

Perform the following configuration steps:

Step 1 Configure an additional subinterface on the existing serial interfaces on the PE and CE routers. Configure IP addresses and DLCIs on this interface using parameters in Table 16. Verify point-to-point connectivity over the new subinterface.

Source router	IP address	DLCI	Destination router	IP address	DLCI
A1	150.1.x1.209/30	312	PE2	150.1.x1.210/30	321
B1	150.1.x2.209/30	313	PE3	150.1.x2.210/30	331

Table 16: Additional Frame Relay PVC parameters

Task 2: Configure BGP as the PE-CE routing protocol

Step 1 Remove RIP and OSPF routing process from WGxA1/WGxB1.

Step 2 Reactivate the BGP neighbor on WGxA1/WGxB1 using the **no neighbor shutdown** command.

Step 3 Add the second BGP neighbor (the other PE router) on WGxA1/WGxB1 using the **neighbor** command.

Step 4 Configure the WGxA1/WGxB1 as a BGP neighbor within the VRF address family on WGxPE2 and WGxPE3.

Verification

- Check BGP connectivity with the **show ip bgp summary** and **show ip bgp neighbor** commands on CE routers.

```
WG2A1#sh ip bgp sum
BGP router identifier 202.1.1.1, local AS number 65021
BGP table version is 105, main routing table version 105
13 network entries and 26 paths using 2197 bytes of memory
12 BGP path attribute entries using 624 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
2 BGP extended community entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP activity 54/243 prefixes, 172/142 paths, scan interval 15 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
150.1.21.1      4     2    1129    1118      105    0    0 00:02:16        11
150.1.21.9      4     2    1078    1063      105    0    0 00:48:21        11
```

- Verify the BGP table on the WGxA1/WGxB1 with the **show ip bgp** command. You should see all the VPN routes:

```
WG2A1#sh ip bgp
BGP table version is 105, local router ID is 202.1.1.1
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
*> 150.1.21.0/30    0.0.0.0                  0         32768 ?
*> 150.1.21.4/30    150.1.21.1                    200      0 2 ?
*                   150.1.21.9               0    100      0 2 ?
*  150.1.21.8/30    150.1.21.1                    200      0 2 ?
*                   150.1.21.9               0    100      0 2 ?
*>                  0.0.0.0                  0         32768 ?
*> 150.1.21.128/30  150.1.21.1                    200      0 2 ?
*                   150.1.21.9                    100      0 2 ?
*> 150.1.21.132/30  150.1.21.1                    200      0 2 ?
*                   150.1.21.9                    100      0 2 ?
*> 202.1.0.0        150.1.21.1                    200      0 2 ?
*                   150.1.21.9               1    100      0 2 ?
*> 202.1.0.1/32     0.0.0.0                  0         32768 ?
*> 202.1.1.0        0.0.0.0                  0         32768 ?
*> 202.1.2.0        150.1.21.1                    200      0 2 ?
*                   150.1.21.9               1    100      0 2 ?
*> 202.1.127.3/32   150.1.21.1                    200      0 2 ?
*                   150.1.21.9                    100      0 2 ?
*> 202.1.127.4/32   150.1.21.1                    200      0 2 ?
*                   150.1.21.9                    100      0 2 ?
*> 202.1.134.0      150.1.21.1                    200      0 2 ?
*                   150.1.21.9                    100      0 2 ?
*> 202.1.135.0      150.1.21.1                    200      0 2 ?
*                   150.1.21.9                    100      0 2 ?
```

Note If you have not disabled OSPF or RIP on the WGxA1/WGxB1, you might see inconsistent BGP tables (for example, the entry for 150.1.21.8/30 in the printout above).

- Verify the per-VRF BGP table on the PE routers with the **show ip bgp vpnv4 vrf** command. You should see the BGP routes coming from the CE routers being selected as the best routes for those destinations:

```
WG2PE2#sh ip bgp vpnv4 vrf wg2a
BGP table version is 272, local router ID is 192.168.2.2
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: 2:10 (default for vrf a)
*  150.1.21.0/30    150.1.21.10           1000             0 65021 ?
*>i                 192.168.2.3              0    100      0 ?
*> 150.1.21.4/30    0.0.0.0                  0         32768 ?
*  150.1.21.8/30    150.1.21.10           1000             0 65021 ?
*>                  0.0.0.0                  0         32768 ?
*>i150.1.21.128/30  192.168.2.1              0    100      0 ?
*>i150.1.21.132/30  192.168.2.1              0    100      0 ?
*> 202.1.0.0        150.1.21.6               1         32768 ?
*>i202.1.0.1/32     192.168.2.3            100    100      0 65021 ?
*                   150.1.21.10           1000             0 65021 ?
*>i202.1.1.0        192.168.2.3            100    100      0 65021 ?
*                   150.1.21.10           1000             0 65021 ?
*> 202.1.2.0        150.1.21.6               1         32768 ?
*>i202.1.127.3/32   192.168.2.1              1    100      0 ?
*>i202.1.127.4/32   192.168.2.1            782    100      0 ?
*>i202.1.134.0      192.168.2.1              1    100      0 ?
*>i202.1.135.0      192.168.2.1            782    100      0 ?
```

Task 3: Select Primary and Backup Link with BGP

Step 1 Use BGP local preference on WGxA1/WGxB1 to select the link to WGxPE2 as the primary link and the link to WGxPE3 as the backup link.

Step 2 Set MED in outgoing routing updates from WGxA1/WGxB1 to make sure that the PE routers prefer the connection between the CE routers and WGxPE2.

Verification:

- Verify the proper setting of local preference on WGxA1/WGxB1 by using the **show ip bgp** command. Make sure that the routes received from WGxPE2 are always selected as the best routes.
- Verify the proper setting of MED by using the **show ip bgp vpnv4 vrf** command on the PE routers. Make sure that the PE routers select routes coming from WGxA1/WGxB1 through WGxPE2 as the best routes.
- Shutdown the subinterface between WGxA1/WGxB1 and WGxPE2 while concurrently performing continuous **ping** from WGxPE1 (WGxPE4 for subgroup B) to the CE router. Count the lost responses and measure the switchover time.
- Re-enable the subinterface between WGxA1/WGxB1 and WGxPE2 and verify whether the connectivity is retained throughout the convergence process using continuous **ping** from WGxPE1 to the CE router.

Task 4: Convergence Time Optimization

Step 1 Change the BGP timers on WGxA1/WGxB1 and WGxPE2 using the **neighbor timers** command. Set the keepalive timer to 5 seconds and the holdtime to 15 seconds. Clear the BGP session using **clear ip bgp** command to establish the new timer values.

Verification

Repeat the convergence time measurements from the previous task.

D

Laboratory Exercises—MPLS VPN Topologies

Overview

This chapter contains a set of exercises to provide you with insight into advanced MPLS VPN topologies.

It includes the following exercises:

- Overlapping VPN topology
- Network Management VPN
- Internet access with packet leaking
- Internet access through a dedicated subinterface
- Internet-in-a-VPN

These exercises support the **MPLS VPN Topologies** and **Internet Access from the VPN** chapters and presumes knowledge of the MPLS VPN infrastructure, already established through the **MPLS VPN Implementation** exercises.

Laboratory Exercise D-1: Overlapping VPN Topology

Objective

Your VPN customers want to exchange data between their central sites. You have decided to implement this request with an overlapping VPN topology.

In this laboratory exercise, you will establish overlapping VPNs with the following connectivity goals:

- WGxA1, WGxA2, WGxA3, and WGxA4 can communicate
- WGxB1, WGxB2, WGxB3, and WGxB3 can communicate
- WGxA1 and WGxB1 can communicate
- WGxA1 cannot reach WGxB2, WGxB3, or WGxB4
- WGxB1 cannot reach WGxA2, WGxB3, or WGxB4

Visual Objective

This figure shows the desired VPN connectivity.

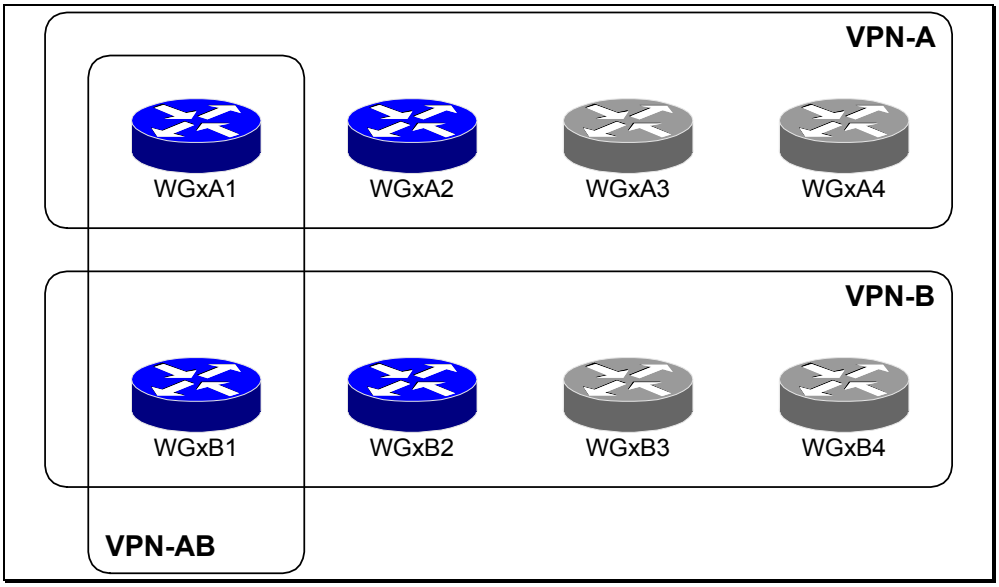

Figure 9: Overlapping VPNs

The logical connectivity between the PE-routers and the CE-routers that were set up during the **Running BGP between the PE and the CE routers** exercise is displayed in the following figure:

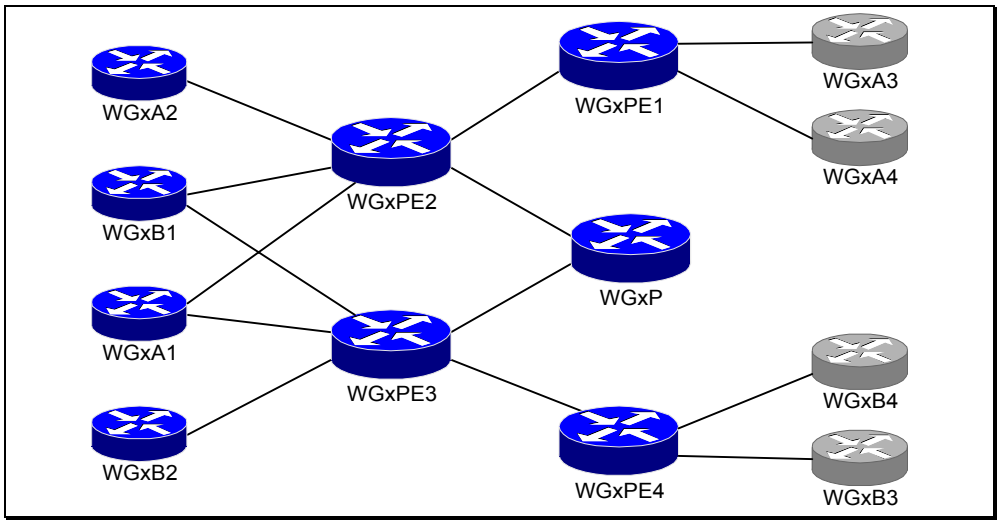

Figure 10: Logical connectivity between the PE-routers and the CE-routers

Command list

Use the following commands to complete this exercise:

Command	Task
router bgp *as-number*	Select BGP configuration
address-family ipv4 vrf *name*	Selects per-VRF instance of a routing protocol
neighbor *ip-address* remote-as *as-number*	Defines a new BGP neighbor
no neighbor *ip-address*	Removes a BGP neighbor
ip vrf *name*	Creates a virtual routing and forwarding table
rd *value*	Assigns a route-distinguisher to a VRF
route-target import|export *value*	Assigns a route target to a VRF
ip vrf forwarding *name*	Assigns an interface to a VRF
show ip vrf detail	Displays detailed VRF information
show ip bgp vpnv4 vrf *name*	Displays VPNv4 routes associated with the specified VRF
show ip route vrf *name*	Displays IP routing table of specified VRF
telnet host /vrf *name*	Telnets to a CE router connected to the specified VRF
ping vrf *name host*	Pings a host reachable through the specified VRF

Table 17: Configuration and monitoring commands used to configure overlapping VPN topology

Task 1: Design your VPN solution

Site WGxA1 cannot belong to the same VRF as the other WGxA sites. Similarly, site WGxB1 cannot belong to the same VRF as the WGxB sites. Also, WGxA1 and WGxB1 cannot share the same VRF.

Step 1 Allocate new route distinguishers for VRFs to which WGxA1 and WGxB1 will be connected.

Step 2 A new route target is needed for VPN-AB. Coordinate the value of this route target with the other subgroup within your workgroup.

Note You could use x:11 as the RD for VRFs connected to WGxA1, x:21 as the RD for VRFs connected to WGxB1, and x:30 as the route-target for the VPN-AB.

Task 2: Remove WGxA1/WGxB1 from existing VRFs

Sites WGxA1 and WGxB1 have to be migrated to new VRFs. All the references to them must be removed from the routing protocol contexts.

Step 1 Remove BGP neighbors WGxA1 and WGxB1 from the PE-routers.

Step 2 Check any other references to WGxA1 or WGxB1 in the PE-router configuration and, if required, remove them.

Task 3: Configure new VRFs for WGxA1 and WGxB1

Step 1 Create VRFs for WGxA1 and WGxB1 on WGxPE2 and WGxPE3 with the **ip vrf** command.

Step 2 Assign new route distinguishers to the newly created VRFs with the **rd** command.

Step 3 Assign proper import and export route-targets to the newly created VRFs with the **route-target** command.

Step 4 Re-establish BGP routing between the PE-routers and the CE-routers. Please refer to the **Running BGP between PE and CE routers** lab exercise if you need more details.

Verification:

- On the PE-router, verify that the interface toward the CE-router is in the proper VRF by using the **show ip vrf interfaces** command. This should result in a printout similar to the one below:

```
WG3PE3#show ip vrf interfaces wg3a1
Interface              IP-Address      VRF                              Protocol
Serial0/0.3            150.1.31.1      wg3a1                            up
```


- Verify the BGP neighbors on the PE-router with the **show ip bgp vpnv4 vrf summary** command. This should give you a printout similar to the one below. Check the status of the WGxA1 or WGxB1 in the printout.

```
WG3PE3#show ip bgp vpnv4 vrf wg3a1 sum
BGP router identifier 192.168.3.3, local AS number 3
BGP table version is 113, main routing table version 113
16 network entries and 20 paths using 3264 bytes of memory
28 BGP path attribute entries using 1456 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
15 BGP AS-PATH entries using 360 bytes of memory
4 BGP extended community entries using 96 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP activity 138/183 prefixes, 551/398 paths, scan interval 15 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
150.1.31.2      4 65031      13      18      113    0    0 00:09:47        4
192.168.3.2     4     3    3170    3140      113    0    0 09:04:00       16
192.168.3.4     4     3    3074    3160      113    0    0 09:03:24        0
```

- Check the BGP routing table in the new VRF (wg3a1 in our example) with the **show ip bgp vpnv4 vrf** command. You should see routes from the WGxA1 or WGxB1 as well as routes imported from other VRFs. Use the AS-path to work out which routes belong to which CE-router. Routes announced by WGxA1 should have 650x1 in the AS-path, routes announced by WGxB1 should have 650x2 in the AS-path and all other routes should have an empty AS-path.

```
WG3PE3#show ip bgp vpnv4 vrf wg3a1
BGP table version is 113, local router ID is 192.168.3.3
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: 3:11 (default for vrf wg3a1)
*  150.1.31.0/30    150.1.31.2             100             0 65031 ?
*>i                 192.168.3.2             50    100      0 65031 ?
*>i150.1.31.4/30    192.168.3.2              0    100      0 ?
*>i150.1.31.128/30  192.168.3.1              0    100      0 ?
*>i150.1.31.132/30  192.168.3.1              0    100      0 ?
*  150.1.31.208/30  150.1.31.2             100             0 65031 ?
*>i                 192.168.3.2             50    100      0 65031 ?
*>i150.1.32.0/30    192.168.3.2              0    100      0 65032 ?
*>i203.1.0.0        192.168.3.2              1    100      0 ?
*  203.1.0.1/32     150.1.31.2             100             0 65031 ?
*>i                 192.168.3.2             50    100      0 65031 ?
*  203.1.1.0        150.1.31.2             100             0 65031 ?
*>i                 192.168.3.2             50    100      0 65031 ?
*>i203.1.2.0        192.168.3.2              1    100      0 ?
*>i203.1.127.3/32   192.168.3.1              1    100      0 ?
*>i203.1.127.4/32   192.168.3.1             65    100      0 ?
*>i203.1.134.0      192.168.3.1              1    100      0 ?
*>i203.1.135.0      192.168.3.1             65    100      0 ?
*>i203.2.0.1/32     192.168.3.2              0    100      0 65032 ?
*>i203.2.1.0        192.168.3.2              0    100      0 65032 ?
```

These routes are coming from WGxA1

These routes are coming from other CE-routers

These routes are coming from WGxB1

- Use the **show ip bgp vpnv4 vrf *name prefix*** command to display details of an individual route and verify that the proper route-targets are attached to the route. Your printout should be similar to the one below:

```
WG3PE3#show ip bgp vpnv4 vrf wg3a1 203.1.1.0
BGP routing table entry for 3:11:203.1.1.0/24, version 107
Paths: (2 available, best #2, table wg3a1)
  Advertised to non peer-group peers:
  150.1.31.2 192.168.3.4
  65031
    150.1.31.2 from 150.1.31.2 (203.1.1.1)
      Origin incomplete, metric 100, localpref 100, valid, external
      Extended Community: RT:3:10 RT:3:30
  65031
    192.168.3.2 (metric 20) from 192.168.3.2 (192.168.3.2)
      Origin incomplete, metric 50, localpref 100, valid, internal, best
      Extended Community: RT:3:10 RT:3:30
```

- Telnet to WGxA1 and perform **ping** and **trace** to the loopback address of WGxB1 (or vice versa). The other router should be reachable. For subgroup B, perform the test in the other direction.

```
WG3PE3#telnet a1 /vrf wg3a1
Trying A1 (203.1.0.1)... Open

WG3A1#ping b1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.2.0.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/64/68 ms
WG3A1#trace b1

Type escape sequence to abort.
Tracing the route to B1 (203.2.0.1)

  1 150.1.31.210 24 msec 20 msec 20 msec
  2 150.1.32.2 [AS 65032] 32 msec *  36 msec
```

- Telnet to WGxA2 and try to **ping** WGxB1 or WGxB2. Those routers should not be reachable from WGxA2. For subgroup B, **ping** WGxA1 and WGxA2 from WGxB2.

```
WG3A1#a2
Trying A2 (203.1.0.2)... Open

WG3A2#ping b1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.2.0.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
WG3A2#trace b1

Type escape sequence to abort.
Tracing the route to B1 (203.2.0.1)

  1 150.1.31.5 44 msec 32 msec 36 msec
  2 150.1.31.5 !H  *  !H
WG3A2#ping b2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.2.0.2, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
```



```
WG3A2#trace b2

Type escape sequence to abort.
Tracing the route to B2 (203.2.0.2)

  1 150.1.31.5 40 msec 32 msec 36 msec
  2 150.1.31.5 !H  *  !H
```

Laboratory Exercise D-2: Common Services VPN

Objective

MPLS VPN infrastructure can be used to implement a new approach to managed CE-router service, where the central Network Management Station (NMS) can monitor all CE-routers through a dedicated Virtual Private Network. The NMS VPN should only provide connectivity between NMS and a single IP address on the CE-router that is used for network management purposes.

In this exercise, you will establish a network management VPN between the loopback interfaces of the CE-routers and the NMS router. You will only establish connectivity between the NMS and the CE-router loopback interfaces with a /32 subnet mask.

To achieve this objective, you will complete the following tasks:

- Design your Network Management VPN
- Configure a VRF for the management LAN
- Establish connectivity between management VRF and customer VRFs by configuring proper route targets
- Establish routing between the PE-router and the NMS router to propagate routes to CE-router loopback interfaces to the NMS router

Background Information

The connectivity between the customer routers, the NMS router and the PE-routers is shown in the figure below:

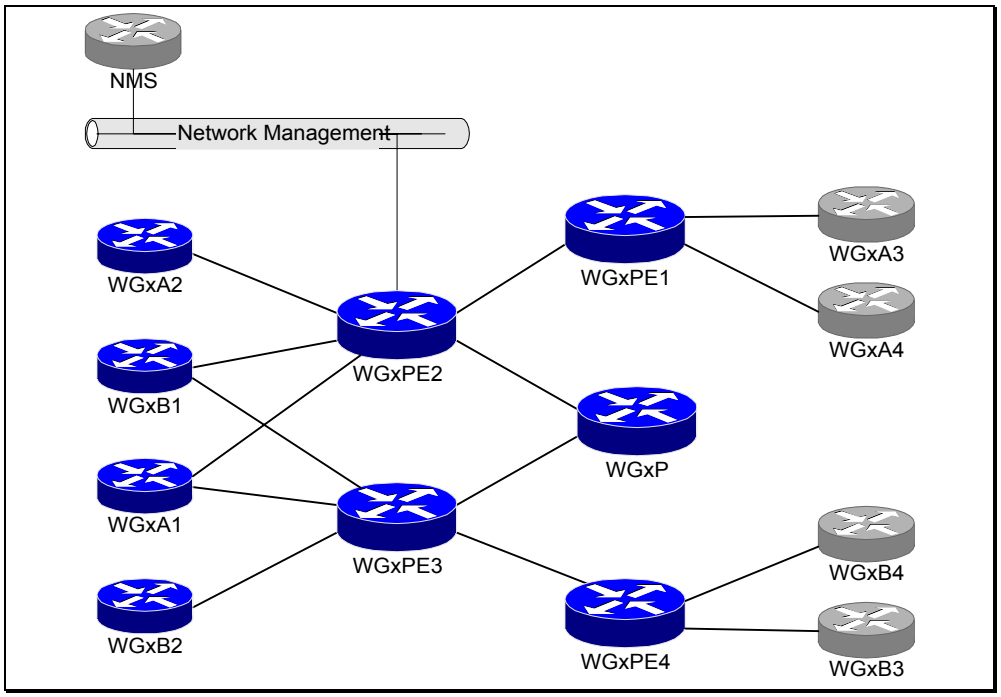

Figure 11: Logical connectivity between the NMS router, CE-routers, and PE-routers

Note	The NMS router is shared between workgroups and is not configurable.

Command list

Use the following commands to complete this exercise:

Command	Task
router bgp *as-number*	Select BGP configuration
address-family ipv4 vrf *name*	Selects per-VRF instance of a routing protocol
ip vrf *name*	Creates a virtual routing and forwarding table
rd *value*	Assigns a route-distinguisher to a VRF
route-target import|export *value*	Assigns a route target to a VRF
ip vrf forwarding *name*	Assigns an interface to a VRF
redistribute bgp *as-number* metric *value*	Redistribute BGP routes into RIP, specifying RIP hop count for the redistributed routes
ip prefix-list *name* permit *address mask* ge *len*	Creates an IP prefix-list that matches all prefixes in specified address space with subnet mask longer or equal to the specified value
route-map *name* permit *seq*	Creates a route-map entry
match ip address prefix-list *list*	Matches a prefix in a route-map with specified IP prefix-list
set extcommunity rt *value* additive	Appends the specified route target to route matched with the **match** command
export map *name*	Specifies a VRF export route-map
show ip vrf detail	Displays detailed VRF information
show ip bgp vpnv4 vrf *name*	Displays VPNv4 routes associated with the specified VRF
show ip route vrf *name*	Displays IP routing table of specified VRF
telnet host /vrf *name*	Telnets to a CE router connected to the specified VRF
ping vrf *name host*	Pings a host reachable through the specified VRF

Table 18: Configuration and monitoring commands used to configure Network Management VPN

Task 1: Design your Network Management VPN

Network management VPN is a *common services* VPN; therefore you need two route-targets for the VPN—the server route-target and the client route-target. You also need a new VRF for the Network Management LAN and associated route distinguisher.

Step 1 Allocate two route targets for a new Network Management VPN and a route distinguisher for the NMS VRF.

Note You could use x:500 as the RD for NMS VRF, route-target x:500 as the *server* route-target and x:501 as the *client* route-target.

Task 2: Create Network Management VRF

Step 1 Create a new VRF on WGxPE2 with the **ip vrf** command and configure RD allocated in the previous step with the **rd** command.

Step 2 Configure route import and export in the new VRF with the **route-target** command. The VRF should import client and server routes and export routes with the server route-target.

Verification

Verify the parameters of the new VRF with the **show ip vrf *name* detail** command on the WGxPE2.

Task 3: Establish connectivity between NMS VRF and other VRFs

To establish connectivity between the NMS VRF and the customer VRF you must attach the *client* route-target to routes toward CE-router loopback addresses when they are exported from the customer VRF. You also have to import route toward NMS router into all customer VRFs.

Step 1 Create an **ip prefix-list** that will match CE-router loopback addresses.

Step 2 Create a **route-map** that will match the CE-router loopback addresses with the prefix-list and append the *client* route-target to those routes.

Step 3 Apply the route-map to routes exported from the customer VRF with the **export route-map** command.

Step 4 Import NMS routes into the customer VRF by specifying the proper import route-target.

Verification

- Verify that the proper route-targets are appended to the routes toward CE-router loopback addresses by using the **show ip bgp vpnv4 vrf *name prefix*** command. This should result in a printout similar to the one below:

```
WG3PE2#show ip bgp vpnv4 vrf wg3a1 203.2.0.1 255.255.255.255
BGP routing table entry for 3:11:203.2.0.1/32, version 66
Paths: (1 available, best #1, table wg3a1)
  Advertised to non peer-group peers:
  150.1.31.209
  65032, imported path from 3:21:203.2.0.1/32
    150.1.32.2 from 150.1.32.2 (203.2.1.1)
      Origin incomplete, metric 0, localpref 100, valid, external, best
      Extended Community: RT:3:20 RT:3:30 RT:3:501
```

- Verify that the route toward CE-router loopback address is inserted into the NMS VRF by using the **show ip route vrf** command on WGxPE2.

```
WG3PE2#show ip route vrf NMS
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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     203.1.0.0 255.255.255.255 is subnetted, 2 subnets
B       203.1.0.2 [20/1] via 150.1.31.6 (wg3a), 00:02:26, Serial0/0.3
B       203.1.0.1 [20/50] via 150.1.31.209 (wg3a1), 00:05:26
     203.2.0.0 255.255.255.255 is subnetted, 1 subnets
B       203.2.0.1 [20/0] via 150.1.32.2 (wg3b1), 00:05:26
C    192.168.22.0 255.255.255.0 is directly connected,
```

- You can also check individual routes imported into the NMS VRF with the **show ip bgp vpnv4 vrf** ***name prefix*** command on WGxPE2.

```
WG3PE2#show ip bgp vpnv4 vrf NMS 203.1.0.2 255.255.255.255
BGP routing table entry for 3:500:203.1.0.2/32, version 82
Paths: (1 available, best #1, table NMS)
  Not advertised to any peer
  Local, imported path from 3:10:203.1.0.2/32
    150.1.31.6 from 0.0.0.0 (192.168.3.2)
      Origin incomplete, metric 1, localpref 100, weight 32768, valid, external,
 best
      Extended Community: RT:3:10 RT:3:501
```

Task 4: Establish routing between WGxPE2 and the NMS router

The routes toward loopback interfaces of the CE-routers are already present in the NMS VRF. These routes need to be announced to the NMS router via RIP.

Step 1 Configure RIP routing with the NMS router on the WGxPE2.

Step 2 Redistribute BGP routes into the RIP routing process. Because these routes are derived from a variety of routing protocols, you have to specify the RIP metric manually with the **redistribute bgp** ***as*** **metric** ***metric*** command.

Note Please refer to exercise **Initial MPLS VPN setup** if you need more information on configuring RIP between the PE-routers and the CE-routers.

Verification

- Verify propagation of routes toward loopback addresses of CE-routers by using the **show ip route** command on the NMS router.

```
WG3PE2#telnet 192.168.22.22 /vrf NMS
Trying 192.168.22.22 ... Open

NMS>show ip route
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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR

Gateway of last resort is not set

     203.1.0.0/32 is subnetted, 2 subnets
R       203.1.0.2 [120/1] via 192.168.22.3, 00:00:05, Ethernet0
R       203.1.0.1 [120/1] via 192.168.22.3, 00:00:05, Ethernet0
     203.2.0.0/32 is subnetted, 1 subnets
R       203.2.0.1 [120/1] via 192.168.22.3, 00:00:05, Ethernet0
C    192.168.22.0/24 is directly connected, Ethernet0
```

- Perform **ping** and **trace** from the NMS router toward individual loopback addresses.
- Perform **ping** and **trace** from the CE-routers toward the NMS router. These operations will fail unless you perform extended **ping** and **trace,** specifying the loopback interface as the source IP address.

Laboratory Exercise D-3: Internet Connectivity Through Route Leaking

Objective

Some customers want to reach global Internet directly from their VPN. The MPLS VPN implementation on Cisco IOS allows you to implement this solution with static routes that facilitate packet propagation between the customer VPN and the global IP routing table.

In this laboratory exercise, you will complete the following task:

- Establish Internet connectivity to the WGxA1 and WGxB1 by configuring global and VRF static routes

Visual Objective

Route leaking between the customer VPN and the global routing table will be established on the emphasized links in the diagram below to enable connectivity between the CE-routers and the Internet destinations (routers *Good, Cheap,* and *Client*).

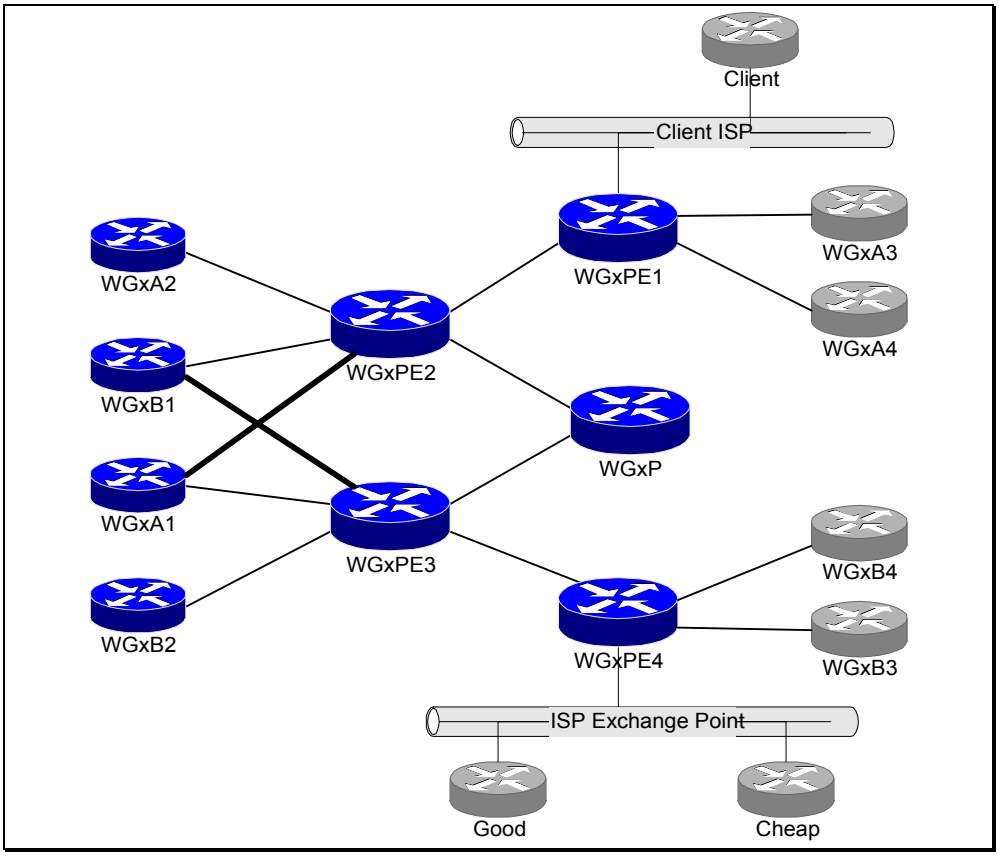

Figure 12: Route leaking between VPN and Internet

Command list

Use the following commands to complete this exercise:

Command	Task
router bgp *as-number*	Select BGP configuration
address-family ipv4 vrf *name*	Selects per-VRF instance of a routing protocol
neighbor *ip-address* default-originate	Announces default route to specified neighbor in BGP updates
ip route *prefix mask interface* [tag *value*]	Creates a global static route
redistribute static	Redistributes static routes into a routing protocol
ip route vrf *name prefix mask next-hop* global	Creates a VRF static route with a global next-hop
no route-target import|export *value*	Removes a route target from a VRF
show ip bgp vpnv4 vrf *name*	Displays VPNv4 routes associated with the specified VRF
show ip route vrf *name*	Displays IP routing table of specified VRF
telnet host /vrf *name*	Telnets to a CE router connected to the specified VRF
ping vrf *name host*	Pings a host reachable through the specified VRF

Table 19: Configuration and monitoring commands used to configure simple VPN with RIP routing

Task 1: Cleanup from the previous VPN exercises

If you have completed all MPLS VPN topology exercises so far, WGxA1 and WGxB1 will now participate in two VPNs. The connectivity between WGxA1 and WGxB1 needs to be broken before you establish Internet connectivity.

Step 1 Break the connectivity between WGxA1 and WGxB1 by removing the route-target corresponding to VPN-AB from the VRFs to which WGxA1 and WGxB1 are connected with the **no route-target** command.

Task 2: Configure route leaking between customer VPN and the Internet

You will establish the route leaking only between WGxB1 – WGxPE3 and WGxA1 – WGxPE2 using the following steps:

Step 1 Subgroup A configures a global static route for 20x.1.0.0/16 toward WGxA1 on WGxPE2. The /16 route is used to cover the whole address space allocated to customer A. Because BGP redistribution is already configured on the WGxPE2, this route will be redistributed into BGP automatically if you assign tag 10 to it with the **ip route … tag 10** command.

Step 2 Similar to the previous step, subgroup B configures a global static route for 20x.2.0.0/15 toward WGxB1 on WGxPE3.

Step 3 Subgroup A configures a default route in the VRF to which the WGxA1 is connected on WGxPE2. The next-hop should be WGxPE4.

Step 4 Similar to the previous step, subgroup B configures a default route in the VRF to which the WGxB1 is connected on WGxPE3. The next-hop should be WGxPE1.

Step 5 Configure propagation of default route toward the WGxA1 and WGxB1 by using the **neighbor *address* default-originate** command on the PE-routers.

Verification

- Verify that the static route toward the customer's address space is inserted in the global routing table with the **show ip route *prefix*** command.

```
WG3PE2#sh ip route 203.1.0.0
Routing entry for 203.1.0.0 255.255.0.0, supernet
  Known via "static", distance 1, metric 0 (connected)
  Tag 10
  Redistributing via bgp 3
  Advertised by bgp 3 route-map TAG
  Routing Descriptor Blocks:
  * directly connected, via Serial0/0.31
      Route metric is 0, traffic share count is 1
```

- Verify that the static route toward the customer's address space gets redistributed into BGP with the **show ip bgp *prefix*** command.

```
WG3PE2#sh ip bgp 203.1.0.0
BGP routing table entry for 203.1.0.0/16, version 115
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  192.168.3.1 192.168.3.3
  Local
    0.0.0.0 from 0.0.0.0 (192.168.3.2)
      Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced,
```

- Verify the presence of the default route in the VRF associated with WGxA1 or WGxB1 with the **show ip route vrf *name*** command.

```
WG3PE2#sh ip route vrf wg3a1
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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is 192.168.3.4 to network 0.0.0.0

B    203.1.2.0 255.255.255.0
           [20/1] via 150.1.31.6 (wg3a), 22:42:37, Serial0/0.3
B    203.1.1.0 255.255.255.0 [20/50] via 150.1.31.209, 22:43:43
B    203.2.1.0 255.255.255.0 [20/0] via 150.1.32.2 (wg3b1), 22:43:37
     203.1.0.0 255.255.255.255 is subnetted, 2 subnets

… rest deleted …
```


- Verify that the next-hop of the default route in the VRF is associated with a global next-hop with the **show ip route vrf *name prefix*** command.

```
WG3PE2#sh ip route vrf wg3a1 0.0.0.0
Routing entry for 0.0.0.0 0.0.0.0, supernet
  Known via "static", distance 1, metric 0, candidate default path
  Routing Descriptor Blocks:
  * 192.168.3.4 (Default-IP-Routing-Table)
      Route metric is 0, traffic share count is 1
```

- Verify connectivity from the Internet routers to WGxA1 and WGxB1 using the **trace** command.

```
Client>trace 203.1.0.1

Type escape sequence to abort.
Tracing the route to 203.1.0.1
  1 192.168.21.3 4 msec 4 msec 4 msec
  2 192.168.3.21 [AS 3] 16 msec 28 msec 24 msec
  3 150.1.31.209 56 msec *  56 msec
```

- Verify the connectivity from the Internet routers to WGxA2 and WGxB2 using the **trace** command.

```
Client>trace 203.1.0.2

Type escape sequence to abort.
Tracing the route to 203.1.0.2
  1 192.168.21.3 4 msec 4 msec 4 msec
  2 192.168.3.21 [AS 3] 20 msec 28 msec 24 msec
  3 150.1.31.209 56 msec 60 msec 60 msec
  4 150.1.31.210 52 msec 56 msec 48 msec
  5  *  *  *
```

Additional exercise: Fix intra-VPN routing

You will probably encounter a display very similar to the one above—the trace proceeds past WGxA1 back to the PE-router and proceeds no further. Try to analyze the routing within the customer VPN and fix the intra-VPN routing so that all customer sites become accessible from the Internet.

Laboratory Exercise D-4: Separate Interface for Internet Connectivity

Objective

The Internet access directly from a VPN site is not acceptable to some customers due to security concerns—these customers want to retain the traditional Internet access model with a firewall between the customer VPN and the global Internet. This request is usually implemented by using dedicated VPN and Internet subinterfaces on the physical PE-CE link.

In this laboratory exercise, you will complete the following tasks:

- Configure a dedicated Internet subinterface between the PE-router and the CE-router
- Establish Internet connectivity from the customer solely through the global routing table

Visual Objective

You will configure additional virtual circuits (emphasized in the diagram below) between the PE-routers and the CE-routers. These circuits will be in the global routing table and you will configure a global BGP session between PE-routers and CE-routers to exchange Internet routes between the Service Provider and the Customer.

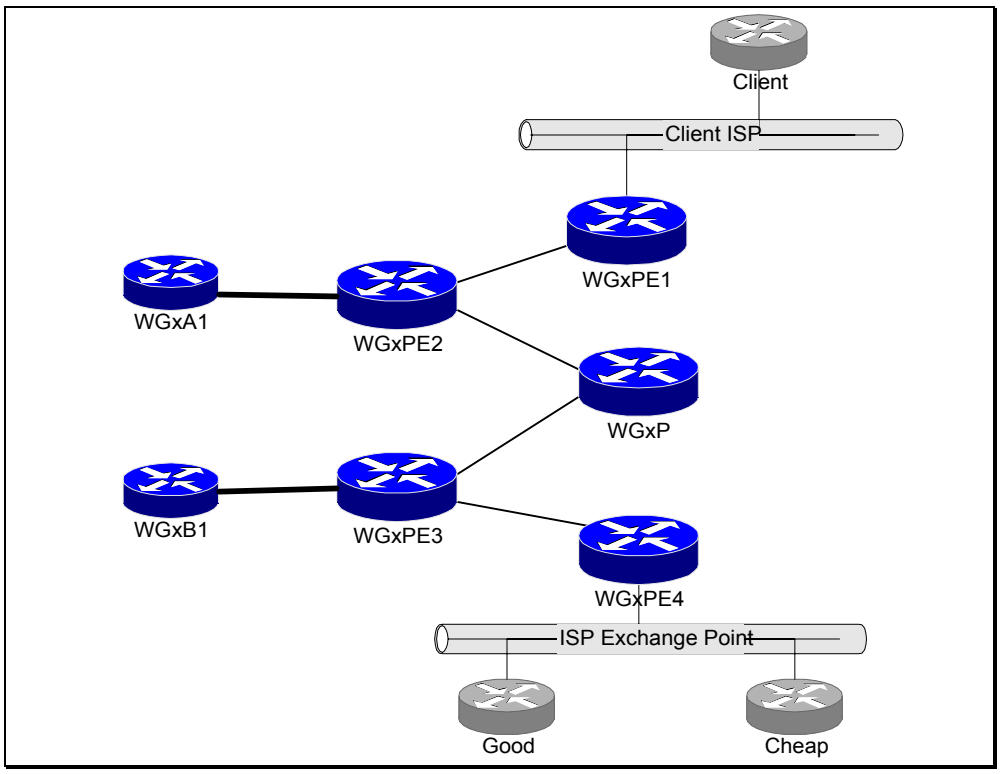

Figure 13: Internet routing in a global routing table

Command list

Use the following commands to complete this exercise:

Command	Task
router bgp *as-number*	Select BGP configuration
address-family ipv4 vrf *name*	Selects per-VRF instance of a routing protocol
neighbor *ip-address* prefix-list *list* out	Filters BGP updates sent to the specified neighbor through the specified prefix-list
neighbor *ip-address* default-originate	Announces default route to specified neighbor in BGP updates
no neighbor *ip-address* default-originate	Stops announcing default route to specified neighbor in BGP updates
no ip route *prefix mask interface* [tag *value*]	Removes a global static route
no ip route vrf *name prefix mask next-hop* global	Removes a VRF static route with a global next-hop
ip route *prefix mask* null 0	Creates a summary route in the IP routing table
network *prefix* mask *mask*	Announces an IP prefix in the BGP process
ip prefix-list *name* permit|deny *address mask* [ge|le *len*]	Creates an IP prefix-list that permits or denies all prefixes in specified address space with subnet mask longer or equal (or shorter or equal) to the specified value
show ip bgp vpnv4 vrf *name*	Displays VPNv4 routes associated with the specified VRF
show ip route vrf *name*	Displays IP routing table of specified VRF
telnet host /vrf *name*	Telnets to a CE router connected to the specified VRF
ping vrf *name host*	Pings a host reachable through the specified VRF

Table 20: Configuration and monitoring commands used to Internet access through dedicated subinterface

Task 1: Cleanup from the previous exercise

Step 1 Remove the static routes from the previous lab to disable connectivity between WGxA1/WGxB1 and the Internet.

Step 2 Disable default route advertising from the PE-routers toward the CE-routers by using the **no neighbor *address* default-originate** command.

Verification

- Verify that there is no default route in the VPN by using the **show ip route vrf** command

```
WG3PE2#sh ip route vrf wg3a1
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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

B    203.1.1.0 255.255.255.0 [20/50] via 150.1.31.209, 22:53:14
     203.1.0.0 255.255.255.255 is subnetted, 1 subnets
B       203.1.0.1 [20/50] via 150.1.31.209, 22:53:14
     203.2.0.0 255.255.255.255 is subnetted, 1 subnets
B    203.1.135.0 255.255.255.0 [200/65] via 192.168.3.1, 22:53:22
B    203.1.134.0 255.255.255.0 [200/1] via 192.168.3.1, 22:53:22
     203.1.127.0 255.255.255.255 is subnetted, 2 subnets
B       203.1.127.4 [200/65] via 192.168.3.1, 22:53:22
B       203.1.127.3 [200/1] via 192.168.3.1, 22:53:22
     150.1.0.0 255.255.255.252 is subnetted, 6 subnets
B       150.1.31.128 [200/0] via 192.168.3.1, 22:53:22
B       150.1.31.132 [200/0] via 192.168.3.1, 22:53:22
C       150.1.31.208 is directly connected, Serial0/0.31
B       150.1.31.0 [20/50] via 150.1.31.209, 22:53:14
B       150.1.31.4 is directly connected, 22:52:07, Serial0/0.3
```

- Verify that the customer's address space is no longer reachable in the global routing table by using the **show ip route** ***prefix*** command

```
WG3PE2#show ip route 203.1.0.0
% Network not in table
```

Task 2: Establishing connectivity in the global routing table

Step 1 Create a separate subinterface on WGxA1 – WGxPE2 and WGxB1 – WGxPE3 and put that subinterface into the global routing using the parameters in Table 21.

Source router	Destination router	DLCI
WGxA1	WGxPE2	612
WGxPE2	WGxA1	621
WGxB1	WGxPE3	613
WGxPE3	WGxB1	631

Table 21: Additional PVCs for Internet connectivity

Task 3: Routing between the PE-router and the CE-router

Step 1 Configure BGP over the newly created subinterface.

Step 2 Announce only the default route from the PE-router to the CE-router using a proper combination of the **neighbor default-originate** and **neighbor prefix-list** commands.

Step 3 Advertise 20x.1.0.0/16 from the WGxA1 and 20x.2.0.0/16 from the WGxB1 to the Internet. Do not advertise any other VPN prefixes to the Internet.

Verification

- Verify BGP connectivity between the PE-router and the CE-router using the **show ip bgp summary** command

```
WG3PE2#show ip bgp summary
BGP router identifier 192.168.3.2, local AS number 3
BGP table version is 118, main routing table version 118
72 network entries and 114 paths using 11088 bytes of memory
38 BGP path attribute entries using 1976 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
16 BGP AS-PATH entries using 384 bytes of memory
8 BGP extended community entries using 224 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP activity 276/2010 prefixes, 788/624 paths, scan interval 15 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
150.1.31.213    4 65031       5       4      118    0    0 00:00:59        1
192.168.3.1     4     3    4552    4642      118    0    0 23:01:12       44
192.168.3.3     4     3    4563    4628      118    0    0 23:01:14       65
```

- Verify that the CE-router only advertises a single prefix to the PE-router by using the **show ip bgp regexp** command

```
WG3PE2#show ip bgp regexp 65031
BGP table version is 118, local router ID is 192.168.3.2
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
*> 203.1.0.0/16     150.1.31.213             0             0 65031 i
```

- Telnet to one of the Internet routers and perform **trace** toward the WGxA1 or WGxB1

```
WG3PE2# telnet 192.168.21.99
Trying 192.168.21.99 ... Open

Client>trace 203.1.0.1

Type escape sequence to abort.
Tracing the route to 203.1.0.1

  1 192.168.21.3 4 msec 0 msec 4 msec
  2 192.168.3.21 [AS 3] 16 msec 28 msec 20 msec
  3 150.1.31.213 [AS 3] 40 msec *  44 msec
```


Laboratory Exercise D-5: Internet in a VPN

Objective

Internet connectivity in MPLS VPN-based networks can be achieved through the global IP routing table or through a dedicated Internet VPN. The dedicated Internet VPN approach gives you better security as it completely isolates the Service Provider core (P-routers) from the Internet. On the other hand, it is also less scaleable—for example, you cannot transport full Internet routing in an Internet VPN.

In this laboratory exercise, you will complete the following task:

- Establish Internet connectivity between the VPN customers and the Internet by creating an Internet VPN, thus isolating Internet routing from the MPLS VPN backbone

Visual Objective

This figure shows all the relevant parts of your workgroup. The links you will move to a new Internet VPN are emphasized.

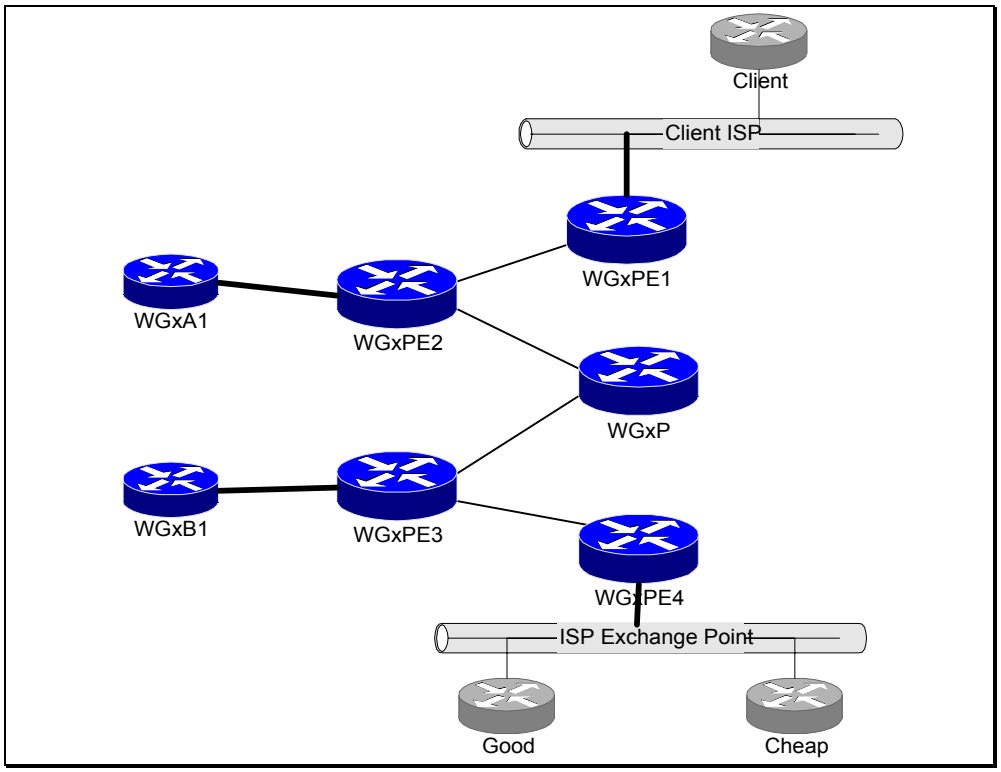

Figure 14: Internet in a VPN

Command list

Use the following commands to complete this exercise:

Command	Task
router bgp *as-number*	Select BGP configuration
address-family ipv4 vrf *name*	Selects per-VRF instance of a routing protocol
neighbor *ip-address* remote-as *as-number*	Configures a BGP neighbor
neighbor *ip-address* activate	Activates the specified BGP neighbor
ip vrf *name*	Creates a virtual routing and forwarding table
rd *value*	Assigns a route-distinguisher to a VRF
route-target import|export *value*	Assigns a route target to a VRF
ip vrf forwarding *name*	Assigns an interface to a VRF
show ip vrf detail	Displays detailed VRF information
show ip bgp neighbor	Displays information on global BGP neighbors
show ip bgp vpnv4 vrf *name*	Displays VPNv4 routes associated with the specified VRF
show ip route vrf *name*	Displays IP routing table of specified VRF
telnet host /vrf *name*	Telnets to a CE router connected to the specified VRF
ping vrf *name host*	Pings a host reachable through the specified VRF

Table 22: Configuration and monitoring commands used to configure Internet VPN

Task 1: Design your Internet VPN

You will need a new route-target for the Internet VPN. You will also need a route-distinguisher for all Internet-related VRFs. Use route-target x:600 and route-distinguisher x:600.

Task 2: Migrate Internet routers in a VPN

Step 1 Create Internet VRF on all PE routers using the **ip vrf**, **rd** and **route-target** commands.

Step 2 Remove all EBGP neighbors from the global BGP process.

Step 3 Migrate links to both Internet backbones and the links toward WGxA1 and WGxB1, which were in the global address space, into the new VPN.

Step 4 Reestablish BGP routing within the new VPN.

---

*This exercise is only used to illustrate the Internet-in-a-VPN principle and does **NOT** contain all the necessary steps. Full Internet routing shall **NEVER** be inserted in a VPN.*

Verification

- Verify the proper setup of Internet VRF on the PE-routers with the **show ip vrf detail** command

```
WG3PE2#sh ip vrf detail
VRF Internet; default RD 3:600
  Interfaces:
    Serial0/0.40
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:3:600
  Import VPN route-target communities
    RT:3:600
  No import route-map
  No export route-map
```

- Verify BGP neighbors in the Internet VRF by using the **show ip bgp vrf *name* summary** command

```
WG3PE2#show ip bgp vpnv4 vrf Internet sum
BGP router identifier 192.168.3.2, local AS number 3
BGP table version is 364, main routing table version 364
60 network entries and 60 paths using 11340 bytes of memory
48 BGP path attribute entries using 2496 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
22 BGP AS-PATH entries using 528 bytes of memory
9 BGP extended community entries using 248 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP activity 622/2864 prefixes, 1318/1135 paths, scan interval 15 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
150.1.31.209    4 65031    1451    1533      364    0    0 00:11:29        7
150.1.31.213    4 65031       6      18      364    0    0 00:02:41        1
150.1.32.2      4 65032    1422    1456      364    0    0 00:14:45        3
192.168.3.1     4     3    4639    4839      364    0    0 00:15:00        6
192.168.3.3     4     3    4725    4797      364    0    0 00:15:00       60
```

- Verify the availability of Internet routes in the Internet VRF by using the **show ip bgp vpnv4 vrf *name*** command

```
WG3PE2#show ip bgp vpnv4 vrf Internet
BGP table version is 364, local router ID is 192.168.3.2
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: 3:600 (default for vrf Internet)
*>i128.20.0.0       192.168.3.4              0    100      0 22 i
*>i128.22.0.0       192.168.3.4              0    100      0 22 i
*>i128.26.0.0       192.168.3.4              0    100      0 22 26 i
*>i128.37.0.0       192.168.3.4              0    100      0 20 42 37 i
*>i128.42.0.0       192.168.3.4              0    100      0 20 42 i
*>i128.51.0.0       192.168.3.4              0    100      0 22 26 51 i
*>i128.213.0.0      192.168.3.4              0    100      0 20 213 i
… rest deleted …
```

- Check connectivity between the Internet routers and the customers using the **trace** or **ping** commands

Additional Task: Direct Internet connectivity for all CE-routers

You might want to explore the Internet-in-a-VPN concept and its flexibility further by joining the Internet VPN and the customer VPN, thus giving all CE-routers direct Internet connectivity. Should you wish to do that, perform this simple step:

Step 1 Configure additional route-targets in the customer VRF to import Internet routes and export customer routes with an Internet route-target.

Verification

- Check the presence of Internet routes in the customer VRF with the **show ip route vrf** command

```
WG3PE2#show ip route vrf wg3a
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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is 150.1.31.209 to network 0.0.0.0

B    201.1.1.0 255.255.255.0 [200/0] via 192.168.3.4, 00:00:14
B    202.2.2.0 255.255.255.0 [200/0] via 192.168.3.4, 00:00:14
B    201.2.1.0 255.255.255.0 [200/0] via 192.168.3.4, 00:00:14
R    203.1.2.0 255.255.255.0 [120/1] via 150.1.31.6, 00:00:25, Serial0/0.3
     201.1.0.0 255.255.255.255 is subnetted, 2 subnets
B       201.1.0.1 [200/0] via 192.168.3.4, 00:00:14
B       201.1.0.2 [200/0] via 192.168.3.4, 00:00:14
     201.2.0.0 255.255.255.255 is subnetted, 2 subnets
B       201.2.0.2 [200/0] via 192.168.3.4, 00:00:14
B       201.2.0.1 [200/0] via 192.168.3.4, 00:00:14
B    203.1.1.0 255.255.255.0 [20/50] via 150.1.31.209 (wg3a1), 00:17:33
     202.1.0.0 255.255.255.255 is subnetted, 2 subnets
B       202.1.0.2 [200/0] via 192.168.3.4, 00:00:14
B       202.1.0.1 [200/0] via 192.168.3.4, 00:00:14
     202.2.0.0 255.255.255.255 is subnetted, 2 subnets
B       202.2.0.1 [200/0] via 192.168.3.4, 00:00:14
B       202.2.0.2 [200/0] via 192.168.3.4, 00:00:14
B    192.213.11.0 255.255.255.0 [200/0] via 192.168.3.4, 00:00:14
B    192.214.11.0 255.255.255.0 [200/0] via 192.168.3.4, 00:00:15
B    192.51.11.0 255.255.255.0 [200/0] via 192.168.3.4, 00:00:15
B    192.37.11.0 255.255.255.0 [200/0] via 192.168.3.4, 00:00:15
B    192.42.11.0 255.255.255.0 [200/0] via 192.168.3.4, 00:00:15
B    192.20.11.0 255.255.255.0 [200/0] via 192.168.3.4, 00:00:15
B    192.22.11.0 255.255.255.0 [200/0] via 192.168.3.4, 00:00:15
B    192.26.11.0 255.255.255.0 [200/0] via 192.168.3.4, 00:00:15
```


- Check the presence of customer routes in the Internet VRF by using the **show ip route vrf** command

```
WG3PE2#show ip route vrf Internet
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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

B    201.1.1.0 255.255.255.0 [200/0] via 192.168.3.4, 00:07:17
B    202.2.2.0 255.255.255.0 [200/0] via 192.168.3.4, 00:07:17
B    201.2.1.0 255.255.255.0 [200/0] via 192.168.3.4, 00:07:17
B    203.1.2.0 255.255.255.0
           [20/1] via 150.1.31.6 (wg3a), 00:00:00, Serial0/0.3
     201.1.0.0 255.255.255.255 is subnetted, 2 subnets
B       201.1.0.1 [200/0] via 192.168.3.4, 00:07:17
B       201.1.0.2 [200/0] via 192.168.3.4, 00:07:17
     201.2.0.0 255.255.255.255 is subnetted, 2 subnets
B       201.2.0.2 [200/0] via 192.168.3.4, 00:07:17
B       201.2.0.1 [200/0] via 192.168.3.4, 00:07:17
… rest deleted …
```

E

Initial Laboratory Configuration

Overview

This chapter contains a series of laboratory exercises that will result in the initial laboratory setup as specified in Introduction to Laboratory Exercises. Following these exercises, you will perform initial router configuration of your backbone routers and customer routers, establish the IP addressing, IGP routing, and BGP routing.

These exercises are review exercises that will help you review the IGP and BGP concepts.

It contains the following exercises:

- Initial Core Router Configuration
- Initial Customer Router Configuration
- Basic ISP Setup

Laboratory Exercise E-1: Initial Core Router Configuration

Objective

In this laboratory exercise you will complete the following task:

- Prepare the core routers for further configuration

Note The group of four students working in one workgroup is split into two subgroups. One subgroup configures WGxPE1 and WGxPE2; the other subgroup configures WGxPE3 and WGxPE4. The faster subgroup also configures the WGxP router.

Task: Configure Initial Router Configuration

Step 1 Configure your WGxP and WGxPE1 through WGxPE4 using the parameters in Table 23.

Parameter	Value
host names	Use hostnames as shown in Figure 2 of the Introduction to Laboratory Exercises chapter.
Enable password	cisco
VTY password	cisco
WAN link encapsulation	Frame Relay
WAN link clock rate	64 kbps (configured on the Frame Relay switch)
Core WAN subnets, P and PE router loopback IP addresses	use IP network 192.168.x.0/24, subnet it as needed
ISP Exchange point subnet [1]	192.168.20.x, subnet mask 255.255.255.0
Client ISP subnet [2]	192.168.21.x, subnet mask 255.255.255.0
Network Management subnet [3]	192.168.22.x, subnet mask 255.255.255.0

Table 23: Initial core router parameters

Step 2 You should also configure **ip host** mappings to ease telnet hopping between core routers.

Step 3 Configure point-to-point Frame Relay subinterfaces on the Frame Relay links. The DLCI values for all Frame Relay virtual circuits are shown in Table 24.

---

[1] Router Good has IP address 192.168.20.20 and router Cheap has IP address 192.168.20.22. They are shared by all workgroups.

[2] Router Client has IP address 192.168.21.99 and is shared by all workgroups.

[3] All workgroups share the same router as Network Management Station. The address of the NMS is 192.168.22.22.

Source router	Destination router	DLCI
P	PE3	103
P	PE2	102
PE1	PE2	112
PE2	P	120
PE2	PE1	121
PE2	A2	212
PE2	B1	211
PE3	PE4	134
PE3	P	130
PE3	A1	231
PE3	B2	232
PE4	PE3	143

Table 24: Initial Core Frame Relay PVC parameters

Verification

- All core router interfaces should be active (line up, line protocol up)
- You should be able to telnet and ping between adjacent core routers

Laboratory Exercise E-2: Initial Customer Router Configuration

Objective

In this laboratory exercise you will complete the following task:

- Prepare the customer routers for further configuration.

Note The group of four students working in one workgroup is split into two subgroups. One subgroup configures WGxAy routers and corresponding customer A parameters (WAN links, routing protocols) on all core routers; the other subgroup configures WGxBy routers and the parameters needed for customer B on all core routers. The same division of work is used for all subsequent exercises.

Task: Configure Customer Routers

Step 1 Configure your WGxAy and WGxBy routers using the parameters in Table 25.

Parameter	Value
host names	Use hostnames as shown in above (x is the number of your workgroup)
Enable password	Cisco
VTY password	Cisco
WAN link encapsulation	Frame Relay
WAN link clock rate	64 kbps (configured on the Frame Relay switch)

Table 25: Initial customer router parameters

Step 2 Configure point-to-point Frame Relay subinterfaces on the Frame Relay links. The DLCI values for all Frame Relay virtual circuits are shown in Table 26.

Source router	Destination router	DLCI
A1	PE3	213
A2	PE2	221
B1	PE2	211
B2	PE3	223

Table 26: Initial Customer Router Frame Relay PVC parameters

Step 3 Configure the loopback interfaces on customer routers and the WAN subnets between the customer routers and the core routers using the address space from Table 27. Configure two loopback interfaces on each customer router—one with the subnet mask of /32 and the other with the subnet mask of /24. The second loopback interface will simulate the LAN subnet of the customer.

Customer	Address space
A (loopbacks)	20x.1.0.0/17
A (WAN links)	150.1.x1.0/25
B (loopbacks)	20x.2.0.0/17
B (WAN links)	150.1.x2.0/25

Table 27: Customer address space

Verification

- All customer router interfaces should be active (line up, line protocol up)
- You should be able to ping between customer routers and adjacent core routers

Laboratory Exercise E-3: Basic ISP Setup

Objective

In this laboratory exercise you will complete the following tasks:

- Configure your core network as a transit autonomous system
- Connect the four customer routers to the ISP backbone.

Task 1: Configure IS-IS in your backbone

Step 1 Configure IS-IS as the IGP between your core routers. Your area is 49.000x. Alternatively, configure OSPF in your backbone, using only OSPF area 0.

Note Please refer to Configuring Cisco Routers for IS-IS course if you need further information on IS-IS configuration. Please refer to Building Scalable Cisco Networks for more information on OSPF configuration.

Task 2: Configure BGP in your backbone

Step 1 Configure BGP within your backbone. Your AS number is x. Use loopback interfaces for internal BGP peering.

Step 2 Configure PE2, PE3 and P as route-reflectors in different clusters. Configure PE1 as the client of PE2 and PE4 as the client of PE3.

Note Please refer to the Configuring BGP on Cisco Routers course if you need further information on BGP configuration.

Task 3: Configure Customer Routing

Configure routing between the core and the customer routers with the following parameters:

Step 1 Use BGP with A1 and B1; customers use private AS numbers; use 650xy range where x is your workgroup number.

Step 2 Use static routing with A2 and B2. Configure static route toward A2 and B2 on the PE routers; configure default routes toward the Internet on A2 and B2. Redistribute routes toward the customers into BGP.

Task 4: Peering with other Service Providers

Configure peering with other service providers as follows:

Step 1 Use BGP with peering service providers' routers Good (192.168.20.20, AS 20), Cheap (192.168.20.22, AS 22), and Client (192.168.21.99, AS 99).

Task 5: Establishing Network Management Connectivity

Step 1 Enable NMS access to your backbone and your customers by announcing (via redistribution) your routes to the NMS router via RIP version 2.

Verification

- All core routers should see all networks in your workgroup
- All core routers should see all customers' networks
- All core and customer routers should be able to reach all networks announced from the Good, Cheap and Client service providers
- All neighboring autonomous systems should see your and the customers' networks
- Make sure you do not propagate AS paths with private AS numbers
- Use ping and trace to verify connectivity

F

Initial Router Configuration

Overview

This chapter contains the router configurations for the initial laboratory setup described in Introduction to Laboratory Exercises.

It includes the following router configurations:

- Router WGxPE1
- Router WGxPE2
- Router WGxPE3
- Router WGxPE4
- Router WGxP
- Router WGxA1
- Router WgxA2
- Router WGxB1
- Router WGxB2

Note Interface names in the attached router configurations may vary from the interface names you will observe in the actual lab due to minor hardware differences.

Router WGxPE1

```
hostname WGxPE1
!
enable password cisco
!
ip subnet-zero
no ip domain-lookup
ip tcp synwait-time 5
ip host P 192.168.x.5
ip host PE1 192.168.x.1
ip host PE2 192.168.x.2
ip host PE3 192.168.x.3
ip host PE4 192.168.x.4
ip host A1 20x.1.0.1
ip host A2 20x.1.0.2
ip host B1 20x.2.0.1
ip host B2 20x.2.0.2
!
ip cef
!
interface Loopback0
 ip address 192.168.x.1 255.255.255.255
!
interface Ethernet0/0
 description *** Client ***
 ip address 192.168.21.x 255.255.255.0
 no shut
!
interface Serial0/0
 no ip address
 clock rate 64000
 encapsulation frame-relay
 no fair-queue
 no shutdown
!
interface Serial0/0.1 point-to-point
 description *** Link to PE2 ***
 ip address 192.168.x.22 255.255.255.252
 ip router isis
 frame-relay interface-dlci 112
!
router isis
 net 49.000x.0000.0000.0001.00
 passive-interface Loopback0
 passive-interface Ethernet0/0
 passive-interface Ethernet1/0
 passive-interface FastEthernet0/0
 passive-interface FastEthernet1/0
 is-type level-2-only
 metric-style wide
!
router bgp x
 no synchronization
 no auto-summary
 network 192.168.x.1 mask 255.255.255.255
 neighbor 192.168.21.99 remote-as 99
 neighbor 192.168.21.99 remove-private-AS
 neighbor 192.168.x.2 remote-as x
```



```
 neighbor 192.168.x.2 update-source Loopback0
!
ip classless
!
no ip http server
!
line con 0
 logging synchronous
 transport input none
 no login
 privilege level 15
 ip netmask-format decimal
 exec-timeout 0
line vty 0 4
 logging synchronous
 no login
 privilege level 15
 ip netmask-format decimal
!
end
```

Router WGxPE2

```
hostname WGxPE2
!
enable password cisco
!
ip subnet-zero
no ip domain-lookup
ip tcp synwait-time 5
ip host P 192.168.x.5
ip host PE1 192.168.x.1
ip host PE2 192.168.x.2
ip host PE3 192.168.x.3
ip host PE4 192.168.x.4
ip host A1 20x.1.0.1
ip host A2 20x.1.0.2
ip host B1 20x.2.0.1
ip host B2 20x.2.0.2
!
ip cef
!
interface Loopback0
 ip address 192.168.x.2 255.255.255.255
!
interface Ethernet0/0
 description *** NMS **
 ip address 192.168.22.x 255.255.255.0
 no shut
!
interface Serial0/0
 no ip address
 clock rate 64000
 encapsulation frame-relay
 no fair-queue
 no shut
!
interface Serial0/0.1 point-to-point
 description *** Link to P ***
 ip address 192.168.x.18 255.255.255.252
 ip router isis
 frame-relay interface-dlci 120
!
interface Serial0/0.2 point-to-point
 description *** Link to PE1 ***
 ip address 192.168.x.21 255.255.255.252
 ip router isis
 frame-relay interface-dlci 121
!
interface Serial0/0.3 point-to-point
 description *** Link to A2 ***
 ip address 150.1.x1.5 255.255.255.252
 frame-relay interface-dlci 212
!
interface Serial0/0.4 point-to-point
 description *** Link to B1 ***
 ip address 150.1.x2.1 255.255.255.252
 frame-relay interface-dlci 211
!
router isis
```



```
 net 49.000x.0000.0000.0002.00
 passive-interface Loopback0
 passive-interface Ethernet0/0
 passive-interface Ethernet1/0
 passive-interface FastEthernet0/0
 passive-interface FastEthernet1/0
 is-type level-2-only
 metric-style wide
!
router bgp x
 no synchronization
 no auto-summary
 redistribute connected
 redistribute static route-map TAG
 neighbor 150.1.x2.2 remote-as 650x2
 neighbor 192.168.x.1 remote-as x
 neighbor 192.168.x.1 update-source Loopback0
 neighbor 192.168.x.1 route-reflector-client
 neighbor 192.168.x.3 remote-as x
 neighbor 192.168.x.3 update-source Loopback0
 neighbor 192.168.x.5 remote-as x
 neighbor 192.168.x.5 update-source Loopback0
!
ip classless
ip route 20x.1.0.2 255.255.255.255 150.1.x1.6 tag 10
ip route 20x.1.2.0 255.255.255.0 150.1.x1.6 tag 10
!
no ip http server
!
route-map TAG permit 10
 match tag 10
!
line con 0
 logging synchronous
 transport input none
 no login
 privilege level 15
 ip netmask-format decimal
 exec-timeout 0
line vty 0 4
 logging synchronous
 no login
 privilege level 15
 ip netmask-format decimal
!
end
```

Router WGxPE3

```
hostname WGxPE3
!
enable password cisco
!
ip subnet-zero
no ip domain-lookup
ip tcp synwait-time 5
ip host P 192.168.x.5
ip host PE1 192.168.x.1
ip host PE2 192.168.x.2
ip host PE3 192.168.x.3
ip host PE4 192.168.x.4
ip host A1 20x.1.0.1
ip host A2 20x.1.0.2
ip host B1 20x.2.0.1
ip host B2 20x.2.0.2
!
ip cef
!
interface Loopback0
 ip address 192.168.x.3 255.255.255.255
!
interface Serial0/0
 no ip address
 clock rate 64000
 encapsulation frame-relay
 no fair-queue
 no shut
!
interface Serial0/0.1 point-to-point
 description *** Link to P ***
 ip address 192.168.x.13 255.255.255.252
 ip router isis
 frame-relay interface-dlci 130
!
interface Serial0/0.2 point-to-point
 description *** Link to PE4 ***
 ip address 192.168.x.10 255.255.255.252
 ip router isis
 frame-relay interface-dlci 134
!
interface Serial0/0.3 point-to-point
 description *** Link to A1 ***
 ip address 150.1.x1.1 255.255.255.252
 frame-relay interface-dlci 231
!
interface Serial0/0.4 point-to-point
 description *** Link to B2 ***
 ip address 150.1.x2.5 255.255.255.128
 frame-relay interface-dlci 232
!
router isis
 net 49.000x.0000.0000.0003.00
 passive-interface Loopback0
 is-type level-2-only
 metric-style wide
!
```



```
router bgp x
 no synchronization
 no auto-summary
 redistribute connected
 redistribute static route-map TAG
 neighbor 150.1.x1.2 remote-as 650x1
 neighbor 192.168.x.2 remote-as x
 neighbor 192.168.x.2 update-source Loopback0
 neighbor 192.168.x.4 remote-as x
 neighbor 192.168.x.4 update-source Loopback0
 neighbor 192.168.x.4 route-reflector-client
 neighbor 192.168.x.5 remote-as x
 neighbor 192.168.x.5 update-source Loopback0
!
ip classless
ip route 20x.2.0.2 255.255.255.255 150.1.x2.6 tag 10
ip route 20x.2.2.0 255.255.255.0 150.1.x2.6 tag 10
!
no ip http server
!
route-map TAG permit 10
 match tag 10
!
line con 0
 logging synchronous
 transport input none
 no login
 privilege level 15
 ip netmask-format decimal
 exec-timeout 0
line vty 0 4
 logging synchronous
 no login
 privilege level 15
 ip netmask-format decimal
!
end
```

Router WGxPE4

```
hostname WGxPE4
!
enable password cisco
!
ip subnet-zero
no ip domain-lookup
ip tcp synwait-time 5
ip host P 192.168.x.5
ip host PE1 192.168.x.1
ip host PE2 192.168.x.2
ip host PE3 192.168.x.3
ip host PE4 192.168.x.4
ip host A1 20x.1.0.1
ip host A2 20x.1.0.2
ip host B1 20x.2.0.1
ip host B2 20x.2.0.2
!
ip cef
!
interface Loopback0
 ip address 192.168.x.4 255.255.255.255
!
interface Ethernet0/0
 description *** Good and Cheap **
 ip address 192.168.20.x 255.255.255.0
 no shut
!
interface Serial0/0
 bandwidth 64
 no ip address
 clock rate 64000
 encapsulation frame-relay
 no fair-queue
 no shut
!
interface Serial0/0.1 point-to-point
 description *** Link to PE3 **
 ip address 192.168.x.9 255.255.255.252
 ip router isis
 frame-relay interface-dlci 143
!
router isis
 net 49.000x.0000.0000.0004.00
 passive-interface Loopback0
 passive-interface Ethernet0/0
 is-type level-2-only
 metric-style wide
!
router bgp x
 no synchronization
 no auto-summary
 network 192.168.x.4 mask 255.255.255.255
 neighbor 192.168.x.3 remote-as x
 neighbor 192.168.x.3 update-source Loopback0
 neighbor 192.168.20.20 remote-as 20
 neighbor 192.168.20.20 remove-private-AS
 neighbor 192.168.20.22 remote-as 22
```



```
 neighbor 192.168.20.22 remove-private-AS
!
ip classless
no ip http server
!
line con 0
 logging synchronous
 transport input none
 no login
 privilege level 15
 ip netmask-format decimal
 exec-timeout 0
line vty 0 4
 logging synchronous
 no login
 privilege level 15
 ip netmask-format decimal
!
end
```

Router WGxP

```
hostname WGxP
!
enable password cisco
!
ip subnet-zero
no ip domain-lookup
ip tcp synwait-time 5
ip host P 192.168.x.5
ip host PE1 192.168.x.1
ip host PE2 192.168.x.2
ip host PE3 192.168.x.3
ip host PE4 192.168.x.4
ip host A1 20x.1.0.1
ip host A2 20x.1.0.2
ip host B1 20x.2.0.1
ip host B2 20x.2.0.2
!
ip cef
!
interface Loopback0
 ip address 192.168.x.5 255.255.255.255
!
interface Serial0/0
 no ip address
 clock rate 64000
 encapsulation frame-relay
 no fair-queue
 no shut
!
interface Serial0/0.1 point-to-point
 description *** Link to PE2 ***
 ip address 192.168.x.17 255.255.255.252
 ip router isis
 frame-relay interface-dlci 102
!
interface Serial0/0.2 point-to-point
 description *** Link to PE3 ***
 ip address 192.168.x.14 255.255.255.252
 ip router isis
 frame-relay interface-dlci 103
!
router isis
 net 49.000x.0000.0000.0005.00
 passive-interface Loopback0
 is-type level-2-only
 metric-style wide
!
router bgp x
 no synchronization
 no auto-summary
 redistribute connected
 neighbor 192.168.x.2 remote-as x
 neighbor 192.168.x.2 update-source Loopback0
 neighbor 192.168.x.3 remote-as x
 neighbor 192.168.x.3 update-source Loopback0
!
ip classless
```



```
!
no ip http server
!
line con 0
 logging synchronous
 transport input none
 no login
 privilege level 15
 ip netmask-format decimal
 exec-timeout 0
line vty 0 4
 logging synchronous
 no login
 privilege level 15
 ip netmask-format decimal
!
end
```

Router WGxA1

```
hostname WGxA1
!
enable password cisco
!
ip subnet-zero
no ip domain-lookup
ip tcp synwait-time 5
ip host P 192.168.x.5
ip host PE1 192.168.x.1
ip host PE2 192.168.x.2
ip host PE3 192.168.x.3
ip host PE4 192.168.x.4
ip host A1 20x.1.0.1
ip host A2 20x.1.0.2
ip host B1 20x.2.0.1
ip host B2 20x.2.0.2
!
interface Loopback0
 ip address 20x.1.0.1 255.255.255.255
!
interface Loopback1
 ip address 20x.1.1.1 255.255.255.0
!
interface Serial0/0
 no ip address
 clock rate 64000
 encapsulation frame-relay
 no fair-queue
 no shut
!
interface Serial0/0.1 point-to-point
 description *** Link to PE3 ***
 ip address 150.1.x1.2 255.255.255.252
 frame-relay interface-dlci 213
!
router bgp 650x1
 no synchronization
 no auto-summary
 redistribute connected
 neighbor 150.1.x1.1 remote-as x
!
ip classless
!
no ip http server
!
line con 0
 logging synchronous
 transport input none
 no login
 privilege level 15
 ip netmask-format decimal
 exec-timeout 0
line vty 0 4
 logging synchronous
 no login
 privilege level 15
 ip netmask-format decimal
```



```
!
end
```

Router WGxA2

```
hostname WGxA2
!
enable password cisco
!
ip subnet-zero
no ip domain-lookup
ip tcp synwait-time 5
ip host P 192.168.x.5
ip host PE1 192.168.x.1
ip host PE2 192.168.x.2
ip host PE3 192.168.x.3
ip host PE4 192.168.x.4
ip host A1 20x.1.0.1
ip host A2 20x.1.0.2
ip host B1 20x.2.0.1
ip host B2 20x.2.0.2
!
interface Loopback0
 ip address 20x.1.0.2 255.255.255.255
!
interface Loopback1
 ip address 20x.1.2.1 255.255.255.0
!
interface Serial0/0
 no ip address
 clock rate 64000
 encapsulation frame-relay
 no fair-queue
 no shut
!
interface Serial0/0.1 point-to-point
 description *** Link to PE2 ***
 ip address 150.1.x1.6 255.255.255.252
 frame-relay interface-dlci 221
!
ip classless
!
ip route 0.0.0.0 0.0.0.0 150.1.x1.5
!
no ip http server
!
line con 0
 logging synchronous
 transport input none
 no login
 privilege level 15
 ip netmask-format decimal
 exec-timeout 0
line vty 0 4
 logging synchronous
 no login
 privilege level 15
 ip netmask-format decimal
!
end
```


Router WGxB1

```
hostname WGxB1
!
enable password cisco
!
ip subnet-zero
no ip domain-lookup
ip tcp synwait-time 5
ip host P 192.168.x.5
ip host PE1 192.168.x.1
ip host PE2 192.168.x.2
ip host PE3 192.168.x.3
ip host PE4 192.168.x.4
ip host A1 20x.1.0.1
ip host A2 20x.1.0.2
ip host B1 20x.2.0.1
ip host B2 20x.2.0.2
!
interface Loopback0
 ip address 20x.2.0.1 255.255.255.255
!
interface Loopback1
 ip address 20x.2.1.1 255.255.255.0
!
interface Serial0/0
 no ip address
 clock rate 64000
 encapsulation frame-relay
 no fair-queue
 no shut
!
interface Serial0/0.1 point-to-point
 description *** Link to PE2 ***
 ip address 150.1.x2.2 255.255.255.252
 frame-relay interface-dlci 211
!
router bgp 650x2
 no synchronization
 no auto-summary
 redistribute connected
 neighbor 150.1.x2.1 remote-as x
!
ip classless
!
no ip http server
!
line con 0
 logging synchronous
 transport input none
 no login
 privilege level 15
 ip netmask-format decimal
 exec-timeout 0
line vty 0 4
 logging synchronous
 no login
 privilege level 15
 ip netmask-format decimal
```

```
!
end
```


Router WGxB2

```
hostname WGxB2
!
enable password cisco
!
ip subnet-zero
no ip domain-lookup
ip tcp synwait-time 5
ip host P 192.168.x.5
ip host PE1 192.168.x.1
ip host PE2 192.168.x.2
ip host PE3 192.168.x.3
ip host PE4 192.168.x.4
ip host A1 20x.1.0.1
ip host A2 20x.1.0.2
ip host B1 20x.2.0.1
ip host B2 20x.2.0.2
!
interface Loopback0
 ip address 20x.2.0.2 255.255.255.255
!
interface Loopback1
 ip address 20x.2.2.1 255.255.255.0
!
interface Serial0/0
 no ip address
 clock rate 64000
 encapsulation frame-relay
 no fair-queue
 no shut
!
interface Serial0/0.1 point-to-point
 description *** Link to PE3 ***
 ip address 150.1.x2.6 255.255.255.252
 frame-relay interface-dlci 223
!
ip classless
!
ip route 0.0.0.0 0.0.0.0 150.1.x2.5
!
no ip http server
!
line con 0
 logging synchronous
 transport input none
 no login
 privilege level 15
 ip netmask-format decimal
 exec-timeout 0
line vty 0 4
 logging synchronous
 no login
 privilege level 15
 ip netmask-format decimal
!
end
```

