



MPLS Virtual Private Networks (VPNs)

Feature History

Release	Modification
12.0(5)T	This feature was introduced.
12.0(21)ST	This feature was update to support the Cisco 10720 Internet router.
12.0(22)S	This feature was updated to support the Cisco 12000 series Internet Router on the following line cards: the 6E3-SMB and 12E3-SMB line cards, the 6-port channelized T3 (6CT3-SMB) line card, the OC-192c/STM-64c Packet-over-SONET (POS) line card, and the Quad OC-48c STM-16c POS line card.

This document describes the Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) feature and includes the following sections:

- [Feature Overview, page 1](#)
- [Supported Platforms, page 7](#)
- [Supported Standards, MIBs and RFCs, page 8](#)
- [Prerequisites, page 8](#)
- [Configuration Tasks, page 9](#)
- [Configuration Examples, page 11](#)

Feature Overview

The IP virtual private network (VPN) feature for Multiprotocol Label Switching (MPLS) allows a Cisco IOS network to deploy scalable IPv4 Layer 3 VPN backbone services. An IP VPN is the foundation companies use for deploying or administering value-added services including applications and data hosting network commerce, and telephony services to business customers. In private local area networks (LANs), IP-based intranets have fundamentally changed the way companies conduct their business. Companies are moving their business applications to their intranets to extend over a wide area network (WAN). Companies are also embracing the needs of their customers, suppliers, and partners by using extranets (an intranet that encompasses multiple businesses). With extranets, companies reduce business process costs by facilitating supply-chain automation, electronic data interchange (EDI), and other forms of network commerce. To take advantage of this business opportunity, service providers must have an IP VPN infrastructure that delivers private network services to businesses over a public infrastructure.

IP Virtual Private Networks

To effectively implement an IP VPN in your facility, ensure your IP VPN meets the following basic requirements:

Privacy—All IP VPNs offer privacy over a shared (public) network infrastructure. Most companies use an encrypted tunnel. This is only one of several ways to provide network and data privacy.

Scalability—For proper service delivery, VPNs must scale to serve hundreds of thousands of sites and users. Besides being a managed service, VPNs are also a management tool for service providers to control access to services. One example is Closed User Groups for data and voice services.

Flexibility—IP VPNs must handle the any-to-any traffic patterns characteristic of corporate intranets and extranets, in which data no longer flows to and from a central location. VPNs must also have the inherent flexibility to add new sites quickly, connect users over different media, and meet the increasingly sophisticated transport and bandwidth requirements of new intranet applications.

Predictable Performance—Performance needs vary widely requiring different classes of service, but the common requirement is that the performance is predictable. Examples of the ranges of performance requirements include:

- Remote access for mobile users—Require widespread connectivity
- Branch offices—Require a sustained performance level because of the interactive nature of the intranet application in a branch office
- Video conferencing—Require specific performance characteristics

MPLS Virtual Private Networks

MPLS VPNs allow service providers to deploy scalable VPNs and build the foundation to deliver value-added services, including:

Connectionless Service—A significant technical advantage of MPLS VPNs is that they are connectionless. The Internet owes its success to its basic technology, TCP/IP. TCP/IP is built on packet-based, connectionless network paradigm. This means that no prior action is necessary to establish communication between hosts, making it easy for two parties to communicate. To establish privacy in a connectionless IP environment, current VPN solutions impose a connection-oriented, point-to-point overlay on the network. Even if it runs over a connectionless network, a VPN cannot take advantage of the ease of connectivity and multiple services available in connectionless networks. When you create a connectionless VPN, you do not need tunnels and encryption for network privacy, thus eliminating significant complexity.

Centralized Service—Building VPNs in Layer 3 allows delivery of targeted services to a group of users represented by a VPN. A VPN must give service providers more than a mechanism for privately connecting users to intranet services. It must also provide a way to flexibly deliver value-added services to targeted customers. Scalability is critical, because customers want to use services privately in their intranets and extranets. Because MPLS VPNs are seen as private intranets, you may use new IP services such as:

- Multicast
- Quality of service (QoS)
- Telephony support within a VPN
- Centralized services including content and web hosting to a VPN

You can customize several combinations of specialized services for individual customers. For example, a service that combines IP multicast with a low-latency service class enables videoconferencing within an intranet.

Scalability—If you create a VPN using connection-oriented, point-to-point overlays, Frame Relay, or ATM virtual connections (VCs), the VPN's key deficiency is scalability. Specifically, connection-oriented VPNs without fully meshed connections between customer sites, are not optimal. MPLS-based VPNs instead use the peer model and Layer 3 connectionless architecture to leverage a highly scalable VPN solution. The peer model requires a customer site to only a peer with one provider edge (PE) router as opposed to all other CPE or customer edge (CE) routers that are members of the VPN. The connectionless architecture allows the creation of VPNs in Layer 3, eliminating the need for tunnels or VCs.

Other scalability issues of MPLS VPNs are due to the partitioning of VPN routes between PE routers and the further partitioning of VPN and IGP routes between PE routers and provider (P) routers in a core network.

- PE routers must maintain VPN routes for those VPNs who are members.
- P routers do not maintain any VPN routes.

This increases the scalability of the provider's core and ensures that no one device is a scalability bottleneck.

Security—MPLS VPNs offer the same level of security as connection-oriented VPNs. Packets from one VPN do not inadvertently go to another VPN.

Security is provided in the following areas:

- At the edge of a provider network, ensuring packets received from a customer are placed on the correct VPN.
- At the backbone, VPN traffic is kept separate. Malicious spoofing (an attempt to gain access to a PE router) is nearly impossible because the packets received from customers are IP packets. These IP packets must be received on a particular interface or subinterface to be uniquely identified with a VPN label.

Easy to Create—To take full advantage of VPNs, it must be easy for customers to create new VPNs and user communities. Because MPLS VPNs are connectionless, no specific point-to-point connection maps or topologies are required. You can add sites to intranets and extranets and form closed user groups. When you manage VPNs in this manner, it enables membership of any given site in multiple VPNs, maximizing flexibility in building intranets and extranets.

Flexible Addressing—To make a VPN service more accessible, customers of a service provider can design their own addressing plan, independent of addressing plans for other service provider customers. Many customers use private address spaces, as defined in RFC 1918, and do not want to invest the time and expense of converting to public IP addresses to enable intranet connectivity. MPLS VPNs allow customers to continue to use their present address spaces without network address translation (NAT) by providing a public and private view of the address. A NAT is required only if two VPNs with overlapping address spaces want to communicate. This enables customers to use their own unregistered private addresses, and communicate freely across a public IP network.

Integrated Class of Service (CoS) Support—CoS is an important requirement for many IP VPN customers. It provides the ability to address two fundamental VPN requirements:

- Predictable performance and policy implementation
- Support for multiple levels of service in a MPLS VPN

Network traffic is classified and labeled at the edge of the network before traffic is aggregated according to policies defined by subscribers and implemented by the provider and transported across the provider core. Traffic at the edge and core of the network can then be differentiated into different classes by drop probability or delay.

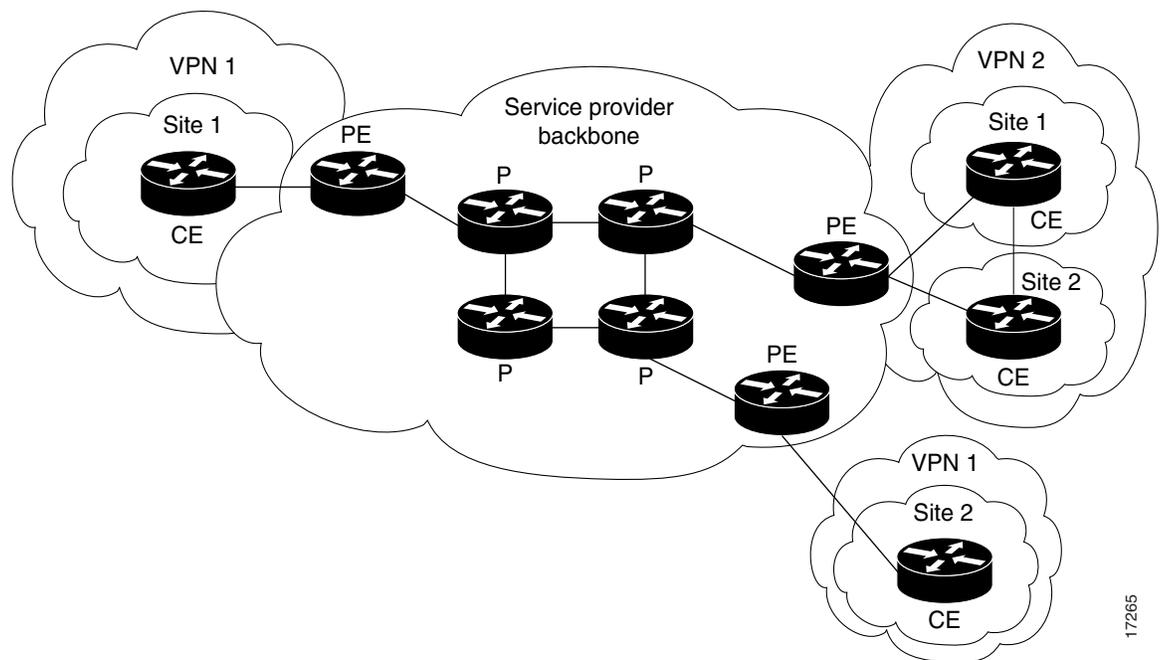
Straightforward Migration—For service providers to quickly deploy VPN services, use a straightforward migration path. MPLS VPNs are unique because you can build them over multiple network architectures, including IP, ATM, Frame Relay, and hybrid networks.

Migration for the end customer is simplified because there is no requirement to support MPLS on the customer edge (CE) router and no modifications are required to a customer's intranet.

For a list of platforms supported by MPLS VPNs, refer to the section Supported Platforms.

See [Figure 1](#) shows an example of a VPN with a service provider (P) backbone network, service provider edge routers (PE), and customer edge routers (CE).

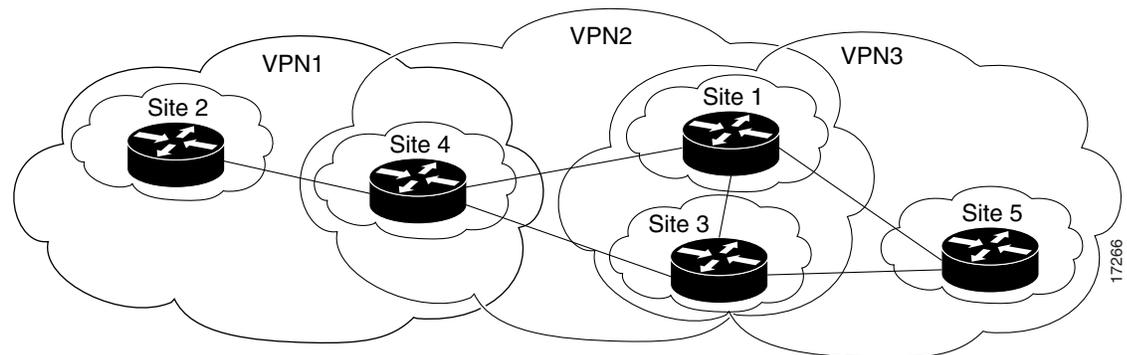
Figure 1 VPNs with a Service Provider Backbone



A VPN contains customer devices attached to the CE routers. These customer devices use VPNs to exchange information between devices. Only the PE routers are aware of the VPNs.

[Figure 2](#) shows five customer sites communicating within three VPNs. The VPNs can communicate with the following sites:

- VPN1—sites 2 and 4
- VPN2—sites 1, 3, and 4
- VPN3—sites 1,3, and 5

Figure 2 Customer Sites within VPNs

VPN Operation

Each VPN is associated with one or more VPN routing/forwarding instances (VRFs). A VRF defines the VPN membership of a customer site attached to a PE router. A VRF consists of an IP routing table, a derived Cisco Express Forwarding (CEF) table, a set of interfaces that use the forwarding table, and a set of rules and routing protocol parameters that control the information that is included into the routing table.

A one-to-one relationship does not necessarily exist between customer sites and VPNs. A given site can be a member of multiple VPNs, as shown in [Figure 2](#). However, a site can only associate with one (and only one) VRF. A customer site's VRF contains all the routes available to the site from the VPNs of which it is a member.

Packet forwarding information is stored in the IP routing table and the CEF table for each VRF. A separate set of routing and CEF tables is maintained for each VRF. These tables prevent information from being forwarded outside a VPN, and also prevent packets that are outside a VPN from being forwarded to a router within the VPN.

VPN Route Target Communities

The distribution of VPN routing information is controlled through the use of VPN route target communities, implemented by border gateway protocol (BGP) extended communities. Distribution of VPN routing information works as follows:

1. When a VPN route learned from a CE router is injected into BGP, a list of VPN route target extended community attributes are associated with it. Typically the list of route target community values is set from an export list of route targets associated with the VRF from which the route was learned.
2. An import list of route target extended communities is associated with each VRF. The import list defines route target extended community attributes a route must have for the route to be imported into the VRF. For example, if the import list for a particular VRF includes route target communities A, B, and C, then any VPN route that carries any of those route target extended communities—A, B, or C—is imported into the VRF.

BGP Distribution of VPN Routing Information

A service provider edge (PE) router can learn an IP prefix from a customer edge (CE) router by static configuration, through a BGP session with the CE router, or through the routing information protocol (RIP) exchange with the CE router. The IP prefix is a member of the IPv4 address family. After it learns the IP prefix, the PE converts it into a VPN-IPv4 prefix by combining it with an 8-byte route distinguisher (RD). The generated prefix is a member of the VPN-IPv4 address family. It serves to uniquely identify the customer address, even if the customer site is using globally nonunique (unregistered private) IP addresses.

The route distinguisher used to generate the VPN-IPv4 prefix is specified by a configuration command associated with the VRF on the PE router.

BGP distributes reachability information for VPN-IPv4 prefixes for each VPN. BGP communication takes place at two levels: within IP domains, known as an autonomous systems (interior BGP or IBGP) and between autonomous systems (external BGP or EBGP). PE-PE or PE-RR (route reflector) sessions are IBGP sessions, and PE-CE sessions are EBGP sessions.

BGP propagates reachability information for VPN-IPv4 prefixes among PE routers by means of the BGP multiprotocol extensions (see RFC 2283, Multiprotocol Extensions for BGP-4) which define support for address families other than IPv4. It does this in a way that ensures the routes for a given VPN are learned only by other members of that VPN, enabling members of the VPN to communicate with each other.

MPLS Forwarding

Based on routing information stored in the VRF IP routing table and VRF CEF table, packets are forwarded to their destination using MPLS.

A PE router binds a label to each customer prefix learned from a CE router and includes the label in the network reachability information for the prefix that it advertises to other PE routers. When a PE router forwards a packet received from a CE router across the provider network it labels the packet with the label learned from the destination PE router. When the destination PE router receives the labeled packet it pops the label and uses it to direct the packet to the correct CE router. Label forwarding across the provider backbone, is based on either dynamic label switching or traffic engineered paths. A customer data packet carries two levels of labels when traversing the backbone:

1. Top label directs the packet to the correct PE router
2. Second label indicates how that PE router should forward the packet to the CE router

Benefits

This section describes the benefits of VPNs in general and MPLS VPNs in particular.

IP VPNs are attractive because they have the following benefits:

- Reduce the cost of connecting branch offices, telecommuters, and mobile users to a corporate intranet, which operate over the public infrastructure of the Internet
- Are more cost-effective than private WANs constructed with leased lines

However, conventional VPNs do not scale well. They are based on creating and maintaining a full mesh of tunnels or permanent virtual circuits among all sites belonging to a particular VPN, using:

- IPsec
- Layer 2 tunneling protocol (L2TP)
- Layer 2 forwarding (L2F) protocol

- Generic routing encapsulation (GRE)
- Frame Relay
- ATM protocols

The overhead required to provision and manage these connection-based schemes cannot be supported in a provider network that must support hundreds or thousands of VPNs, each with tens or hundreds or thousands of sites and thousands or tens of thousands of routes.

MPLS VPNs, which are created in Layer 3, are connectionless, and therefore substantially more scalable and easier to build and manage than conventional VPNs. In addition, you can add value-added services, such as application and data hosting, network commerce, and telephony services to a particular MPLS VPN because the service provider's backbone recognizes each MPLS VPN as a separate, connectionless IP network.

MPLS VPNs offer the following benefits:

- A platform for rapid deployment of additional value-added IP services, including intranets, extranets, voice, multimedia, and network commerce
- Privacy and security equal to that provided by Layer-2 VPNs by limiting the distribution of a VPN's routes to only those routers that are members of the VPN Seamless integration with customer intranets
- Increased scalability over current VPN implementations, with thousands of sites per VPN and hundreds of thousands of VPNs per service provider IP Class of Service (CoS), with support for multiple classes of service and priorities within VPNs, as well as between VPNs
- Management of VPN membership and provisioning of new VPNs for rapid deployment
- Scalable any-to-any connectivity for extended intranets and extranets that encompass multiple businesses

Related Features and Technologies

VPNs may be used with the Class of Service (CoS) feature for MPLS.

Related Documents

- *MPLS Class of Service*
- *Cisco IOS Release 12.0 Network Protocols Command Reference, Part 1*

Supported Platforms

The following is a list of supported router platforms:

- Cisco 7200 series routers
- Cisco 7500 series routers
- Cisco 12000 series routers
- Cisco 10720 Internet router

Determining Platform Support Through Cisco Feature Navigator

Cisco IOS software is packaged in feature sets that support specific platforms. To get updated information regarding platform support for this feature, access Cisco Feature Navigator. Cisco Feature Navigator dynamically updates the list of supported platforms as new platform support is added for the feature.

Cisco Feature Navigator is a web-based tool that enables you to quickly determine which Cisco IOS software images support a specific set of features and which features are supported in a specific Cisco IOS image. You can search by feature or release. Under the release section, you can compare releases side by side to display both the features unique to each software release and the features in common.

To access Cisco Feature Navigator, you must have an account on Cisco.com. If you have forgotten or lost your account information, send a blank e-mail to cco-locksmith@cisco.com. An automatic check will verify that your e-mail address is registered with Cisco.com. If the check is successful, account details with a new random password will be e-mailed to you. Qualified users can establish an account on Cisco.com by following the directions at <http://www.cisco.com/register>.

Cisco Feature Navigator is updated regularly when major Cisco IOS software releases and technology releases occur. For the most current information, go to the Cisco Feature Navigator home page at the following URL:

<http://www.cisco.com/go/fn>

Availability of Cisco IOS Software Images

Platform support for particular Cisco IOS software releases is dependent on the availability of the software images for those platforms. Software images for some platforms may be deferred, delayed, or changed without prior notice. For updated information about platform support and availability of software images for each Cisco IOS software release, refer to the online release notes or, if supported, Cisco Feature Navigator.

Supported Standards, MIBs and RFCs

MIBs

No new or modified MIBs are supported by this feature.

RFCs

- *RFC 1163, A Border Gateway Protocol*
- *RFC 1164, Application of the Border Gateway Protocol in the Internet*
- *RFC 2283, Multiprotocol Extensions for BGP-4*
- *RFC 2547, BGP/MPLS VPNs*

Standards

No new or modified standards are supported by this feature.

Prerequisites

Your network must be running the following Cisco IOS services before you configure VPN operation:

- MPLS in provider backbone routers, or GRE tunnel connectivity among all provider edge (PE) routers
- MPLS with VPN code in provider routers with VPN edge service (PE) routers
- BGP in all routers providing a VPN service
- CEF switching in every MPLS-enabled router
- CoS feature (optional)

Configuration Tasks

Perform the following tasks to configure and verify VPNs:

- [Defining VPNs](#)
- [Configuring BGP PE to PE or PE to CE Routing Sessions](#)
- [Configuring RIP PE to CE Routing Sessions](#)
- [Configuring Static Route PE to CE Routing Sessions](#)
- [Verifying VPN Operation](#)

Defining VPNs

To define VPN routing instances, perform the following steps on the PE router:

	Command	Purpose
Step 1	Router(config)# ip vrf <i>vrf-name</i>	Enter VRF configuration mode and define the VPN routing instance by assigning a VRF name.
Step 2	Router(config-vrf)# rd <i>route-distinguisher</i>	Create routing and forwarding tables.
Step 3	Router(config-vrf)# route-target {import export both} <i>route-target-ext-community</i>	Create a list of import and/or export route target communities for the specified VRF.
Step 4	Router(config-vrf)# import map <i>route-map</i>	(Optional) Associate the specified route map with the VRF.
Step 5	Router(config-if)# ip vrf forwarding <i>vrf-name</i>	Associate a VRF with an interface or subinterface.

Configuring BGP PE to PE or PE to CE Routing Sessions

To configure BGP PE to PE or PE to CE routing sessions in a provider network, perform the following steps on the PE routers:

	Command	Purpose
Step 1	Router(config)# router bgp <i>autonomous-system</i>	Configures the IBGP or EGBP routing process with the autonomous system number passed along to other IBGP or EGBP routers.
Step 2	Router(config-router)# neighbor { <i>ip-address</i> <i>peer-group-name</i> } remote-as <i>number</i>	Specifies a neighbor's IP address or IBGP/EBGP peer group identifying it to the local autonomous system.
Step 3	Router(config-router)# neighbor ip-address activate	Activates the advertisement of the IPv4 address family.

Configuring RIP PE to CE Routing Sessions

To configure RIP PE to CE routing sessions perform the following steps on the PE router:

	Command	Purpose
Step 1	Router(config)# router rip	Enables RIP.
Step 2	Router(config-router)# address-family ipv4 [unicast] vrf <i>vrf-name</i>	Defines RIP parameters for PE to CE routing sessions. The default is Off for auto-summary and synchronization in the VRF address-family submode.
Step 3	Router(config-router)# network prefix	Enables RIP on the PE to CE link.

Configuring Static Route PE to CE Routing Sessions

To configure static route PE to CE routing sessions perform the following steps on the PE router:

	Command	Purpose
Step 1	Router(config)# ip route vrf <i>vrf-name</i>	Defines static route parameters for every PE to CE session.
Step 2	Router(config-router)# address-family ipv4 [unicast] vrf <i>vrf-name</i>	Defines static route parameters for every BGP PE to CE routing session. The default is Off for auto-summary and synchronization in the VRF address-family submode.
Step 3	Router(config-router-af)# redistribute static	Redistributes VRF static routes into the VRF BGP table.
Step 4	Router(config-router-af)# redistribute static connected	Redistributes directly connected networks into the VRF BGP table.
Step 5	Router(config-router-af)# exit-address-family	Exits address family configuration mode.
Step 6	Router(config-router)# end	(Optional) Exits to privileged EXEC mode.

Verifying VPN Operation

To verify VPN operation, perform the following steps:

	Command	Purpose
Step 1	Router# show ip vrf	Displays the set of defined VRFs and interfaces.
Step 2	Router# show ip vrf [{ brief detail interfaces }] <i>vrf-name</i>	Displays information about defined VRFs and associated interfaces.
Step 3	Router# show ip route vrf <i>vrf-name</i>	Displays the IP routing table for a VRF.
Step 4	Router# show ip protocols vrf <i>vrf-name</i>	Displays the routing protocol information for a VRF.
Step 5	Router# show ip cef vrf <i>vrf-name</i>	Displays the CEF forwarding table associated with a VRF.
Step 6	Router# show ip interface <i>interface-number</i>	Displays the VRF table associated with an interface.
Step 7	Router# show ip bgp vpnv4 all [labels]	Displays information about all BGP.
Step 8	Router# show mpls forwarding vrf <i>vrf-name</i> [<i>prefix</i> <i>mask/length</i>] [detail]	Displays label forwarding entries that correspond to VRF routes advertised by this router.

Configuration Examples

This section provides a sample configuration file from a PE router.

```

ip cef distributed          ! CEF switching is pre-requisite for label Switching
frame-relay switching
!
ip vrf vrf1                ! Define VPN Routing instance vrf1
rd 100:1
route-target both 100:1   ! Configure import and export route-targets for vrf1
!
ip vrf vrf2                ! Define VPN Routing instance vrf2
rd 100:2
route-target both 100:2   ! Configure import and export route-targets for vrf2
route-target import 100:1 ! Configure an additional import route-target for vrf2
import map vrf2_import    ! Configure import route-map for vrf2
!
interface lo0
ip address 10.13.0.13 255.255.255.255
!
interface atm9/0/0        ! Backbone link to another Provider router
!
interface atm9/0/0.1 tag-switching
ip unnumbered loopback0
no ip directed-broadcast
mpls atm vpi 2-5
mpls ip

interface atm5/0
no ip address
no ip directed-broadcast
atm clock INTERNAL
no atm ilmi-keepalive

```

```

interface Ethernet1/0
ip address 3.3.3.5 255.255.0.0
no ip directed-broadcast
no ip mroute-cache
no keepalive

interface Ethernet5/0/1
ip vrf forwarding vrf1
ip address 10.20.0.13 255.255.255.0
!
interface hssi 10/1/0
hssi internal-clock
encaps fr
frame-relay intf-type dce
frame-relay lmi-type ansi
!
interface hssi 10/1/0.16 point-to-point
ip vrf forwarding vrf2
ip address 10.20.1.13 255.255.255.0
frame-relay interface-dlci 16
!
router bgp 1
no synchronization
no bgp default ipv4-activate
neighbor 10.15.0.15 remote-as 1
neighbor 10.15.0.15 update-source lo0
!
address-family vpnv4 unicast
neighbor 10.15.0.15 activate
exit-address-family
!
address-family ipv4 unicast vrf vrf1
redistribute static
redistribute connected
neighbor 10.20.0.60 remote-as 65535
neighbor 10.20.0.60 activate
no auto-summary
exit-address-family
!
address-family ipv4 unicast vrf vrf2
redistribute static
redistribute connected
neighbor 10.20.1.11 remote-as 65535
neighbor 10.20.1.11 update-source h10/1/0.16
neighbor 10.20.1.11 activate
no auto-summary
exit-address-family
!
! Define a VRF static route
ip route vrf vrf1 12.0.0.0 255.0.0.0 e5/0/1 10.20.0.60
!
route-map vrf2_import permit 10 ! Define import route-map for vrf2.
...

```