Ethernet Routing Switch

8600

Engineering

IP VPN-Lite for Ethernet Routing Switch 8600 Technical Configuration Guide

Avaya Data Solutions

Document Date: July 2010

Document Number: NN48500-562

Document Version: 1.1

Notices
While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer
Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of this documentation unless such modifications, additions, or deletions were performed by Avaya. End User agree to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer
Avaya is not responsible for the contents or reliability of any linked Web sites referenced within this site or documentation(s) provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty
Avaya provides a limited warranty on this product. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language, as well as information regarding support for this product, while under warranty, is available to Avaya customers and other parties through the Avaya Support Web site: http://www.avaya.com/support
Please note that if you acquired the product from an authorized reseller, the warranty is provided to you by said reseller and not by Avaya.

Licenses
THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO/ ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA AUTHORIZED RESELLER, AND AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE ("AVAYA").

Copyright
Except where expressly stated otherwise, no use should be made of the Documentation(s) and Product(s) provided by Avaya. All content in this documentation(s) and the product(s) provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third Party Components
Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements ("Third Party Components"), which may contain terms that expand or limit rights to use certain portions of the Product ("Third Party Terms"). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: http://support.avaya.com/Copyright.

Trademarks
The trademarks, logos and service marks ("Marks") displayed in this site, the documentation(s) and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the documentation(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party. Avaya is a registered trademark of Avaya Inc. All non-Avaya trademarks are the property of their respective owners.

Downloading documents
For the most current versions of documentation, see the Avaya Support. Web site: http://www.avaya.com/support.

Contact Avaya Support
Avaya provides a telephone number for you to use to report problems or to ask questions about your product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: http:// www.avaya.com/support.

Abstract

This Technical Configuration Guide provides a brief summary for the configuration of IP VPN-Lite for the Avaya Ethernet Routing Switch 8600.

Table of Contents

Figures

Document Updates

July 2010

Conventions

This section describes the text, image, and command conventions used in this document.

Symbols:

Tip – Highlights a configuration or technical tip.

Note – Highlights important information to the reader.

Warning – Highlights important information about an action that may result in equipment damage, configuration or data loss.

Text:

Bold text indicates emphasis.

Italic text in a Courier New font indicates text the user must enter or select in a menu item, button or command:

ERS5520-48T# ***show running-config***

Output examples from Avaya devices are displayed in a Lucinda Console font:

ERS5520-48T# ***show running-config***

```
! Embedded ASCII Configuration Generator Script
! Model = Ethernet Routing Switch 5520-24T-PWR
! Software version = v5.0.0.011
enable
configure terminal
```

1. Overview: IP VPN-Lite

1.1 RFC 4364 MP-BGP Based IP VPNs Overview

IP VPN-Lite, based on RFC 4364, provides the ability to run IP virtual private networks across a normal IP backbone running any type of IGP protocol using MP-BGP – please note that RFC 4364 obsoletes RFC 2547. The RFC framework defines two node types inside the provider backbone: Provider (P) node & Provider Edge (PE) node. Also one node type located at the Customer premises: Customer Edge (CE) node.

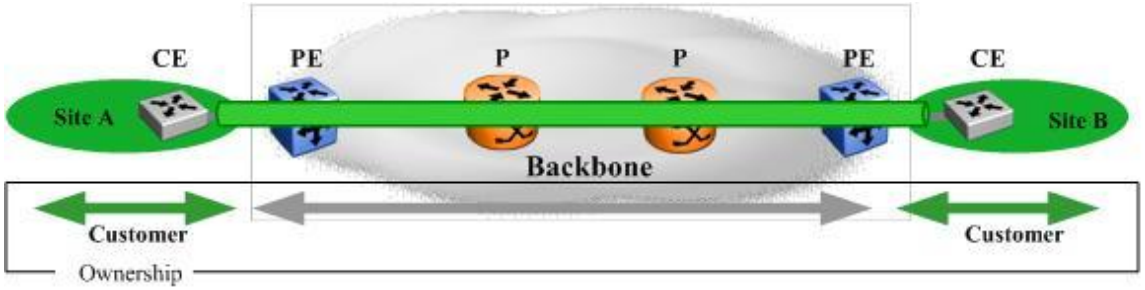

Figure 1: IP VPN-Lite Framework

In a service provider environment the CE nodes are typically owned by the end customers and are L3 IP routers. However, in a typical Enterprise environment, the CE nodes are more likely to be owned by the same Enterprise which owns the backbone. Typically, the CE devices might also be simple L2 switches. We will thus refer to those CE devices as L2 CEs.

The RFC framework defines an architecture where the backbone P nodes are completely unaware of the VPN services they are carrying. This means that they can scale to an unlimited number of IP VPNs and also that no configuration change is required on the P nodes when IP VPN services are added or removed. This is achieved by defining the IP VPNs at the edge of the backbone network in the PE nodes and allowing the PE nodes to form iBGP peering among them selves. This is ensured at the data plane via the use of some form of encapsulation to hide the customer IP forwarding headers from the P nodes.

The PE nodes must be able to exchange IPv4 customer routes with the locally attached CE devices. These routes need to be kept separate from the Backbone IGP routes as well as from other IPv4 routes from different CE devices belonging to different customers. The framework also ensures that PE nodes can handle the same IPv4 route (e.g. using private address space) in use in multiple locally attached CE devices belonging to different customers. This is all achieved via the use of VPN Routing & Forwarding (VRF) Tables which essentially equate to creating virtual routing instances on the PE nodes. Each VRF is thus associated with one single Customer, connecting to one or more CE devices but all belonging to the same Customer. These VRFs will thus exchange routes with the locally connected CE devices using any suitable routing protocol (eBGP, OSPF, RIP, Static Routes) if the CE is a L3 device. If on the other hand the CE is a L2 switch then the Customer routes will be local (direct) routes configured directly on the relevant VRF of the PE node.

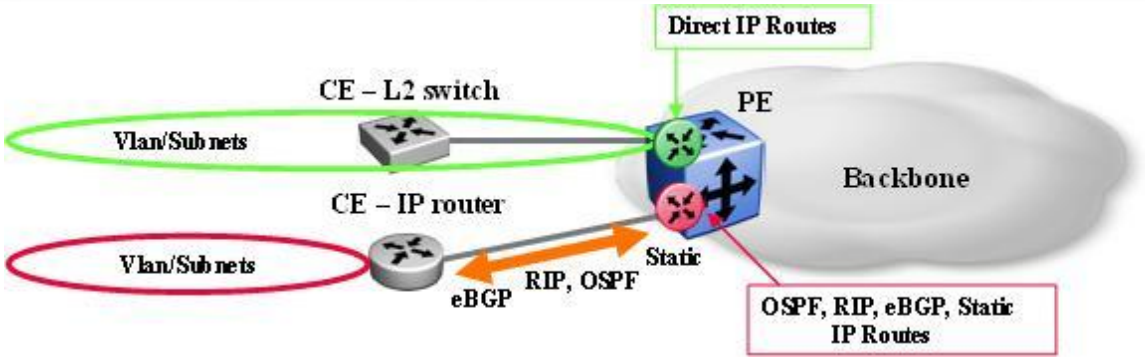

Figure 2: CE to PE Connectivity

The PE nodes must also be able to exchange these local VRF customer IPv4 routes with other remote PE nodes who also own a VRF for the same customer (same IP VPN) while still ensuring that routes from different customers/IP VPNs are kept separate and, likewise, any identical IPv4 routes originating from two different customers can both be advertised and kept separate. This is achieved via the use of iBGP peering between the PE nodes. These iBGP sessions are terminated on a single Circuitless interface on the PE nodes (which belongs to the Backbone Global Routing Table (GRT)). BGP runs over TCP and can therefore be run directly between the PE nodes across the backbone (there is no BGP requirement on the P nodes).

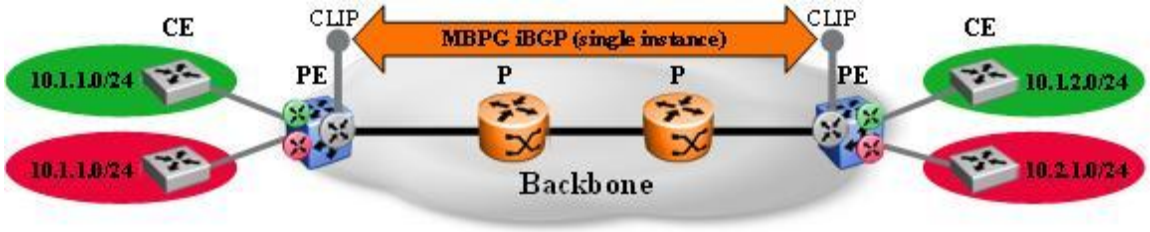

Figure 3: PE to PE Connectivity

A full iBGP peering mesh is required between all PEs. In order to scale to a large number of PE devices, the use of BGP Route reflectors is recommended.

The RFC framework mandates the use on Multi-Protocol (MP) BGP. This requirement stems from the need to (a) make identical IPv4 routes originating from different customers unique for BGP and (b) the need to associate a given customer IPv4 route with a given VPN-id.

The former is achieved via the use of Route Distinguishers (RD) and the use of a new type of address to be advertised in BGP Updates. A new Address Family Identifier (AFI) and Subsequent Address Family Identifier (SAFI) are defined in MP-BGP for VPN-IPv4 routes where a VPN-IPv4 route is the customer IPv4 route pre-pended with a configurable RD. The RD is configured on each and every VRF created on the PE nodes and must be configured in such a way that no other VRF on any other PE in the backbone has the same value. RDs are encoded as part of the Network Layer Reachability Information (NLRI) in the BGP Update messages.

The latter is achieved via the use Route Targets (RT) which are encoded as MP-BGP extended communities. RTs are configured on the PE nodes, within the VRFs, as either import or export or both. RTs constitute the glue which determines whether a customer VPN-IPv4 route being advertised by one PE node is to be accepted by another remote PE node thus forming a logical IP VPN end to end. The export RT is added to VPN-IPv4 BGP Updates originating from the PE which is advertising those customer routes. These routes will only be accepted by remote PE nodes if these have a matching import RT configured on one of their VRFs and furthermore the routes will then only be installed in those VRFs

with matching import RTs. RTs must therefore be configured in such a way as to be unique for each IP VPN. This mechanism also enhances the PE scaling capability since a given PE node will only accept VPN-IPv4 routes for which it has local VRFs belonging to that IP VPN; any other VPN-IPv4 routes will not be accepted.

Since each VRF can be configured with any number of RTs (either as import, export or both) this allows each VRF to be part of any number of overlapping IP VPNs. The use of RT can also be exploited to achieve a number of different IP VPN topologies, from any-to-any (meshed) where all VRFs in the same IP VPN have the same import and export RT, to hub & spoke topologies where the hub nodes use one export RT (configured as import RT on spokes) and a different import RT (configured as export RT on the spokes); topologies with multiple hub sites can also be easily achieved.

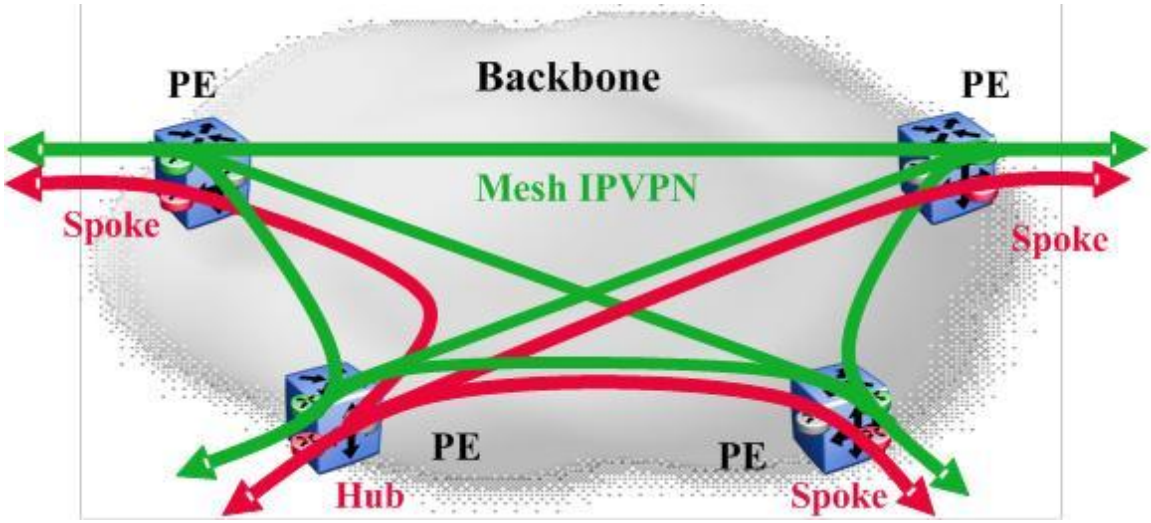

Figure 4: Hub and Spoke

In terms of configuration both the RD and the RT are configured with the same format and are usually configured in the same VRF context on the PE device.

The following diagram illustrates the format used for the RD.

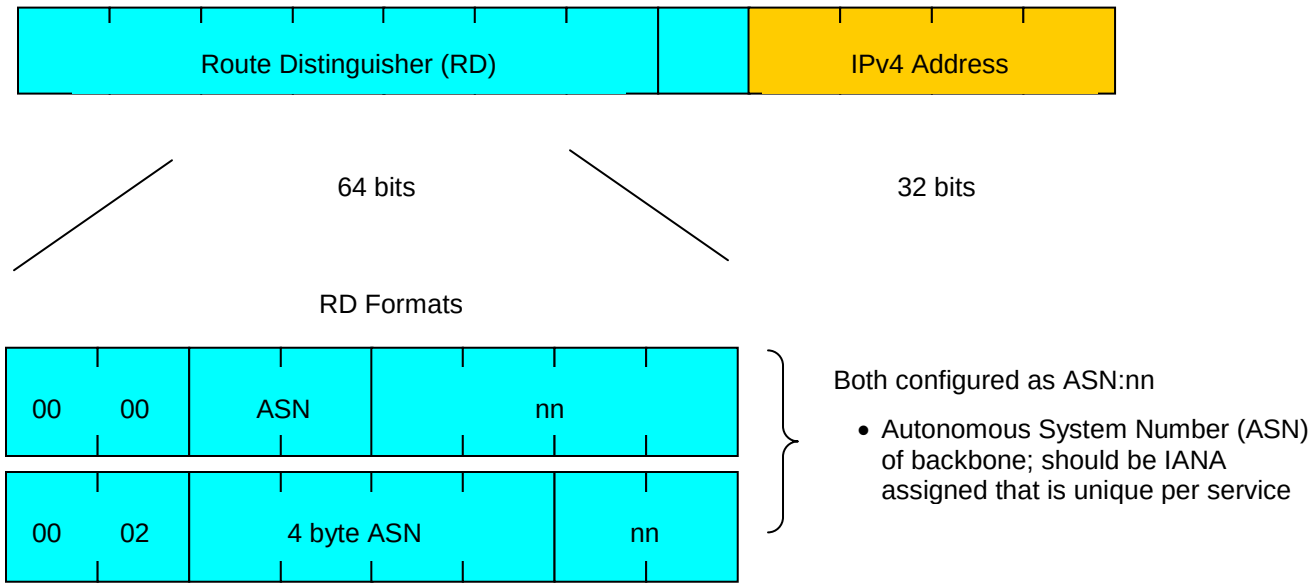

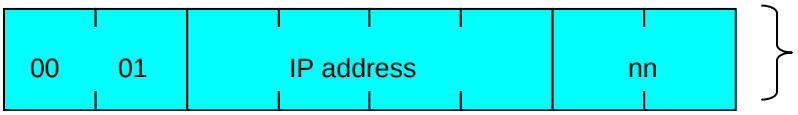

Both configured as IP-address:nn

- Defined if backbone uses private AS number
- Mandatory format on ERS8600

Figure 5: VPN-IPv4 Address

So far we have covered how the RFC framework works at the control plane, i.e. how the customer IP VPN routes are re-advertised over the provider backbone.

In terms of Data Plane packet forwarding across the same backplane, RFC 4364 defines an implementation based on MPLS where the Backbone must be MPLS capable and a full mesh of MPLS Label Switched Paths (LSPs) must already be in place between the PE nodes.

While still leveraging the same identical RFC 4364 framework at the control plane level, Avaya IP VPN-Lite delivers the same IP VPN capabilities over a simple IP routed backbone using simple IP in IP encapsulation with no requirement for MPLS and the complexities involved with running and maintaining an MPLS backbone.

1.2 Delivering RFC4364 IP VPNs over an IP Backbone with Avaya IP VPN-Lite

With IP VPN-Lite a second Circuitless IP address is configured on the PE nodes (again in the Backbone GRT and re-advertised across the Backbone by the IGP). This second Circuitless address will be used to provide address space for the outer header of IP-in-IP encapsulation for all IP VPNs packets terminating to and originating from the PE. This second Circuitless address is therefore ideally configured as a network route (i.e. not as a 32 bit mask host route) with enough address space to accommodate every VRF configured on the PE. A 24 bit mask will thus provide sufficient address space for 252 VRFs. Furthermore, since these networks only need to be routed within the provider backbone and no further, public address space can be used. When this second Circuitless IP address is configured it must also be enabled for IP VPN services.

With Avaya IP VPN-Lite, the RD is now used to convey one extra piece of information over and above its intended use within the RFC4364 framework. In the RFC, the only purpose of the RD is to ensure that identical IPv4 routes from different customers are rendered unique so that BGP can treat them as separate VPN-IPv4 routes. With IP VPN-Lite, the RD is now also used to advertise to remote PE devices what IP address needs to be used as the outer IP-in-IP encapsulation when those remote PE devices need to deliver a customer packet over the IP VPN back to the PE node which owns the destination route to which the packet is addressed.

Therefore when configuring RD for IP VPN-Lite the RD must (a) always be configured as Type 1 format (IPaddress:number) and (b) the IP address configured in the RD must allocate one host IP address defined by the second Circuitless interface for each VRF on the PE. Again, the RD must still be configured to ensure that no other VRF on any other PE has the same RD.

In the following example the second circuitless IP interface is configured as a private address, with a 24 bit mask, where the third octet identifies the PE node-id and the fourth octet (the host portion) defines the VRF on that PE node. The number following the IP address is then simply allocated to uniquely identify the VPN-id.

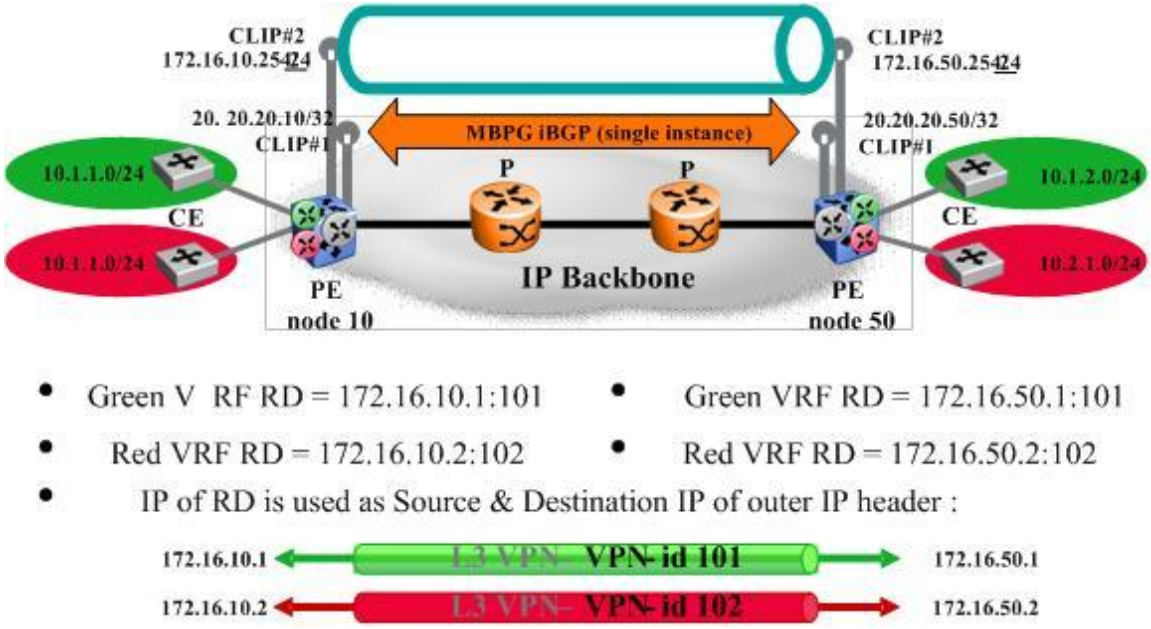

Figure 6: RD Value

IP VPN-Lite can therefore easily be deployed on any enterprise existing IP routed network and automatically leverage the existing backbone architecture in terms of load balancing and sub-second failover resiliency. While MPLS struggles to achieve these goals and only does so by bringing in exponential complexity, Avaya IP VPN-Lite can simply leverage these capabilities from either a pure IP OSPF routed core where ECMP is enabled or a network core designed with Avaya SMLT/RSMLT clustering.

Furthermore, PEs can be just as easily deployed with SMLT clustering towards the CE edge devices thus delivering a very attractive clustered PE solution. This is easily achievable whether the CE is a L2 device (using SMLT Clustering) or an L3 device (where the SMLT cluster needs to be RSMLT enabled).

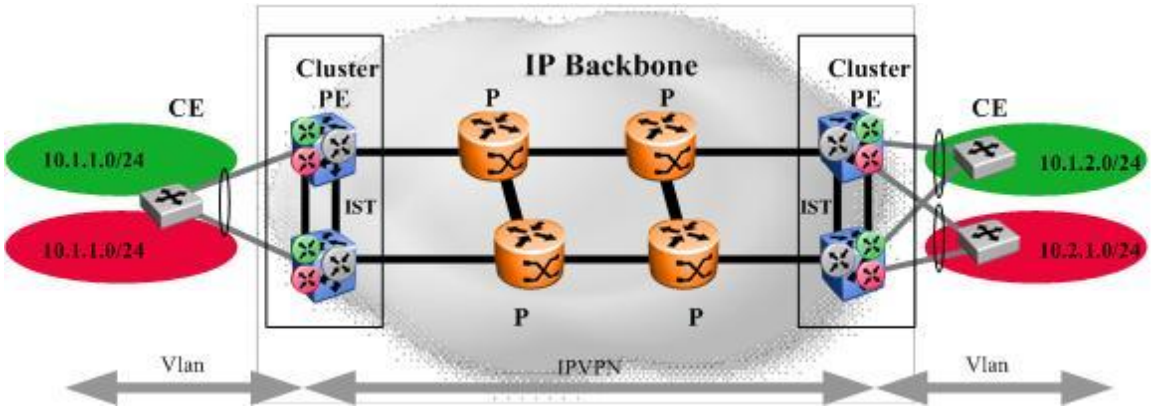

Figure 7: SMLT Deployment

Overall, IP VPN-Lite provides support for the following:

- 256 VPNs per each system
- Filtering support (UNI side)
- Overlapping addresses

- MP-BGP extensions
- BGP route refresh
- BGP route reflection
- Peering to multiple route reflectors
- Route reflection server (NNI side)
- Full mesh and hub and spoke designs
- Extended community Type 0 and 1
- import and export route targets and route distinguishers
- IP-BGP extensions
- IEEE 802.3ad/MLT
- Split MultiLink Trunking (SMLT) and Routed Split MultiLink Trunking
- (RSMLT) for CE connectivity
- ECMP
- VRF-based ping and traceroute
- UNI packet classification (port, VLAN, IP, VRF, and VPN)
- VRF UNI routing protocols (RIP, OSPF, eBGP)

An IP VPN-Lite PE device provides four functions:

- An IGP protocol, such as OSPF, across the core network to connect remote PE devices
- VRFs to provide traffic separation
- MP-BGP to exchange VPN routes and service IP addresses with remote PE devices
- The forwarding plane to encapsulate the customer IP packet into the revise IP header

1.3 Hardware Requirements

IP VPN-Lite works only with R or RS modules and requires a Super Mezzanine daughter card for the 8692SF. This feature can operate in mixed mode but only R and RS modules can be used.

2. Configuration – IP VPN-Lite

The following is a configuration example using IP VPN-Lite between four ERS8600's. These four ERS8600 are deployed across an OSPF backbone as shown in the diagram below. A total of two VRF's will be configured, VRF blue and VRF red, allowing for over-lapping OSPF networks.

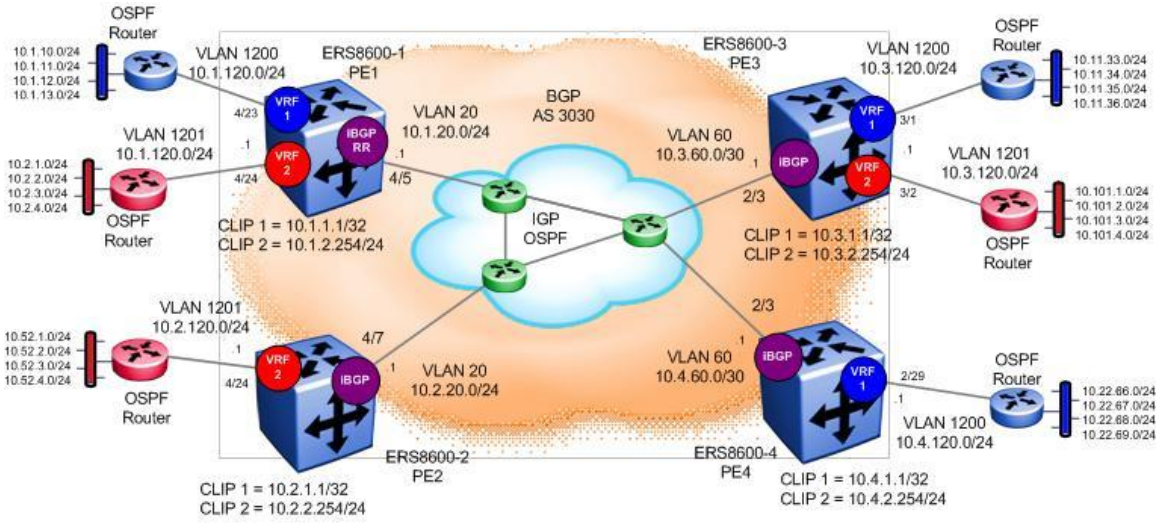

Figure 8: IP VPN-Lite Configuration Example

To ease the iBGP peering requirements, we will configure ERS8600-1 as a route reflector with ERS8600-2, ERS8600-3, and ERS8600-4 as clients. Typically, for added resilience, a second route reflector could be deployed in a cluster – for this example, we will only deploy one route reflector. For more details of this topic and other BGP related items, please refer to the BGP-4 TCG (part number NN48500-538).

In regards to the IP VPN-Lite related items, the following settings will be used.

Item	ERS8600-1	ERS8600-2	ERS8600-3	ERS8600-4
Circuitless IP Addresses (CLIP)				
CLIP 1 - Global	10.1.1.1/32	10.2.1.1/32	10.3.1.1/32	10.4.1.1/32
CLIP 2 - VRF	10.1.2.254/24	10.2.2.254/24	10.3.2.254/24	10.4.2.254/24
BGP				
Type	iBGP with Route Reflection (RR) enabled	iBGP, RR client to 8600-1	iBGP, RR client to 8600-1	iBGP, RR client to 8600-1
AS	3030			
IP VPN-Lite	Enabled for each neighbor			
IP VPN-Lite				

RD – blue (vrf1)	10.1.2.1:1	-	10.3.2.1:1	10.4.2.1:1
RT – blue	3030:1			
RD – red (vrf2)	10.1.2.2:2	10.2.2.2:2	10.3.2.2:2	-
RT – red	3030:2			

Although a separate CLIP address could be configure for each VRF, it is recommended to configure just one CLIP address for all VRF's with a 24-bit mask and configure the RD address from the CLIP address space using the network portion and using the VRF ID as the host portion. In this example for 8600-1, CLIP #2 is configured as 10.1.2.254/24 while the RD value for the blue VRF is configured as 10.1.2.1:1, and the RD value for the red VRF is configured as 10.1.2.2:2

2.1 Configuration Steps

The configuration steps required are as follows:

1. **Base Configuration**
   - Base VLAN configuration to core with IGP protocol
     - In our example, OSPF is used
   - It recommended to use a unique address scheme per zone, i.e. 10.<zone>.<vlan id>.d, i.e. for this example, we will use the following:
     - 8600-1 will use zone 1, 8600-2 will use zone 2, 8600-3 will use zone 3, while 8600-4 will use zone 4
     - For core VLAN's id's, use 2-99
     - For customer VLAN's id's, use 1000-3999
   - CLIP addresses
     - A CLIP address is required for iBGP peering between the PE routers. It is recommended to use a CLIP address with a 32-bit netmask
     - Another CLIP address is required for the VRF's. It is recommended to configure one CLIP address with a 24-bit mask for VRF's and then configure the RD value from this CLIP address space using the same network space and a host address set to the same value as that of the VRF ID. Please note that the network portion must be unique per PE.
2. **BGP Configuration**
   - iBGP peering is required between the PE routers, 8600-1, 8600-2, 8600-3, an d8600-4
   - To help with scaling, route reflectors should be used
   - Between each BGP peer, IP VPN-Lite capability must be enabled
3. **VRF configuration**
   - Add VRF(s) and assigning VLAN, IP address, and IGP protocol or static routing to each VRF
4. **IP VPN-Lite Configuration**
   - Add a IP VPN-Lite Router Identifier (RD) derived from the CLIP address configured in step 1:
     - Make the RD = IPaddress:nn where the IPaddress is derived from the CLIP network address portion and the host is set to the VRF and nn equals the VPN-id
     - For each PE router, the RD must be unique

- Add a IP VPN-Lite Route Target (RT)
  - Create a common BGP extended community route target (RT) address that is used to identify the IP-VRF between the ERS8600 PE switches
  - Is it recommended to use an RT value in the following format: ASN:VPN-id

2.1.1 ERS8600 – Base Configuration

2.1.1.1 Create the Core VLANs

ERS8600-1 Step 1 – Create VLAN 20

```
ERS8600-1:5# config ethernet 4/5 perform-tagging enable
ERS8600-1:5# config vlan 1 port remove 4/5
ERS8600-1:5# config vlan 20 create byport 1
ERS8600-1:5# config vlan 20 ports add 4/5
ERS8600-1:5# config vlan 20 ip create 10.1.20.1/24
ERS8600-1:5# config vlan 20 ip ospf enable
ERS8600-1:5# config ethernet 4/5 stg 1 stp disable
ERS8600-1:5# config ethernet 4/5 state disable
ERS8600-1:5# config ethernet 4/5 state enable
```

ERS8600-2 Step 1 – Create VLAN 20

```
ERS8600-2:5# config ethernet 4/7 perform-tagging enable
ERS8600-2:5# config vlan 1 port remove 4/7
ERS8600-2:5# config vlan 20 create byport 1
ERS8600-2:5# config vlan 20 ports add 4/7
ERS8600-2:5# config vlan 20 ip create 10.2.20.1/24
ERS8600-2:5# config vlan 20 ip ospf enable
ERS8600-2:5# config ethernet 4/7 stg 1 stp disable
ERS8600-2:5# config ethernet 4/7 state disable
ERS8600-2:5# config ethernet 4/7 state enable
```

ERS8600-3 Step 1 – Create VLAN 60

```
ERS8600-3:5# config vlan 60 create byport 1
ERS8600-3:5# config vlan 60 port add 2/3
ERS8600-3:5# config vlan 60 ip create 10.3.60.1/30
ERS8600-3:5# config vlan 60 ip ospf enable
ERS8600-3:5# config ethernet 2/3 stg 1 stp disable
ERS8600-3:5# config ethernet 2/3 state disable
ERS8600-3:5# config ethernet 2/3 state enable
```

ERS8600-4 Step 1 – Create VLAN 60

```
ERS8600-4:5# config vlan 60 create byport 1
ERS8600-4:5# config vlan 60 port add 2/3
```

```
ERS8600-4:5# config vlan 60 ip create 10.4.60.1/30
ERS8600-4:5# config vlan 60 ip ospf enable
ERS8600-4:5# config ethernet 2/3 stg 1 stp disable
ERS8600-4:5# config ethernet 2/3 state disable
ERS8600-4:5# config ethernet 2/3 state enable
```

2.1.1.2 Create Circuitless IP Addresses (CLIP Addresses)

ERS8600-1 Step 1 – Create CLIP 1 and CLIP 2 with OSPF enabled plus enable IPVPN-lite for CLIP 2 only

```
ERS8600-1:5# config ip circuitless-ip-int 1 create 10.1.1.1/32
ERS8600-1:5# config ip circuitless-ip-int 1 ospf enable
ERS8600-1:5# config ip circuitless-ip-int 2 create 10.1.2.254/24
ERS8600-1:5# config ip circuitless-ip-int 2 ospf enable
ERS8600-1:5# config ip circuitless-ip-int 2 ipvpn-lite-capability enable
```

ERS8600-2 Step 1 – Create CLIP 1 and CLIP 2 with OSPF enabled plus enable IPVPN-lite for CLIP 2 only

```
ERS8600-2:5# config ip circuitless-ip-int 1 create 10.2.1.1/32
ERS8600-2:5# config ip circuitless-ip-int 1 ospf enable
ERS8600-2:5# config ip circuitless-ip-int 2 create 10.2.2.254/24
ERS8600-2:5# config ip circuitless-ip-int 2 ospf enable
ERS8600-2:5# config ip circuitless-ip-int 2 ipvpn-lite-capability enable
```

ERS8600-3 Step 1 – Create CLIP 1 and CLIP 2 with OSPF enabled plus enable IPVPN-lite for CLIP 2 only

```
ERS8600-3:5# config ip circuitless-ip-int 1 create 10.3.1.1/32
ERS8600-3:5# config ip circuitless-ip-int 1 ospf enable
ERS8600-3:5# config ip circuitless-ip-int 2 create 10.3.2.254/24
ERS8600-3:5# config ip circuitless-ip-int 2 ospf enable
ERS8600-3:5# config ip circuitless-ip-int 2 ipvpn-lite-capability enable
```

ERS8600-4 Step 1 – Create CLIP 1 and CLIP 2 with OSPF enabled plus enable IPVPN-lite for CLIP 2 only

```
ERS8600-4:5# config ip circuitless-ip-int 1 create 10.4.1.1/32
ERS8600-4:5# config ip circuitless-ip-int 1 ospf enable
ERS8600-4:5# config ip circuitless-ip-int 2 create 10.4.2.254/24
ERS8600-4:5# config ip circuitless-ip-int 2 ospf enable
ERS8600-4:5# config ip circuitless-ip-int 2 ipvpn-lite-capability enable
```

2.1.1.3 Enable OSPF

ERS8600-1 Step 1 – Enable OSPF globally and set OSPF router-id to CLIP 1 address

```
ERS8600-1:5# config ip ospf admin-state enable
ERS8600-1:5# config ip ospf router-id 10.1.1.1
ERS8600-1:5# config ip ospf enable
```

ERS8600-2 Step 2 – Enable OSPF globally and set OSPF router-id to CLIP 1 address

```
ERS8600-2:5# config ip ospf admin-state enable
ERS8600-2:5# config ip ospf router-id 10.2.1.1
ERS8600-2:5# config ip ospf enable
```

ERS8600-3 Step 1 – Enable OSPF globally and set OSPF router-id to CLIP 1 address

```
ERS8600-3:5# config ip ospf admin-state enable
ERS8600-3:5# config ip ospf router-id 10.3.1.1
ERS8600-3:5# config ip ospf enable
```

ERS8600-4 Step 1 – Enable OSPF globally and set OSPF router-id to CLIP 1 address

```
ERS8600-4:5# config ip ospf admin-state enable
ERS8600-4:5# config ip ospf router-id 10.4.1.1
ERS8600-4:5# config ip ospf enable
```

2.1.2 BGP Configuration with IP VPN-Lite Capability Enabled

2.1.2.1 Enable BGP Globally and Add Neighbors

ERS8600-1 Step 1 – Enable BGP globally

```
ERS8600-1:5# config ip bgp auto-summary disable
ERS8600-1:5# config ip bgp synchronization disable
ERS8600-1:5# config ip bgp local-as 3030
ERS8600-1:5# config ip bgp enable
```

ERS8600-2 Step 1 – Enable BGP globally

```
ERS8600-2:5# config ip bgp auto-summary disable
ERS8600-2:5# config ip bgp synchronization disable
ERS8600-2:5# config ip bgp local-as 3030
ERS8600-2:5# config ip bgp enable
```

ERS8600-3 Step 1 – Enable BGP globally

```
ERS8600-3:5# config ip bgp auto-summary disable
ERS8600-3:5# config ip bgp synchronization disable
ERS8600-3:5# config ip bgp local-as 3030
ERS8600-3:5# config ip bgp enable
```

ERS8600-4 Step 1 – Enable BGP globally

```
ERS8600-4:5# config ip bgp auto-summary disable
ERS8600-4:5# config ip bgp synchronization disable
ERS8600-4:5# config ip bgp local-as 3030
ERS8600-4:5# config ip bgp enable
```

ERS8600-1 Step 2 – Add BGP neighbors using the CLIP address , enable BGP route reflector for each neighbor and enable IPVPN-lite

```
ERS8600-1:5# config ip bgp neighbor 10.2.1.1 create
ERS8600-1:5# config ip bgp neighbor 10.2.1.1 remote-as 3030
ERS8600-1:5# config ip bgp neighbor 10.2.1.1 route-reflector-client enable
ERS8600-1:5# config ip bgp neighbor 10.2.1.1 update-source-interface 10.1.1.1 add
ERS8600-1:5# config ip bgp neighbor 10.2.1.1 address-family vpnv4 enable
ERS8600-1:5# config ip bgp neighbor 10.2.1.1 ipvpn-lite-capability enable
ERS8600-1:5# config ip bgp neighbor 10.2.1.1 admin-state enable
ERS8600-1:5# config ip bgp neighbor 10.3.1.1 create
ERS8600-1:5# config ip bgp neighbor 10.3.1.1 remote-as 3030
ERS8600-1:5# config ip bgp neighbor 10.3.1.1 route-reflector-client enable
ERS8600-1:5# config ip bgp neighbor 10.3.1.1 update-source-interface 10.1.1.1 add
ERS8600-1:5# config ip bgp neighbor 10.3.1.1 address-family vpnv4 enable
ERS8600-1:5# config ip bgp neighbor 10.3.1.1 ipvpn-lite-capability enable
ERS8600-1:5# config ip bgp neighbor 10.3.1.1 admin-state enable
ERS8600-1:5# config ip bgp neighbor 10.4.1.1 create
ERS8600-1:5# config ip bgp neighbor 10.4.1.1 remote-as 3030
ERS8600-1:5# config ip bgp neighbor 10.4.1.1 route-reflector-client enable
ERS8600-1:5# config ip bgp neighbor 10.4.1.1 update-source-interface 10.1.1.1 add
ERS8600-1:5# config ip bgp neighbor 10.4.1.1 address-family vpnv4 enable
ERS8600-1:5# config ip bgp neighbor 10.4.1.1 ipvpn-lite-capability enable
ERS8600-1:5# config ip bgp neighbor 10.4.1.1 admin-state enable
```

ERS8600-2 Step 2 – Add BGP neighbors and enable IPVPN-lite. As this is a route reflector client, it only needs to peer with 8600-1.

```
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 create
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 remote-as 3030
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 update-source-interface 10.2.1.1 add
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 address-family vpnv4 enable
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 ipvpn-lite-capability enable
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 admin-state enable
```

ERS8600-3 Step 2 – Add BGP neighbors and enable IPVPN-lite, As this is a route reflector client, it only needs to peer with 8600-1.

```
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 create
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 remote-as 3030
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 update-source-interface 10.3.1.1 add
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 address-family vpnv4 enable
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 ipvpn-lite-capability enable
ERS8600-3:5# config ip bgp neighbor 10.1.1.1 admin-state enable
```

ERS8600-4 Step 2 – Add BGP neighbors and enable IPVPN-lite. As this is a route reflector client, it only needs to peer with 8600-1.

```
ERS8600-4:5# config ip bgp neighbor 10.1.1.1 create
ERS8600-4:5# config ip bgp neighbor 10.1.1.1 remote-as 3030
ERS8600-4:5# config ip bgp neighbor 10.1.1.1 update-source-interface 10.4.1.1 add
ERS8600-4:5# config ip bgp neighbor 10.1.1.1 address-family vpnv4 enable
ERS8600-4:5# config ip bgp neighbor 10.1.1.1 ipvpn-lite-capability enable
ERS8600-4:5# config ip bgp neighbor 10.1.1.1 admin-state enable
```

2.1.3 VRF Configuration

2.1.3.1 Create VRF Instances

Please note the VRF instance ID is of local significance to the switch. Enter an ID number from 1 to 255.

ERS8600-1 Step 1 – Create VRF instances

```
ERS8600-1:5# config ip vrf blue create id 1
ERS8600-1:5# config ip vrf blue ospf create
ERS8600-1:5# config ip vrf red create id 2
ERS8600-1:5# config ip vrf red ospf create
```

ERS8600-2 Step 1 – Create VRF instances

```
ERS8600-2:5# config ip vrf red create id 2
ERS8600-2:5# config ip vrf red ospf create
```

ERS8600-3 Step 1 – Create VRF instances

```
ERS8600-3:5# config ip vrf blue create id 1
ERS8600-3:5# config ip vrf blue ospf create
ERS8600-3:5# config ip vrf red create id 2
ERS8600-3:5# config ip vrf red ospf create
```

ERS8600-4 Step 1 – Create VRF instances

```
ERS8600-4:5# config ip vrf blue create id 1
ERS8600-4:5# config ip vrf blue ospf create
```

2.1.3.2 Create VLANs for VRF

ERS8600-1 Step 1 – Create VLAN 1200 and 1201

```
ERS8600-1:5# config vlan 1200 create byport 1 name blue
ERS8600-1:5# config vlan 1200 vrf blue
ERS8600-1:5# config vlan 1200 ports add 4/23
ERS8600-1:5# config vlan 1200 ip create 10.1.120.1/24
ERS8600-1:5# config vlan 1200 ip ospf enable
ERS8600-1:5# config vlan 1201 create byport 1 name red
ERS8600-1:5# config vlan 1201 vrf red
ERS8600-1:5# config vlan 1201 ports add 4/24
ERS8600-1:5# config vlan 1201 ip create 10.1.120.1/24
ERS8600-1:5# config vlan 1201 ip ospf enable
```

ERS8600-2 Step 2 – Create VLAN 1201

```
ERS8600-2:5# config vlan 1201 create byport 1 name red
ERS8600-2:5# config vlan 1201 vrf red
ERS8600-2:5# config vlan 1201 ports add 4/24
ERS8600-2:5# config vlan 1201 ip create 10.2.120.1/24
ERS8600-2:5# config vlan 1201 ip ospf enable
```

ERS8600-3 Step 1 – Create VLAN 1200 and 1201

```
ERS8600-3:5# config vlan 1200 create byport 1 name blue
```

```
ERS8600-3:5# config vlan 1200 vrf blue
ERS8600-3:5# config vlan 1200 ports add 3/1
ERS8600-3:5# config vlan 1200 ip create 10.3.120.1/24
ERS8600-3:5# config vlan 1200 ip ospf enable
ERS8600-3:5# config vlan 1201 create byport 1 name red
ERS8600-3:5# config vlan 1201 vrf red
ERS8600-3:5# config vlan 1201 ports add 3/2
ERS8600-3:5# config vlan 1201 ip create 10.3.120.1/24
ERS8600-3:5# config vlan 1201 ip ospf enable
```

ERS8600-4 Step 1 – Create VLAN 1200

```
ERS8600-4:5# config vlan 1200 create byport 1 name blue
ERS8600-4:5# config vlan 1200 vrf blue
ERS8600-4:5# config vlan 1200 ports add 3/1
ERS8600-4:5# config vlan 1200 ip create 10.4.120.1/24
ERS8600-4:5# config vlan 1200 ip ospf enable
```

ERS8600-1 Step 2 – Enable OSPF globally for VRF instances

```
ERS8600-1:5# config ip vrf blue ospf enable
ERS8600-1:5# config ip vrf red ospf enable
```

ERS8600-2 Step 2 – Enable OSPF globally for VRF instances

```
ERS8600-2:5# config ip vrf red ospf enable
```

ERS8600-3 Step 2 – Enable OSPF globally for VRF instances

```
ERS8600-3:5# config ip vrf blue ospf enable
ERS8600-3:5# config ip vrf red ospf enable
```

ERS8600-4 Step 2 – Enable OSPF globally for VRF instances

```
ERS8600-4:5# config ip vrf blue ospf enable
```

ERS8600-1 Step 3 – Disable BGP auto-summarization and synchronization

```
ERS8600-1:5# config ip vrf blue bgp auto-summary disable
ERS8600-1:5# config ip vrf blue bgp synchronization disable
ERS8600-1:5# config ip vrf red bgp auto-summary disable
ERS8600-1:5# config ip vrf red bgp synchronization disable
```

ERS8600-2 Step 3 – Disable BGP auto-summarization and synchronization

```
ERS8600-2:5# config ip vrf red bgp auto-summary disable
ERS8600-2:5# config ip vrf red bgp synchronization disable
```

ERS8600-3 Step 3 – Disable BGP auto-summarization and synchronization

```
ERS8600-3:5# config ip vrf blue auto-summary disable
ERS8600-3:5# config ip vrf blue synchronization disable
ERS8600-3:5# config ip vrf red bgp auto-summary disable
ERS8600-3:5# config ip vrf red bgp synchronization disable
```

ERS8600-4 Step 3 – Disable BGP auto-summarization and synchronization

```
ERS8600-4:5# config ip vrf blue auto-summary disable
ERS8600-4:5# config ip vrf blue synchronization disable
```

2.1.4 IP VPN-Lite Configuration

2.1.4.1 IP VPN-Lite Configuration

For this configuration example, the Route Distinguisher (RD) network address is derived from the CLIP 2 address configured above with a host address equal to VRF ID while the route target (RT) of AS number:VPN ID

ERS8600-1 Step 1 – For the blue VRF, the Route Distinguisher (RD) is configured as 10.1.2.1:1 while the route target (RT) is configured as 3030:1

```
ERS8600-1:5# config ip vrf blue ipvpn create
ERS8600-1:5# config ip vrf blue ipvpn rd 10.1.2.1:1
ERS8600-1:5# config ip vrf blue ipvpn rt add import 3030:1
ERS8600-1:5# config ip vrf blue ipvpn rt add export 3030:1
ERS8600-1:5# config ip vrf blue ipvpn enable
```

ERS8600-1 Step 2 – For the red VRF, the Route Distinguisher (RD) is configured as 10.1.2.2:2 while the route target (RT) is configured as 3030:2

```
ERS8600-1:5# config ip vrf red ipvpn create
ERS8600-1:5# config ip vrf red ipvpn rd 10.1.2.2:2
ERS8600-1:5# config ip vrf red ipvpn rt add import 3030:2
ERS8600-1:5# config ip vrf red ipvpn rt add export 3030:2
ERS8600-1:5# config ip vrf red ipvpn enable
```

ERS8600-2 Step 1 – For the red VRF, the Route Distinguisher (RD) is configured as 10.2.2.2:2 while the route target (RT) is configured as 3030:2

```
ERS8600-2:5# config ip vrf red ipvpn create
ERS8600-2:5# config ip vrf red ipvpn rd 10.2.2.2:2
ERS8600-2:5# config ip vrf red ipvpn rt add import 3030:2
ERS8600-2:5# config ip vrf red ipvpn rt add export 3030:2
ERS8600-2:5# config ip vrf red ipvpn enable
```

ERS8600-3 Step 1 – For the blue VRF, the Route Distinguisher (RD) is configured as 10.3.2.1:1 while the route target (RT) is configured as 3030:1

```
ERS8600-3:5# config ip vrf blue ipvpn create
ERS8600-3:5# config ip vrf blue ipvpn rd 10.3.2.1:1
ERS8600-3:5# config ip vrf blue ipvpn rt add import 3030:1
ERS8600-3:5# config ip vrf blue ipvpn rt add export 3030:1
ERS8600-3:5# config ip vrf blue ipvpn enable
```

ERS8600-3 Step 2 – For the red VRF, the Route Distinguisher (RD) is configured as 10.3.2.2:2 while the route target (RT) is configured as 3030:2

```
ERS8600-3:5# config ip vrf red ipvpn create
ERS8600-3:5# config ip vrf red ipvpn rd 10.3.2.2:2
ERS8600-3:5# config ip vrf red ipvpn rt add import 3030:2
ERS8600-3:5# config ip vrf red ipvpn rt add export 3030:2
ERS8600-3:5# config ip vrf red ipvpn enable
```

ERS8600-4 Step 1 – For the blue VRF, the Route Distinguisher (RD) is configured as 10.4.2.1:1 while the route target (RT) is configured as 3030:1

```
ERS8600-4:5# config ip vrf blue ipvpn create
ERS8600-4:5# config ip vrf blue ipvpn rd 10.4.2.1:1
ERS8600-4:5# config ip vrf blue ipvpn rt add import 3030:1
ERS8600-4:5# config ip vrf blue ipvpn rt add export 3030:1
ERS8600-4:5# config ip vrf blue ipvpn enable
```

2.2 Verification

2.2.1 BGP and VRF

2.2.1.1 Verify BGP Peers

Please note that ERS8600-1 is configured as a route reflector, it will be peered with all switches used in this example. In the case of ERS8600-2, ERS8600-3, and ERS8600-4, they will only be peered with ERS8600-1.

Step 1 – Assuming we are logged into ERS8600-1, verify the BGP peers are all established

```
ERS8600-1:5# show ip bgp sum
```

Result:

```
================================================================================
                          BGP Summary - GlobalRouter
================================================================================

                                  BGP version - 4
                                     local-as - 3030
                                   Identifier - 10.1.1.1
                                Decision state - Idle
                  The total number of routes is   13

BGP NEIGHBOR INFO :

   NEIGHBOR       RMTAS     STATE     HLDTM  KPALV  HLDCFG KPCFG   WGHT  CONRTY ADVINT
--------------------------------------------------------------------------------
10.2.1.1          3030   Established  180    60     180    60     100    120    5
10.3.1.1          3030   Established  180    60     180    60     100    120    5
10.4.1.1          3030   Established  180    60     180    60     100    120    5

Total bgp neighbors: 3
```

Via ERS8600-1, verify the following information:

Option	Verify
NEIGHBOR STATE	Verify that the state between each iBGP peer is displayed as ***Established*** for iBGP peers ***10.2.1.1***, ***10.3.1.1***, and ***10.4.1.1***. As ERS8600-1 is configured as a route-reflector, only this switch will have a state established to each peer.
RMTAS	Displays as ***3030***, which is the BGP AS number used for this example for the iBGP between the four switches.

local-as Identifier	The local-as should be displayed as ***3030*** while the Identifier should be displayed as ***10.1.1.1*** which could be the CLIP address assigned as the OSPF router-id.

2.2.1.2 Verify MP-BGP RD and RT

Assuming we are logged into ERS8600-1, use the following command to verify the BGP extended community attribute to each BGP peer.

Step 1 –Verify the BGP extended community attribute, in this example to iBGP peer ERS8600-3

```
ERS8600-1:5# show ip bgp neighbor route-vpnv4 10.3.1.1 ext-community
or, for each individual VRF, enter the following commands:
show ip bgp neighbor route-vpnv4 10.3.1.1 ext-community vrf blue
show ip bgp neighbor route-vpnv4 10.3.1.1 ext-community vrf red
```

Result:

```
================================================================================
                      IPVPN BGP Neighbor Routes - VRF blue
================================================================================

The total number of accepted routes from the neighbor is 1

NETWORK/MASK       PEER REM ADDR   NEXTHOP ADDRESS ORG LOC PREF   STATUS
--------------------------------------------------------------------------------
10.3.120.0/24      10.3.1.1        10.3.2.1        INC 100        Used
        Svc Label: 524290
        RD -Type: IPADDR  Value: 10.3.2.1:1
        RT -Type: IPADDR  Value: 3030:1
        AS_PATH : path-is-empty

================================================================================
                      IPVPN BGP Neighbor Routes - VRF red
================================================================================

The total number of accepted routes from the neighbor is 1

NETWORK/MASK       PEER REM ADDR   NEXTHOP ADDRESS ORG LOC PREF   STATUS
--------------------------------------------------------------------------------
10.3.120.0/24      10.3.1.1        10.3.2.2        INC 100        Used
        Svc Label: 524291
        RD -Type: IPADDR  Value: 10.3.2.2:2
```

```
            RT -Type: IPADDR  Value: 3030:2
            AS_PATH : path-is-empty
```

Via ERS8600-1, verify the following information:

Option	Verify
NETWORK/MASK	For this example, for ERS8600-3, VRF blue and VRF red both are configured with the same over-lapping subnet and should be displayed as***10.3.120.0/24***.
PEER REM ADDR	Displays as ***10.3.1.1***, which is the iBGP peer address of ERS8600-3.
NEXTHOP	The next-hop address for VRF blue should be displayed as ***10.3.2.1*** (CLIP address for VRF blue) while the next-hop for VRF red should be displayed as ***10.3.2.2*** (CLIP address for VRF red)
RD –Type: IPADDR	Should be displayed as ***10.3.2.1:1*** for VRF blue and ***10.3.2.2:2*** for VRF red.
RT –Type: IPADDR	Should be displayed as ***3030:1*** for VRF blue and ***3030:2*** for VRF red.

2.2.1.3 Verify VRF Routes

Assuming we are logged into ERS8600-1 and wish to verify the VRF, BGP, and neighbor routes.

Step 1 –Verify VRF blue routes by using the following command

```
ERS8600-1:5# show ip route info vrf blue
```

Result:

```
================================================================================
                              IP Route - VRF blue
================================================================================
                                              NH          INTER
DST              MASK            NEXT         VRF    COST FACE   PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.1.120.0      255.255.255.0   10.1.120.1      -     1    1200   LOC  0   DB   0
10.1.10.0       255.255.255.0   10.1.120.2     blue   20   1200   OSPF 0   IB   25
10.1.11.0       255.255.255.0   10.1.120.2     blue   20   1200   OSPF 0   IB   25
10.1.12.0       255.255.255.0   10.1.120.2     blue   20   1200   OSPF 0   IB   25
10.1.13.0       255.255.255.0   10.1.120.2     blue   20   1200   OSPF 0   IB   25
10.11.11.0      255.255.255.0   10.3.60.2      Glob~  0    20   BGP  0   IBV  175
10.11.33.0      255.255.255.0   10.3.60.2      Glob~  0    20   BGP  0   IBV  175
10.11.34.0      255.255.255.0   10.3.60.2      Glob~  0    20   BGP  0   IBV  175
10.11.35.0      255.255.255.0   10.3.60.2      Glob~  0    20   BGP  0   IBV  175
10.11.36.0      255.255.255.0   10.3.60.2      Glob~  0    20   BGP  0   IBV  175
10.22.22.0      255.255.255.0   10.4.60.2      Glob~  0    20   BGP  0   IBV  175
```

```
10.22.66.0      255.255.255.0   10.4.60.2       Glob~ 0    20  BGP  0    IBV  175
10.22.67.0      255.255.255.0   10.4.60.2       Glob~ 0    20  BGP  0    IBV  175
10.22.68.0      255.255.255.0   10.4.60.2       Glob~ 0    20  BGP  0    IBV  175
10.22.69.0      255.255.255.0   10.4.60.2       Glob~ 0    20  BGP  0    IBV  175

15 out of 15 Total Num of Route Entries, 15 Total Num of Dest Networks displayed.
--------------------------------------------------------------------------------
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route
```

Step 2 –Verify VPN blue BGP routes by using the following command

```
ERS8600-1:5# show ip bgp route vrf blue
```

Result:

```
================================================================================
                            BGP Routes - VRF blue
================================================================================
The total number of routes is 10

NETWORK/MASK       PEER REM ADDR   NEXTHOP ADDRESS ORG LOC PREF
--------------------------------------------------------------------------------
10.11.11.0/24      10.3.1.1        10.3.2.1        INC 100
        AS_PATH: path-is-empty
10.11.33.0/24      10.3.1.1        10.3.2.1        INC 100
        AS_PATH: path-is-empty
10.11.34.0/24      10.3.1.1        10.3.2.1        INC 100
        AS_PATH: path-is-empty
10.11.35.0/24      10.3.1.1        10.3.2.1        INC 100
        AS_PATH: path-is-empty
10.11.36.0/24      10.3.1.1        10.3.2.1        INC 100
        AS_PATH: path-is-empty
10.22.22.0/24      10.4.1.1        10.4.2.1        INC 100
        AS_PATH: path-is-empty
10.22.66.0/24      10.4.1.1        10.4.2.1        INC 100
        AS_PATH: path-is-empty
10.22.67.0/24      10.4.1.1        10.4.2.1        INC 100
        AS_PATH: path-is-empty
10.22.68.0/24      10.4.1.1        10.4.2.1        INC 100
        AS_PATH: path-is-empty
```

8600-3 CLIP 2 Address

8600-4 CLIP 2 Address

```
10.22.69.0/24      10.4.1.1         10.4.2.1         INC 100
        AS_PATH: path-is-empty
```

Step 3 –Verify VPN blue routes from BGP neighbor ERS8600-3

```
ERS8600-1:5# show ip bgp neighbor route-vpnv4 10.3.1.1 vrf blue
```

Result:

```
================================================================================
                      IPVPN BGP Neighbor Routes - VRF blue
================================================================================

The total number of accepted routes from the neighbor is 5

NETWORK/MASK       PEER REM ADDR   NEXTHOP ADDRESS ORG LOC PREF   STATUS
--------------------------------------------------------------------------------
10.11.11.0/24      10.3.1.1         10.3.2.1         INC 100        Used
        AS_PATH : path-is-empty
10.11.33.0/24      10.3.1.1         10.3.2.1         INC 100        Used
        AS_PATH : path-is-empty
10.11.34.0/24      10.3.1.1         10.3.2.1         INC 100        Used
        AS_PATH : path-is-empty
10.11.35.0/24      10.3.1.1         10.3.2.1         INC 100        Used
        AS_PATH : path-is-empty
10.11.36.0/24      10.3.1.1         10.3.2.1         INC 100        Used
        AS_PATH : path-is-empty
```

Step 4 –Verify VPN blue routes from BGP neighbor ERS8600-4

```
ERS8600-1:5# show ip bgp neighbor route-vpnv4 10.4.1.1 vrf blue
```

Result:

```
================================================================================
                      IPVPN BGP Neighbor Routes - VRF blue
================================================================================

The total number of accepted routes from the neighbor is 5

NETWORK/MASK       PEER REM ADDR   NEXTHOP ADDRESS ORG LOC PREF   STATUS
--------------------------------------------------------------------------------
10.22.22.0/24      10.4.1.1         10.4.2.1         INC 100        Used
        AS_PATH : path-is-empty
```

```
10.22.66.0/24      10.4.1.1          10.4.2.1          INC 100        Used
        AS_PATH : path-is-empty
10.22.67.0/24      10.4.1.1          10.4.2.1          INC 100        Used
        AS_PATH : path-is-empty
10.22.68.0/24      10.4.1.1          10.4.2.1          INC 100        Used
        AS_PATH : path-is-empty
10.22.69.0/24      10.4.1.1          10.4.2.1          INC 100        Used
        AS_PATH : path-is-empty
```

2.2.1.4 Packet Capture – BGP Route Update

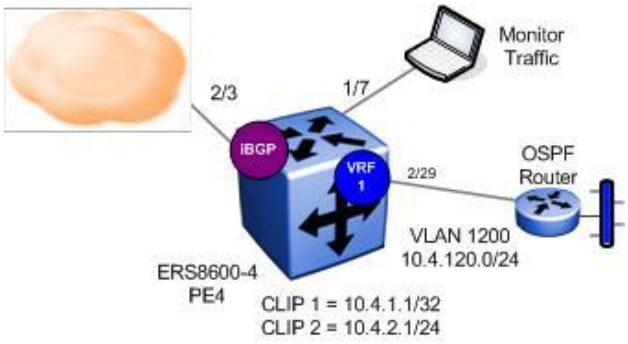

Assuming we have enable port mirroring on ERS8600-4 from in-port 2/3 to a mirrored port 1/7 as shown in the diagram above, shown below is a packet capture displaying a route update from ERS8600-1 to ERS86000-4 with ERS8600-1 advertising routes 10.1.10.0/24 to 10.1.13.0/24 and the local interface from the blue VRF. The command shown was used to enable port mirroring for this example:

- `ERS8600-4:5#` ***config diag mirror-by-port 1 create in-port 2/3 out-port 1/7 mode both enable true***

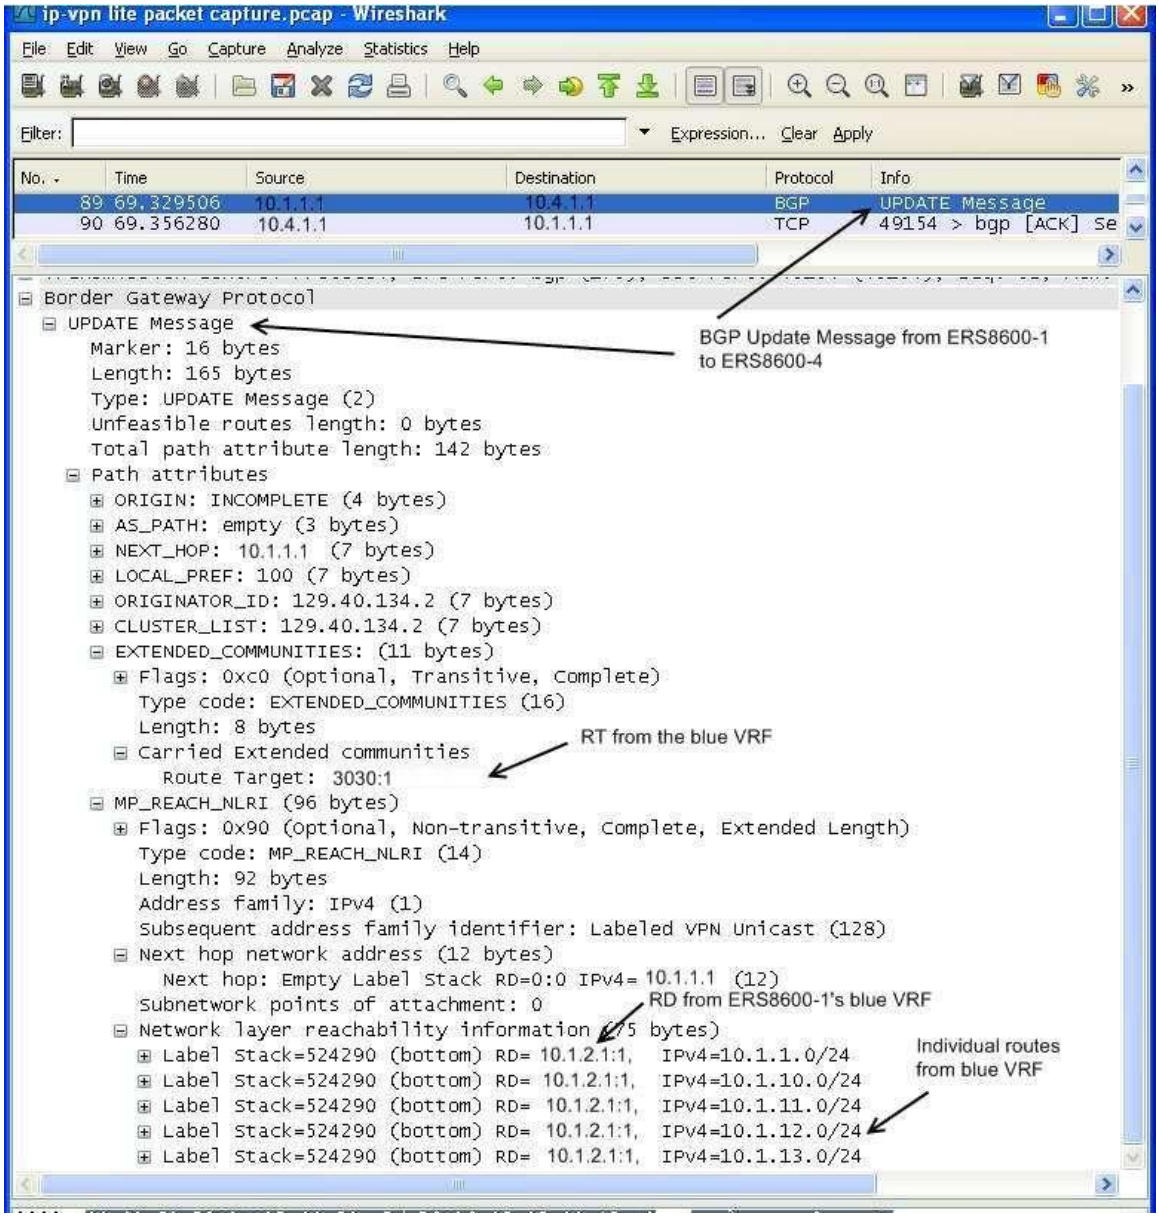
ip-vpn lite packet capture.pcap - Wireshark
File Edit View Go Capture Analyze Statistics Help
Filter: Expression... Clear Apply
No. Time Source Destination Protocol Info
89 69.329506 10.1.1.1 10.4.1.1 BGP UPDATE Message
90 69.356280 10.4.1.1 10.1.1.1 TCP 49154 > bgp [ACK] Se
Border Gateway Protocol
UPDATE Message
Marker: 16 bytes
Length: 165 bytes
Type: UPDATE Message (2)
Unfeasible routes length: 0 bytes
Total path attribute length: 142 bytes
Path attributes
ORIGIN: INCOMPLETE (4 bytes)
AS_PATH: empty (3 bytes)
NEXT_HOP: 10.1.1.1 (7 bytes)
LOCAL_PREF: 100 (7 bytes)
ORIGINATOR_ID: 129.40.134.2 (7 bytes)
CLUSTER_LIST: 129.40.134.2 (7 bytes)
EXTENDED_COMMUNITIES: (11 bytes)
Flags: 0xc0 (Optional, Transitive, Complete)
Type code: EXTENDED_COMMUNITIES (16)
Length: 8 bytes
Carried Extended communities
Route Target: 3030:1
MP_REACH_NLRI (96 bytes)
Flags: 0x90 (Optional, Non-transitive, Complete, Extended Length)
Type code: MP_REACH_NLRI (14)
Length: 92 bytes
Address family: IPv4 (1)
Subsequent address family identifier: Labeled VPN Unicast (128)
Next hop network address (12 bytes)
Next hop: Empty Label Stack RD=0:0 IPv4= 10.1.1.1 (12)
Subnetwork points of attachment: 0
Network layer reachability information (75 bytes)
Label Stack=524290 (bottom) RD= 10.1.2.1:1, IPv4=10.1.1.0/24
Label Stack=524290 (bottom) RD= 10.1.2.1:1, IPv4=10.1.10.0/24
Label Stack=524290 (bottom) RD= 10.1.2.1:1, IPv4=10.1.11.0/24
Label Stack=524290 (bottom) RD= 10.1.2.1:1, IPv4=10.1.12.0/24
Label Stack=524290 (bottom) RD= 10.1.2.1:1, IPv4=10.1.13.0/24
BGP Update Message from ERS8600-1 to ERS8600-4
RT from the blue VRF
RD from ERS8600-1's blue VRF
Individual routes from blue VRF

2.2.1.5 Packet Capture – Ping Trace

Shown below is a packet capture from ERS8600-1 to ERS8600-4 via the blue VRF VLAN 1200 captured on ERS8600-4.

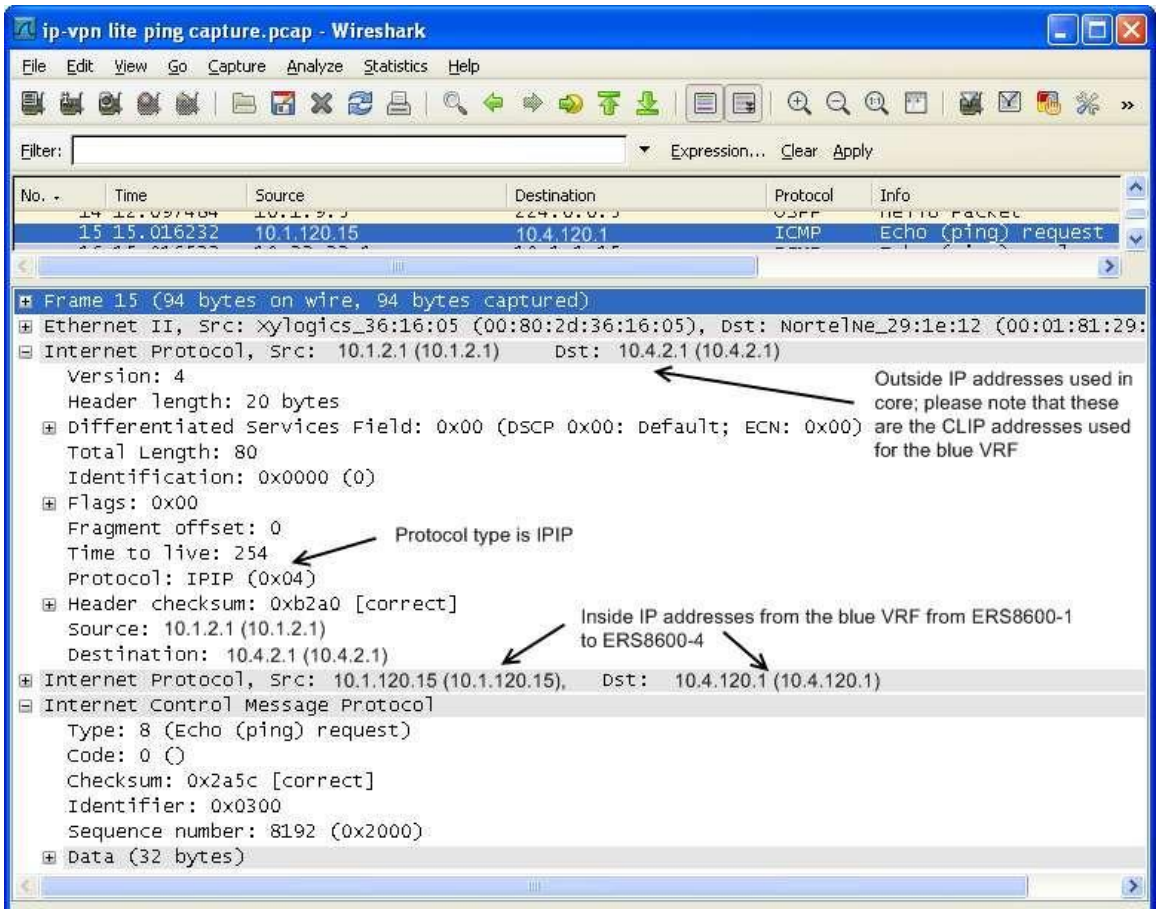

3. Software Baseline

Software revision 5.0 or higher is required.

4. Reference Documentation

Document Title	Publication Number	Description
Configuration – IP VPN	NN46205-520 (323790-A)	
Configuration – MPLS Services	NN46205-519 (323614-1)	
IP-VPN (MPLS) for ERS 8600 Technical Configuration Guide	NN48500-569	
VRF-Lite for Ethernet Routing Switch 8600 Technical Configuration Guide	NN48500-570	
Configuration — BGP Services	NN46205-510	
Border Gateway Protocol (BGP-4) Technical Configuration Guide	NN48500-538	

5. Customer service

Visit the Avaya Web site to access the complete range of services and support that Avaya provides. Go to www.avaya.com or go to one of the pages listed in the following sections.

5.1 Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to www.avaya.com/support.

5.2 Getting product training

Ongoing product training is available. For more information or to register, you can access the Web site at www.avaya.com/support. From this Web site, you can locate the Training contacts link on the left-hand navigation pane.

5.3 Getting help from a distributor or reseller

If you purchased a service contract for your Avaya product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

5.4 Getting technical support from the Avaya Web site

The easiest and most effective way to get technical support for Avaya products is from the Avaya Technical Support Web site at www.avaya.com/support.