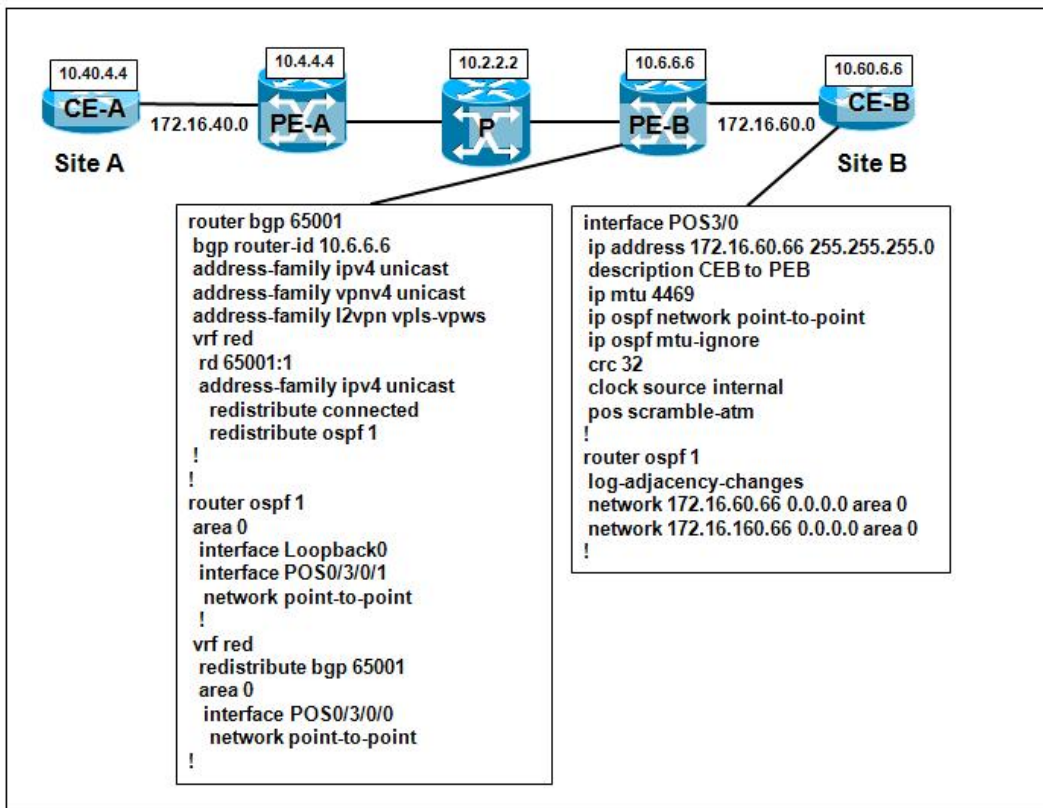


Cisco 642-780

**Maintaining Cisco Service Provider VPNs and MPLS
Networks (MSPVM)**

Version: 4.5

QUESTION NO: 1



Refer to the exhibit.

The customer has an MPLS Layer 3 VPN service CE-A is not able to ping the CE-B loopback address 10.60.6.6 CE-A is able to ping the CE-B network interface address 172.16.60.66. What must be added to the configuration to allow the loopback ping to work?

- A. The PE-B configuration needs a neighbor 10.60.6.6 command added
- B. The PE-B configuration needs interface loopback0 added under vrf red area 0
- C. The CE-B configuration needs to static route added for the PE-B connected interface
- D. The CE-B configuration needs network 10.60.6.6 under router ospf 1

Answer: D

Explanation:

QUESTION NO: 2

Refer to the exhibit.

```
ip vrf VPNA
  rd 23456:100
  route-target both 23456:100
  route-target export 23456:300
  route-target import 23456:301
!
ip vrf VPNB
  rd 23456:200
  route-target both 23456:200
  route-target export 23456:300
  route-target import 23456:301
!
ip vrf VPNX
  rd 23456:300
  route-target both 23456:300
  route-target export 23456:301
!
```

What type of MPLS LAYER 3 VPN configurations is represented?

- A. Simple two-VPN scenario

- B. Overlapping VPNs
- C. Central services VPNs
- D. Extranet VPNs

Answer: C

Explanation: Explanation/Reference:

The Problem: Duplicate Customer Address Ranges

When an SP connects to a wide variety of customers using a Layer 2 WAN service such as Frame Relay or ATM, the SP does not care about the IP addressing and subnets used by those customers. However, in order to migrate those same customers to a Layer 3 WAN service, the SP must learn address ranges from the various customers and then advertise those routes into the SP's network. However, even if the SP wanted to know about all subnets from all its customers, many enterprises use the same address ranges—namely, the private IP network numbers, including the ever-popular network 10.0.0.0.

If you tried to support multiple customers using MPLS unicast IP routing alone, the routers would be confused by the overlapping prefixes, as shown in **Figure 19-11**. In this case, the network shows five of the SP's routers inside a cloud. Three customers (A, B, and C) are shown, with two customer routers connected to the SP's network. All three customers use network 10.0.0.0, with the three customer sites on the right all using subnet 10.3.3.0/24.




Figure 19-11 The Main Challenge with Supporting Layer 3 VPNs

The first and most basic goal for a Layer 3 VPN service is to allow customer A sites to communicate with customer A sites—and only customer A sites. However, the network in **Figure 19-11** fails to meet this goal for several reasons. Because of the overlapping address spaces, several routers would be faced with the dilemma of choosing one customer's route to 10.3.3.0/24 as the best route, and ignoring the route to 10.3.3.0/24 learned from another customer. For example, PE2 would learn about two different 10.3.3.0/24 prefixes. If PE2 chooses one of the two possible routes—for example, if PE2 picked the route to CE-A2 as best—then PE2 could not forward packets to customer B's 10.3.3.0/24 off router CE-B2. Also, a possibly worse effect is that hosts in one customer site may be able to send and receive packets with hosts in another customer's network. Following this same example, hosts in customer B and C sites could forward packets to subnet 10.3.3.0/24, and the routers might forward these packets to customer A's CE-A2 router.

The MPLS VPN Control Plane

The MPLS VPN control plane defines protocols and mechanisms to overcome the problems created by overlapping customer IP address spaces, while adding mechanisms to add more functionality to an MPLS VPN, particularly as compared to traditional Layer 2 WAN services. To understand the mechanics, you need a good understanding of BGP, IGPs, and several new concepts created by both MP-BGP RFCs and MPLS RFCs. In particular, this section introduces and explains the concepts behind three new concepts created for MPLS VPNs:

- VRFs
- Route Distinguishers (RDs)
- Route Targets (RTs)

Each VRF has three main components, as follows:

- An IP routing table (RIB)
- A CEF FIB, populated based on that VRF's RIB
- A separate instance or process of the routing protocol used to exchange routes with the CEs that need to be supported by the VRF

QUESTION NO: 3

In a service provider layer 3 MPLS VPN implementations, what is the minimum number of routing on the PE routers?

- A. Three
- B. Four
- C. Five
- D. Six

Answer: A

Explanation:

QUESTION NO: 4

Refer to the exhibit.

```
RP/0/RSP1/CPU0:VKG-3#ping 10.11.11.11
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.11.11.11, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/2 ms

RP/0/RSP1/CPU0:VKG-3#ping mpls ipv4 10.11.11.11/32

Sending 5, 100-byte MPLS Echos to 10.11.11.11/32,
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/7 ms

RP/0/RSP1/CPU0:VKG-3#ping mpls pseudowire 10.11.11.11 100
Sending 5, 100-byte MPLS Echos to 10.11.11.11 VC: 100,
QQQQ
Success rate is 0 percent (0/5)
```


The commands in the figure were executed from a cisco ASR 9000 series router. The remote end of the ping is a cisco XR12000 series router. Which statement is true?

- A. The XR 12000 has an access list that blacks the MPLS pseudowire ping
- B. The interface connected to the XR 12000 is not running LDP
- C. These results will occur during convergence when MPLS LDP sync is enabled
- D. The MPLS pseudowire ping was not sent

Answer: D


Explanation: Explanation

Explanation/Reference:

http://www.cisco.com/en/US/docs/ios/12_4t/12_4t11/ht_lspng.html

QUESTION NO: 5

Refer to the exhibit.



The output from the debug command is from the Cisco XR12000 series Router, which statement is true?

- A. The output is the result of single iteration of mpls ping executed on the cisco ASR 9000 Series Router'
- B. The output is the result of single iteration of mpls ping executed on the cisco ASR 9000 series Router
- C. The packets did not reach its destination because the control word was not found
- D. The output is the result of single iteration of mpls ping executed on the cisco CRS-1 carrier routing system.


Answer: C

Explanation: Explanation/Reference:

The Control Word

The control word is a 32-bit field that is inserted between the VC label and the transported Layer 2 frame in the case of AToM. It is required for some Layer 2 protocols but optional for others. The control word carries extra information such as protocol control information and a sequence number and does this in a compressed format. This information is needed to correctly and efficiently carry the Layer 2 protocol across the MPLS network. The ingress PE router adds the control word, and the egress PE router strips it off after processing it. Whether the control word is used can be signaled by the PE routers or can be configured. In both cases, the egress PE router knows that it should expect the presence of the control word after the MPLS label stack. **Figure 10-7** shows the generic format of the control word.

Figure 10-7. Generic Format of the Control Word



QUESTION NO: 6

Refer to the exhibit.

```
RP/0/9/CPU0:GSR-1#show cef vrf red 10.22.22.22 detail
10.22.22.22/32, version 1, internal 0x40040001 (0x9cdb88f8) [1], 0x0 (0x0), 0x4100 (0x9d3658f0)
Updated May 26 14:01:44.060
Prefix Len 32, traffic index 0, precedence routine (0)
gateway array (0x9cc5b668) reference count 5, flags 0x80700, source rib (3),
[1 type 1 flags 0x901101 (0x9d6092e4) ext 0x0 (0x0)]
LW-LDI[type=0, refC=0, ptr=0x0, sh-ldi=0x0]
via 192.168.252.3, 5 dependencies, recursive
next hop 192.168.252.3 via 16000/0/21
next hop 10.100.100.2 Gi0/0/1/1 labels imposed {ImplNull 143989}

Load distribution: 0 (refcount 1)

Hash OK Interface Address
0 Y Unknown 16000/0

RP/0/9/CPU0:GSR-1#show bgp vprn4 unicast labels
BGP router identifier 192.168.253.1, local AS number 65000
<snip>
Status codes: s suppressed, d damped, h history, * valid, > best
i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network Next Hop Rcvd Label Local Label
Route Distinguisher: 65000:9 (default for vrf red)
*>i10.1.1.1/32 192.168.252.3 143988 noLabel
*>i10.9.9.0/31 192.168.252.3 143989 noLabel
*> 10.109.109.109/32 0.0.0.0 noLabel 16014
*>i10.170.170.0/24 192.168.252.3 143990 noLabel
*>i10.22.22.22/32 192.168.252.3 143989 noLabel
*>i192.1.99.0/24 192.168.252.3 143991 noLabel
*> 192.168.99.0/24 0.0.0.0 noLabel 16014
```

The commands were executed on a cisco XR12000 PE. Which statement is true?

- A. The outer label for prefix 10.22.22.22 is 143989.
- B. The outer label for prefix 10.22.22.22 is 16000.
- C. The outer label for prefix 10.22.22.22 is implNull.
- D. Prefix 10.22.22.22 has no outer label.

Answer: A

Explanation: The outer label for prefix 10.22.22.22 is 143989.

```

0      Y      Unknown      16000/0
P/0/9/CPU0:GSR-1#show bgp vpnv4 unicast labels
BGP router identifier 192.168.252.1, local AS number 65000
snip>
Status codes: s - suppressed, d - damped, h - history, * - valid, > - best
              i - internal, S - stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network      Next Hop      Rcvd Label      Local Label
Route Distinguisher: 65000:9 (default for vrf red)
> 10.1.1.1/32      192.168.252.3      143988      noLabel
> 10.9.9.0/31      192.168.252.3      143989      noLabel
> 10.100.100.100/32 0.0.0.0      noLabel
> 10.170.170.0/24 192.168.252.3      143990      noLabel
> 10.22.22.22/32 192.168.252.3      143989      noLabel
> 192.168.99.0/24 192.168.252.3      143991      noLabel
> 192.168.99.0/24 192.168.252.3      16014

```

QUESTION NO: 7

You have configured a MPLS-TE tunnel with a predefined explicit path as primary and dynamic path as the backup. The tunnel was designed to carry customer traffic from site A to site B. although the MPLS-TE tunnel is set up. It is not carrying the traffic. Traffic has to go to net work 171.68.0.0/16 connected to site B.

Which three show commands would you use to identify and resolve the issue?

- A. Show mpls traffic-eng-link-management bandwidth to check the bandwidth is available.
- B. Show ip cef network-number to check if tags are imposed
- C. show mpls traffic-eng topology igp-id ospf NETWORK-NUMBER brief to check routing issue
- D. Show mpls traffic-eng autoroute to check auto routing is enabled
- E. Show mpls forwarding-table LABEL-number details to outgoing tags
- F. Show ip cef tunnel1 to check traffic is passing through the tunnel

Answer: B,D,F

Explanation: Explanation/Reference:

show ip cef

To display entries in the Forwarding Information Base (FIB) or to display a summary of the FIB, use the **show ip cef** command in user EXEC or privileged EXEC mode.

show ip cef [*vrf vrf-name*] [*unresolved* [*detail*] | [*detail* | *summary*]]

Specific FIB Entries Based on IP Address Information

show ip cef [*vrf vrf-name*] [*network* [*mask*]] [*longer-prefixes*] [*detail*]

Specific FIB Entries Based on Interface Information

show ip cef [*vrf vrf-name*] [*type number*] [*detail*]

Specific FIB Entries Based on Nonrecursive Routes

show ip cef [*vrf vrf-name*] *non-recursive* [*detail*]

Syntax Description

vrf	(Optional) A Virtual Private Network (VPN) routing/forwarding (VRF) instance.
<i>vrf-name</i>	(Optional) Name assigned to the VRF.
unresolved	(Optional) Displays unresolved FIB entries.
detail	(Optional) Displays detailed FIB entry information.
summary	(Optional) Displays a summary of the FIB.
<i>network</i>	(Optional) Network number for which to display a FIB entry.
<i>mask</i>	(Optional) Network mask to be used with the specified <i>network</i> .
longer-prefixes	(Optional) Displays FIB entries for more specific destinations.
<i>type number</i>	(Optional) Interface type and number for which to display FIB entries.
non-recursive	Displays only nonrecursive routes.

Command Modes

User EXEC

Privileged EXEC

Field	Description
Destination	Multiprotocol Label Switching (MPLS) TE tail-end router ID.
traffic share	A factor based on bandwidth, indicating how much traffic this tunnel should carry, relative to other tunnels, to the same destination. If two tunnels go to a single destination, one with a traffic share of 200 and the other with a traffic share of 100, the first tunnel carries two-thirds of the traffic.
Nexthop	Next-hop router ID of the MPLS-TE tunnel.
absolute metric	Metric with mode absolute for the MPLS-TE tunnel.
relative metric	Metric with mode relative for the MPLS-TE tunnel.

show mpls traffic-eng autoroute

To display tunnels that are announced to the Interior Gateway Protocol (IGP), including information about next hop and destinations, use the **show mpls traffic-eng autoroute** command.

show mpls traffic-eng autoroute [*IP-address*]

Syntax Description

<i>IP-address</i>	(Optional) Tunnel leading to this address.
-------------------	--

Command Default

No default behavior or values

QUESTION NO: 8

Refer to the exhibit.

```
RP/0/9/CPU0:GSR4#show running-config mpls ldp
Mon Jun 7 13:25:27.490 UTC
mpls ldp
router-id 10.10.10.99
discovery targeted-hello accept
graceful-restart
session protection
!
```

Is nonstop forwarding enabled or disabled, and which part of show command can be used to determine this?

- A. Enabled, graceful-restart
- B. Enabled, session protection
- C. Disabled, graceful-restart
- D. Disabled, session protection

Answer: A

Explanation:

QUESTION NO: 9

What two problems are associated with this show MPLS forwarding command output;