



Exam : 642-691

Title : BGP + MPLS Exam

Ver : 09.28.07

QUESTION 1:

Every time a flap occurs on a route, the route receives _____

- A. 750 per-flap penalty points which are user configurable
- B. 1500 per-flap penalty points which are user configurable
- C. 200 per-flap penalty points which are user configurable
- D. 1000 per-flap penalty points which are not user configurable
- E. 2000 per-flap penalty points which are not user configurable

Answer: D

Explanation: Everytime a flap occurs a penalty of 1000 is made and is not configurable. See sections on route dampening in Internet routing architectures guide.

QUESTION 2:

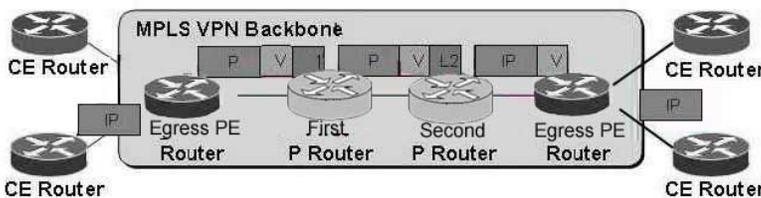
Which component of MPLS architecture uses protocols such as the label distribution (LDP) and tag distribution protocol (TDP) to exchange labels?

- A. Control pane
- B. Data plane
- C. Routing plane
- D. Forwarding plane

Answer: A

QUESTION 3:

Exhibit:



Which statement is true about packet forwarding across an MPLS VPN backbone?

- A. Penultimate hop popping (PHP) on the LDP label is performed by the egress PE router.
- B. Penultimate hop popping (PHP) on the VPN label is performed by the second P router.
- C. Penultimate hop popping (PHP) on the VPN label is performed by the egress PE router.
- D. Penultimate hop popping (PHP) on the LDP label is performed by the second P router.

Answer: D

QUESTION 4:

What is a good definition of Cisco's implementation of an ATM Virtual Path as related to LC-ATM MPLS support?

- A. Designed to establish switch-to-switch connectivity between parts of a public ATM network over a private network.
- B. Designed to establish switch-to-switch connectivity between parts of a private ATM network over a public network.
- C. Designed to establish router-to-router connectivity between parts of a public ATM network over a public network.
- D. Designed to establish router-to-router connectivity between parts of a public ATM network over a private network.

Answer: B

QUESTION 5:

What are two drawbacks when the Internet is accessed through a dedicated subinterface implementation? (Select two.)

- A. The VRF will leak routes into the global routing table.
- B. The P router must run both IGP and BGP.
- C. The PE routers must perform Internet routing
- D. A limited set of Internet services will be implemented.
- E. A separate physical link or specific WAN encapsulation is required.

Answer: C, E

QUESTION 6:

Which three statements are correct regarding Layer 2 overlay VPNs and peer-to-peer VPNs? (Choose three.)

- A. Peer-to-peer VPNs require the establishment of virtual circuits to connect the different customer sites together.
- B. Peer-to-peer VPNs require the service provider to participate in the customer routing, accepting customer routes, transporting them across the service provider backbone, and finally propagating them to other customer sites.
- C. With peer-to-peer VPNs, the service provider is responsible for transport of Layer 2 frames between customer sites, and the customer takes responsibility for all higher layers.
- D. The implementation of Layer 2 overlay VPNs is the traditional switch-WAN model, implemented with technologies like X.25, Frame Relay or ATM.

E. With Layer 2 overlay VPNs, the service provider is not aware of customer routing has no information about customer routers.

F. It is simple to implement Layer 2 overlay VPNs because the Customer Edge (CE) router just needs a connection to the Service Provider's Edge (PE) router.

Answer: B, D, E

QUESTION 7:

Exhibit:

```
router bgp 65123
neighbor 10.1.1.1 remote-as 65111
neighbor 10.1.1.1 route-map test out

route-map test permit 10
match ip address 1
set metric 100
!
route-map test permit 20
set metric 200
!
access-list 1 permit 10.0.0.0
|
```

What will the BGP configuration accomplish? (Choose two.)

- A. Prefix 10.0.0.0/8 announced to neighbor 10.1.1.1. will have a MED of 100.
- B. Prefix 10.0.0.0/8 announced to neighbor 10.1.1.1. will have a MED of 200.
- C. All prefixes except the 10.0.0.0/8 prefix will not be announced to neighbor 10.1.1.1.
- D. All prefixes except the 10.0.0.0/8 prefix announced to neighbor 10.1.1 will have a MED of 0.
- E. All prefixes except the 10.0.0.0/8 prefix announced to neighbor 10.1.1 will have a MED of 200.

Answer: A, E

QUESTION 8:

Exhibit:

```
router bgp 65111
neighbor 172.16.1.1 remote-as 65111
neighbor 172.16.2.1 remote-as 65112
network 192.168.0.0
network 10.0.0.0

ip route 192.168.0.0 255.255.0.0 null0
```

What is wrong with the BGP configuration in the exhibit?

- A. The auto-summary configuration is missing.
- B. The network 10.0.0.0 statement is missing mask 0.255.255.255.

- C. The network 10.0.0.0 statement is missing mask 255.0.0.0
- D. The network 192.168.0.0 statement is missing mask 255.255.0.0.
- E. The network 192.168.0.0 statement is missing mask 0.0.255.255

Answer: D

QUESTION 9:

Why does a modern transit AS not need to have BGP synchronization enabled?

- A. A modern transit AS does not rely on redistribution of BGP routes into an IGP.
- B. A modern transit AS does not rely on redistribution of IGP routes into an IGP.
- C. A modern transit AS rely on redistribution of BGP routes into an IGP.
- D. A modern transit AS rely on redistribution of IGP routes into an IGP.
- E. A modern transit AS uses peer group to reduce the BGP configuration complexity.
- F. Modern transit AS internal (core) routers do not run BGP. Internal (core) routers normally use a default route to reach all external networks.

Answer: A

QUESTION 10:

Which of the following are standard filtering-oriented BGP communities? (Choose four.)

- A. No-export
- B. No-advertise
- C. Remote-as
- D. Local-as
- E. No-reflect
- F. Internet

Answer: A, B, D, F

QUESTION 11:

Between which types of routers are VPNv4 BGP routes propagated?

- A. CE and PE
- B. PE and P
- C. CE and P
- D. P and P
- E. PE and PE
- F. CE and CE

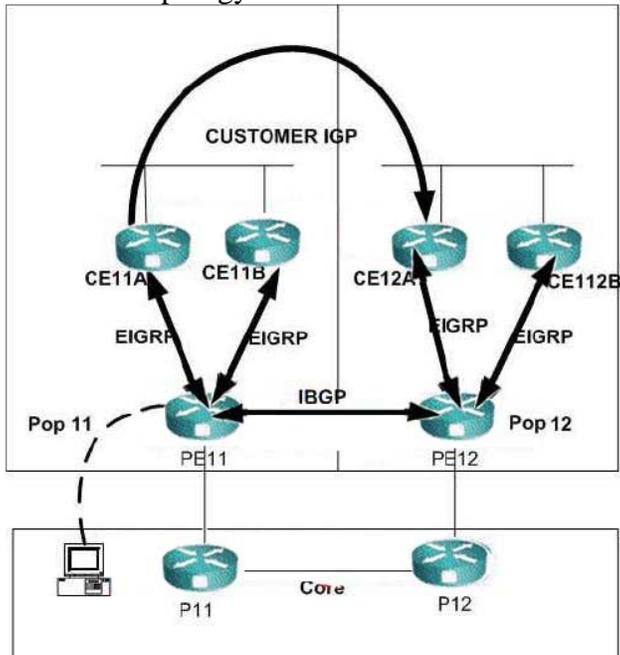
Answer: E

QUESTION 12:

In a managed CE router VPN implementation, what information does the service provider need access to on the customer's CE router?

- A. Access to all customer routes
- B. Access to all the customer's simple VPN VRF
- C. Access to the customer's CE router loopback addresses
- D. Access to the static routes of the customer's CE router

Answer: C

QUESTION 13:**SIMULATION****Network Topology Exhibit**

Certkiller .com is a stock broker service that is using EIGRP AS 1 as the IGP in its network. Certkiller .com has decided to establish connectivity between the New York site and the Los Angeles site. NiceNet, a service provider, has been selected by Certkiller .com to provide the connectivity between the two sites.

For this implementation NiceNet is using a MPLS VPN solution. NiceNet has already established MPLS connectivity between all of its PE routers. MPBGP connectivity has also been established on all PE routers.

Certkiller .com has successfully completed all necessary tasks on the CE routers, and NiceNet has successfully completed all necessary MPLS VPN configuration tasks on PE12. Your task is to successfully complete the implementation of the VPN connection by configuring the MPLS VPN on PE11 using the following information:

Privileged mode password: Certkiller

EIGRP is to be used as the PE - CE protocol for the VPN
MPBGP is being used to propagate VPN routing information between PE routers.
The interface supporting the Certkiller .com is Serial 0/0.101 and will have an IP address of 150.1.119.18 255.255.255.240
The following has been assigned to Certkiller .com's VPN:
The VRF_name is Certkiller _22
The RD is 22:10.
The route target for import is 21:10.
The route target for export is 21:10.
Configure EIGRP between CE11A and PE11 using AS number 10.
Redistribute from BGP to EIGRP AS 10 using a metric string of 10000 100 255 1 1500.
Redistribute from EIGRP AS 10 to BGP using metric string of 1.
VPN is operational when you see the routes for both the local (10.1.11.0) and the remote (10.1.12.0) sites in the VRF routing table.
Click on the host computer to start the simulation.

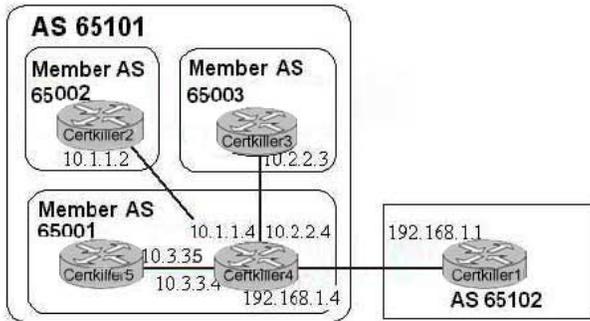
Answer:

```
ip cef
ip vrf Customer_6
rd 6:10
route-target import 14:10 !! these lines can be replaced by route-target
both 14:10
route-target export 14:10
Interface s0/0
ip vrf forwarding Customer_6
ip address 150.1.244.18 255.255.255.240
router bgp xx
address-family vpnv4
address-family ipv4 vrf Customer_6
neighbor x.x.x.x remote-as xx activate
redistribute eigrp metric 1
router eigrp 10
address-family ipv4 vrf Customer_6
network 10.1.11.0
network 10.1.12.0
redistribute BGP XX ! not enough info in lab sim
default-metric 10000 100 255 1 1500
```

QUESTION 14:

Network topology exhibit

Confederations



Based on the exhibit, what is the correct BGP confederation required on Certkiller 4?

- A. router bgp 65001
 bgp confederation identifier 65101
 bgp confederation peers 65002 65003
 neighbor 10.3.3.5 remote-as 65001
 neighbor 10.1.1.2 remote-as 65002
 neighbor 10.2.2.3 remote-as 65003
 neighbor 192.168.1.1 remote-as 65102
- B. router bgp 65001
 bgp confederation identifier 65001
 bgp confederation peers 65002 65003
 neighbor 10.3.3.5 remote-as 65001
 neighbor 10.1.1.2 remote-as 65002
 neighbor 10.2.2.3 remote-as 65003
 neighbor 192.168.1.1 remote-as 65102
- C. router bgp 65101
 bgp confederation peers 65001 65002 65003
 neighbor 10.3.3.5 remote-as 65001
 neighbor 10.1.1.2 remote-as 65002
 neighbor 10.2.2.3 remote-as 65003
 neighbor 192.168.1.1 remote-as 65102
- D. router bgp 65001
 bgp confederation peers 65002 65003
 neighbor 10.3.3.5 remote-as 65001
 neighbor 10.1.1.2 remote-as 65002
 neighbor 10.2.2.3 remote-as 65003
 neighbor 192.168.1.1 remote-as 65102

Answer: A

Explanation:

router bgp 65001 % Confed AS number, not the real BGP AS number as seen by real EBGP neighbor

bgp confederation identifier 65101 % Actual AS number as seen by real EBGP peer.

The confederation identifier refers to external AS number

neighbor 10.3.3.5 remote-as 65001

For B,C and D, they all have the incorrect bgp confederation identifier

QUESTION 15:

How would you implement an overlapping VPN when the sites contained in the overlapping portion of two simple VPNs contain an overlapping address space?

- A. Combine the overlapping VPN with a central service VPN.
- B. Implement a NAT service to provide unidirectional address translation.
- C. Implement a dual NAT service with a registered address to be imported and exported between the two central sites.
- D. Disable routing updates between the overlapping spaces and use a DNS implementation.
- E. Disable routing updates between the overlapping spaces and use static routing

Answer: C

QUESTION 16:

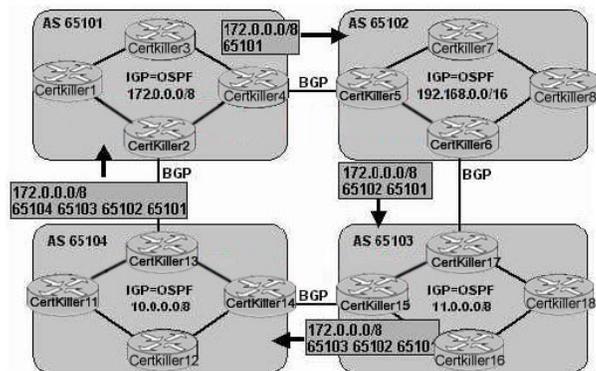
What are two characteristics of BGP confederations? (Choose two)

- A. All routers within the BGP confederation must support BGP confederations.
- B. The member AS numbers used with the confederation are visible from outside the confederation.
- C. An Intra-Confederation EBGP session behaves like an EBGP session when propagating routing updates.
- D. Intra-Confederation EBGP neighbors must be directly connected or ebgp-multihop must be configured.

Answer: A, D

QUESTION 17:

Network topology exhibit.



Why would Certkiller 2 accept or reject the BGP update from Certkiller 13 about the

172.0.0.0/8 prefix?

- A. Certkiller 2 will reject BGP update to avoid a potential routing loop.
- B. Certkiller 2 will accept the BGP update it is selected as the best path by the BGP route selection process.
- C. Certkiller 2 will reject the BGP update because it already has an OSPF (IGP) route to reach the 172.0.0.0/8 prefix
- D. Certkiller 2 will accept the BGP update and it will be used as a backup path to reach the 172.0.0.0/8 prefix if the OSPF (IGP) route is not available.

Answer: A

QUESTION 18:

Exhibit

Neighbor 2.2.2.2 maximum-prefix 1000

What is the purpose of the BGP configuration command in the exhibit?

- A. It limits the number of prefixes that can be sent to neighbor 2.2.2.2 to 1000.
- B. It limits the number of prefixes that can be sent to and received from the 2.2.2.2 neighbor to 1000.
- C. If the number of prefixes sent to the 2.2.2.2 neighbor exceeds 650, the router starts to generate a warning message.
- D. If the number of prefixes received from the 2.2.2.2 neighbor exceeds 650, the router starts to generate a warning message.
- E. If the number of prefixes received from the 2.2.2.2 neighbor exceeds 1000, the neighbor relationship to 2.2.2.2 will be dropped.

Answer: E

QUESTION 19:

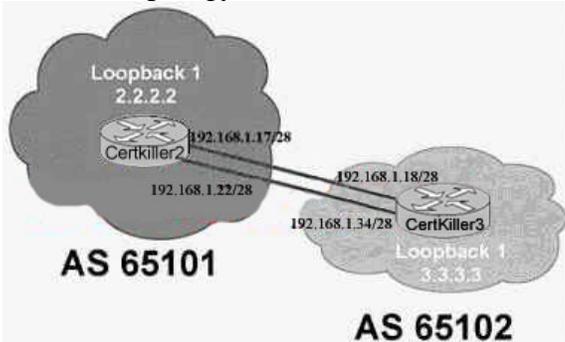
Which two statements are true about an EBGP session or an IBGP session? (Choose two.)

- A. IBGP uses AS-Path to detect routing information loops within the AS.
- B. EBGP routes have a default Admin Distance of 20 and IBGP routes have a default Admin Distance of 200.
- C. No BGP attributes are changed in EBGP updates except for the next-hop attribute if next-hop-self is configured.
- D. Routes learned from an EBGP peer not advertised to another EBGP peer to prevent routing information loops.
- E. IBGP uses split horizon to prevent routing information loops; routes learned from an IBGP peer are not advertised to another IBGP peer.

Answer: B, E

QUESTION 20:

Network topology exhibit



What is the proper configuration on Certkiller 2 to support a single BGP session between Certkiller 2 and Certkiller 3 and to enable load sharing across the two links?

- A. router bgp 65101
neighbor 192.168.1.18 remote-as 65102
neighbor 192.168.1.34 remote-as 65102
neighbor 192.168.1.38 update-source loopback1
neighbor 192.168.1.34 update-source loopback1
maximum-paths 2
- B. router bgp 65101
neighbor 192.168.1.18 remote-as 65102
neighbor 192.168.1.34 remote-as 65102
neighbor 192.168.1.38 ebgp-multihop 2
neighbor 192.168.1.34 ebgp-multihop 2
maximum-paths 2
- C. router bgp 65101
neighbor 3.3.3.3 remote-as 65102
neighbor 3.3.3.3 ebgp-multihop 2
neighbor 3.3.3.3 update-source loopback1
!
ip route 3.3.3.3 255.255.255.255 192.168.1.18
ip route 3.3.3.3 255.255.255.255 192.168.1.34
- D. router bgp 65101
neighbor 3.3.3.3 remote-as 65102
neighbor 3.3.3.3 update-source loopback1
maximum-paths 2
!
ip route 3.3.3.3 255.255.255.255 192.168.1.18
ip route 3.3.3.3 255.255.255.255 192.168.1.34

Answer: C

QUESTION 21:

Which statement is true about implementing a hub-and-spoke overlay VPN?

- A. A hub-and-spoke overlay VPN has no single point of failure
- B. A spoke site has direct connectivity to other spoke sites.
- C. If a dynamic routing protocol like RIP is used, split horizon must be disabled at the hub router, or point-to-point sub-interfaces must be used.
- D. The total number of virtual circuits required to implement a hub-and-spoke overlay VPN is $(N(N-1))/2$ where N = number of sites.

Answer: C

QUESTION 22:

What is the most typical use of an overlapping VPN?

- A. Service providers supporting managed CE Services
- B. Service providers supporting carrier-over-carrier MPLS implementations
- C. Simple intranet VPN implementations
- D. Companies that use MPLS VPNs to implement both intranet and extranet services

Answer: D

QUESTION 23:

Exhibit

```
! Certkiller2 has IP address 10.1.1.2
hostname Certkiller3
|
|Output omitted
|
router bgp 65101
 redistribute static route-map test
 neighbor 10.1.1.2 remote-as 65101
 neighbor 10.1.1.2 send community
|
ip route 11.2.3.0 255.255.255.0 serial0 tag 100
|
route-map test permit 10
 match tag 100
 set community no-export 65101:123
```

What will the configuration on Certkiller 3 accomplish (Choose three.)

- A. Certkiller 3 will use a static route to reach the customer 11.2.3.0/24 prefix and tag that static route with a route tag of "100".
- B. The "no-export" options on the community indicate that the BGP updates should not be sent to any IBGP or EBGP neighbors.

- C. Certkiller 3 will establish an IBGP session with Certkiller 2 and Certkiller 1 should send BGP community information to Certkiller 2 in the IBGP updates.
- D. All BGP routes learned from Certkiller 2 will be tagged with a route tag of 100 and a BGP community number of "65101:123"
- E. Certkiller 1 will redistribute the static route into BGP and use a route-map named "test" to match on route tag "100," then setting the community to "no-export" with a community number of "65101:123".

Answer: A, C, E

QUESTION 24:

Which three statements about MPLS VPNs are true? (Choose three.)

- A. PE routers do not participate in customer routing.
- B. MPLS VPN is similar to using the peer-to-peer dedicated PE router approach.
- C. Customer can use overlapping addresses.
- D. Each customer is assigned an independent routing table (virtual routing and forwarding table - VRF).
- E. The P Router's routing table contains both the global IP routes and the customer routes.
- F. CE routers connect directly to the service provider's P routers.

Answer: B, C, D

QUESTION 25:

What condition must be met for the receiving PE router to install VPNv4 routes into a VRF?

- A. If at least one RD attached to the VPNv4 route matches at least one import RD configured in the VRF
- B. If at least one RT attached to the VPNv4 route matches at least one import RT configured in the VRF
- C. If at least one RD attached to the VPNv4 route matches at least one export RD configured in the VRF
- D. If at least one RT attached to the VPNv4 route matches at least one export RT configured in the VRF

Answer: B

QUESTION 26:

Which description is NOT a topology where a central services VPN would be used?

- A. A service provider offers services to all customers by allowing them access to a common

VPN.

B. Two or more companies want to exchange information by sharing a common set of servers.

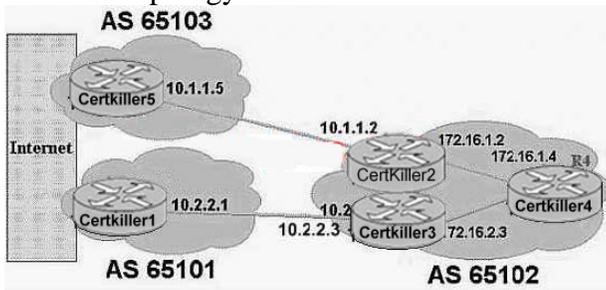
C. A security-conscious company separates its departments and allows them to access only to common servers.

D. Several interconnected subsidiaries of a company require optimal communication between many sites.

Answer: D

QUESTION 27:

Network topology exhibit



What the correct BGP configuration on Certkiller 3 and Certkiller 2 to have AS 65102 prefer the path to AS 65101 over the path to AS 65103?

A. hostname Certkiller 3

```
!  
router bgp 65102  
neighbor 172.16.2.4 remote-as 65102  
neighbor 10.2.2.1 remote-as 65101  
bgp default local-preference 140  
hostname Certkiller 2
```

```
!  
router bgp 65102  
neighbor 172.16.1.4 remote-as 65102  
neighbor 10.1.1.5 remote-as 65103  
bgp default local-preference 120
```

B. hostname Certkiller 3

```
!  
router bgp 65102  
neighbor 172.16.2.4 remote-as 65102  
neighbor 10.2.2.1 remote-as 65101  
neighbor 10.2.2.1 local-preference 140  
hostname Certkiller 2
```

```
!  
router bgp 65102  
neighbor 172.16.1.4 remote-as 65102  
neighbor 10.1.1.5 remote-as 65103
```

neighbor 10.1.1.5 local-preference 120

Answer: A

QUESTION 28:

Which prefix-list will permit the 172.16.1.32 255.255.255.240 prefix?

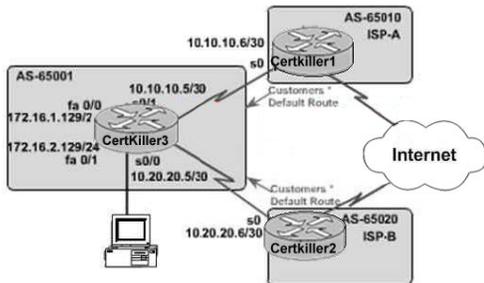
- A. ip prefix-list test permit 172.16.1.0/28
- B. ip prefix-list test permit 172.16.1.0/8 le 30
- C. ip prefix-list test permit 172.16.1.0/16 eq 28
- D. ip prefix-list test permit 172.16.1.0/8 ge 29

Answer: B

QUESTION 29:

SIMULATION

Network topology exhibit



Certkiller .com recently issued a contract to ISP (AS 65020) to provide a second Internet connection their Boston location. A new router, Certkiller 3, will need to be configured to connect to the Internet Service Providers. It is necessary to configure this router to exchange routes with both ISPs, and to filter the routing advertisements out to the ISPs. The basic configuration for the new router has been entered. It is only necessary to complete the BGP configuration steps outlined below.

On the Certkiller 3 router:

Configure BGP to exchange routes with ISP-A (AS 65010).

Configure BGP to exchange routes with ISP-B (AS 65020).

Only announce the two local Ethernet LAN prefixes (172.16.1.0/24 and 172.16.2.0/24) to the ISPs.

Configure the appropriate as-path access-list to permit only locally originated routes to be advertised in outbound routing updates to the ISPs.

Use as-path ACL number 10 containing only one statement.

Start the simulation by clicking on the local host icon.

Answer:

hostname Certkiller 3

```
!  
interface fastethernet 0/0  
ip address 172.16.1.129 255.255.255.0  
no shutdown  
!  
interface fastethernet 0/1  
ip address 172.16.2.129 255.255.255.0  
no shutdown  
!  
router bgp 65001  
network 172.16.1.0 mask 255.255.255.0  
network 172.16.2.0 mask 255.255.255.0  
neighbor 10.10.10.6 remote-as 65010  
neighbor 10.20.20.6 remote-as 65020  
neighbor 10.10.10.6 filter-list 10 out  
neighbor 10.20.20.6 filter-list 10 out  
!  
ip as-path access-list 10 permit ^$  
!
```

QUESTION 30:

DRAG DROP

Label allocation and distribution in a frame-mode MPLS network is accomplished by the following actions. While some of the actions do occur periodically in a network, initially some of these actions require that other actions have already completed. Drag the actions on the left to the boxes on the right in which these actions must initially occur.

- IP routing protocols build IP routing tables
- Local labels are generated for the LIB and LFIB
- LSRs announces their assigned labels to all other LSRs
- Next-hop labels are inserted into the LIB, FIB, and LFIB tables
- Routing information is exchanged using IP routing protocols.

Place here

- Place first action here
- Place second action here
- Place third action here
- Place fourth action here
- Place last action here

Answer:

Place here

- Local labels are generated for the LIB and LFIB
- Routing information is exchanged using IP routing protocols.
- IP routing protocols build IP routing tables
- Next-hop labels are inserted into the LIB, FIB, and LFIB tables
- LSRs announces their assigned labels to all other LSRs

QUESTION 31:

Exhibit, show ip bgp neighbor 10.1.1.1

Output

```
Certkiller3#show ip bgp neighbor 10.1.1.1
BGP neighbor is 10.1.1.1, remote AS 65102, external link
BGP version 4, remote router ID 0.0.0.0
BGP state = Idle
Last read 00:00:04, hold time is 180, keepalive interval is 60 seconds
Received 0 messages 0 notification 0 in queue
Send 0 messages 0 notification 0 in queue
Route refresh request: received 0, sent 0
Default minimum time between advertisement runs is 30 seconds
For address family: IPv4 Unicast
BGP table version 1, neighbor version 0
Index 2, Offset 0, Mask 0x4
0 accepted prefixes consume 0 bytes
Prefix advertised 0 suppressed 0 withdrawn 0
Connection established 0, dropped 0
Last reset never
External BGP neighbor not directly connected.
No active TCP connection
```

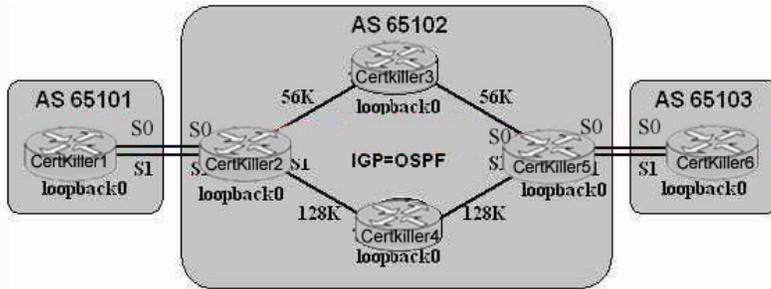
What is causing the BGP session to the 10.1.1.1 neighbor to remain in the Idle state?

- A. There is an AS number configuration error
- B. The EBGp-multihop value for neighbor 10.1.1.1 is set to the default value.
- C. There is an access-list blocking TCP port 179 traffic between the two BGP neighbors.
- D. The BGP session is using the loopback interface but the update-source is not set to specify the loopback interface.

Answer: B

QUESTION 32:

Network topology exhibit



What can prevent the corresponding BGP session from being successfully established?

- A. Certkiller 2 and Certkiller 5 cannot establish the IBGP session because Certkiller 2 and Certkiller 5 are not directly connected.
- B. Certkiller 1 and Certkiller 2 establish the EBGP session if the BGP holddown timers between the two EBGP neighbors are different.
- C. Certkiller 2 and Certkiller 5 cannot establish the IBGP session using the loopback0 interface if the EBGP-multihop value is set to the default value.
- D. Certkiller 1 and Certkiller 2 cannot establish the EBGP session using the loopback0 interface if the EBGP-multihop value is set to the default value.

Answer: D

QUESTION 33:

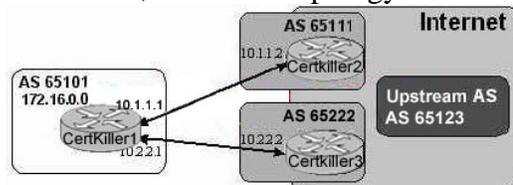
Exhibit #1, BGP configuration

```

router bgp 65101
neighbor 10.1.1.2 remote-as 65111
neighbor 10.2.2.2 remote-as 65222
neighbor 10.1.1.2 weight 900
neighbor 10.1.1.2 route-map test in
!
route-map test permit 10
match as-path 1
set weight 999
!
route-map test permit 20
!
ip as-path access-list 1 permit _65123$

```

Exhibit #2, Network Topology



Which two paths does the router in AS 65101 prefer? (Choose two.)

- A. AS 65111 for reaching all BGP prefixes

- B. AS 65222 for reaching all BGP prefixes
- C. AS 65222 for reaching all BGP prefixes that originate in AS 65123
- D. AS 65111 for reaching all BGP prefixes that do not originate in AS 65123
- E. AS 65222 for reaching all BGP prefixes except those matching AS-Path access list 1

Answer: C, D

QUESTION 34:

Exhibit

```
route-map test permit 10
match ip address prefix-list def
match as-path 1
set weight 111

route-map test permit 20
match ip address prefix-list def
set weight 110
|
route-map test permit 30
set weight 109
|
ip as-path access-list 1 permit _65111$
|
ip prefix-list def seq 10 permit 0.0.0.0/0
```

What does the route map in the exhibit accomplish? Select three.

- A. The weight of a default route that originates in AS65111 will be set to 111.
- B. The weight of a default route that originates in AS65111 will be set to 110.
- C. The weight of all routes that originates in AS65111 will be set to 111.
- D. The weight of all routes that originates in any AS will be set to 109.
- E. The weight of all routes that originates in any AS will be set to 110.
- F. The weight of all routes other than a default that originates in any AS will be set to 109.

Answer: A, D, F

QUESTION 35:

Network topology exhibit

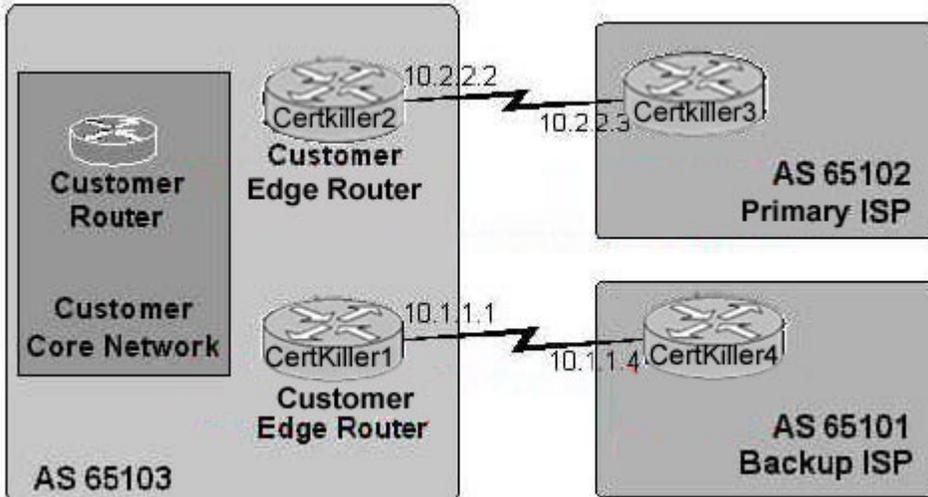


Exhibit #2

```
router bgp 65103
 neighbor 10.1.1.4 remote-as 6510
```

```
!
 route-map test permit 10
!
```

Which two configuration commands will complete the configuration on Certkiller 1 so that when Certkiller 1 send the BGP updates to Certkiller 4, three copies of the customer's AS number (65103) will be prepended to the AS-Path? (Choose two)

- A. Set as-path prepend 3
- B. Set as-path tag 65103 65103 65103
- C. Neighbor 10.1.1.4 route-map test in
- D. Neighbor 10.1.1.4 route-map test out
- E. Set as-path prepend 65103 65103 65103

Answer: D, E

QUESTION 36:

By default, which BGP path attribute is stripped in the outgoing IBGP updates?

- A. Origin
- B. AS-path
- C. Next Hop
- D. Communities
- E. Local Preference

Answer: D

Explanation: Community is stripped in outgoing BGP updates. by default is

no_export

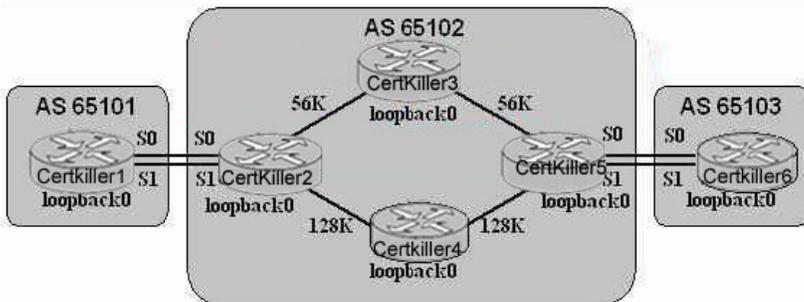
Note: Key to answer is IBGP

Wrong Answer

1. A Well known mandatory (No stripped)
2. B Well known mandatory (No stripped)
3. C Well known mandatory (No stripped)
4. E Local Preference is local to AS, because this is IBGP it is propagated through the AS

QUESTION 37:

Network topology Exhibit



How should the EBGP session between Certkiller 1 and Certkiller 2 be established to enable load balancing? (Choose three.)

- A. Use the maximum-paths 2 option
- B. Use the ebgp-multihop 2 option
- C. Use static routes on Certkiller 1 and Certkiller 2 to reach the other router's loopback
- D. Establish a single EBGP session using the loopback0 interface IP address on Certkiller 1 and Certkiller 2
- E. Establish two EBGP sessions using both the S0 and D1 interface IP address on Certkiller 1 and Certkiller 2

Answer: B, C, D

Explanation:

1. B R1 - R2 requires a EBGP ebgp-multihop of 2.
2. C Requires two static routes going to each other router loopback interface
3. D A BGP session of course is required. Only a single connection is required. Use the loopback is what allows things to work even during a failure

Wrong Answer

1. A This is not required because there is only on EBGP path. Would need two if static routes were not define and were going between different routers
2. E Not an efficient way of doing things requires double the tables and bandwidth for updates

QUESTION 38:

Network topology Exhibit

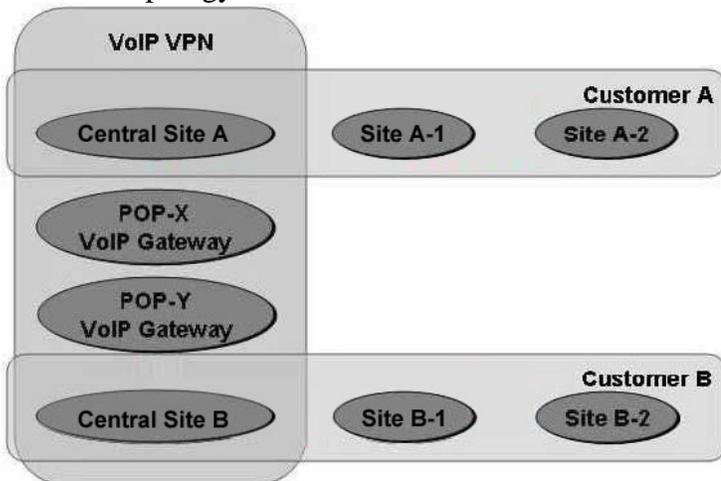


Exhibit #2, Requirements

- All sites of Customer A (except from the central site) can share the same routing information.
- All sites of Customer B (except from the central site) can share the same routing information.
- The two POP VoIP gateway sites can share the same routing information.
- Central A site can connect to all of Customer A sites and the VoIP sites.
- Central B site can connect to all of Customer B sites and the VoIP sites.
- Sites A-1 and A-2 can not connect to Sites B-1 and B-2.
- Central A site can not connect to Sites A-1 and A-2.
- Central A site can not connect to Sites B-1 and B-2.
- Sites B-1 and B-2 can not connect to Sites A-1 and A-2.

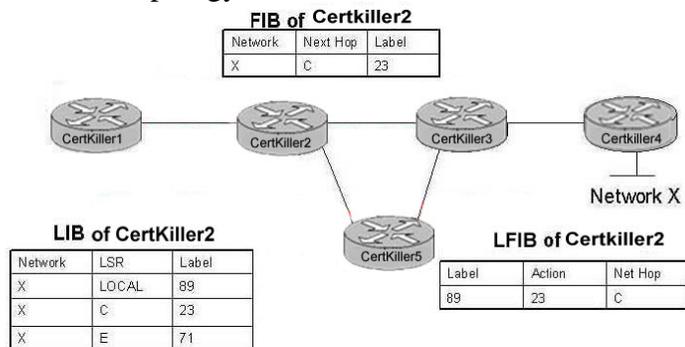
How many VRF tables are needed to support three VPNs (Customer A, Customer B, and a VoIP VPN) with the requirements?

- A. 1 VRF
- B. 3 VRFs
- C. 4 VRFs
- D. 5 VRFs
- E. 7 VRFs

Answer: D

QUESTION 39:

Network topology Exhibit



The relevant entries in the FIB, LIB, and LFIB tables for a frame-mode MPLS network is shown. If the link between Routers B and C was lost, what would be the relevant entries in the LFIB table for label, action, and next hop?

- A. 89, 71, Certkiller 5
- B. 71, 89, Certkiller 5
- C. 89,23, Certkiller 4
- D. 71,23, Certkiller 4

Answer: A

QUESTION 40:

Which well-known tunneling technology provides network layer authentication and optional encryption to make data transfer over the Internet secure?

- A. Layer 2 tunneling
- B. GRE
- C. IPSec
- D. Peer-to-peer
- E. Layer 2 forwarding

Answer: C

QUESTION 41:

Since route distinguishers (RD) can not identify participation in more than one VPN, what is required to support complex MPLS VPNs?

- A. Route target (RT)
- B. Route tags
- C. Route maps
- D. Dedicated PE routers
- E. Site of origin (SOO)

Answer: A

QUESTION 42:

When using BGP to carry external routes, what is the ideal interaction between BGP and IGP running within a Transit AS?

- A. The external routes should be redistributed into an IGP.
- B. The IGP should carry the routes to reach the BGP next-hops and the core subnets within the AS.
- C. The IGP should carry the core subnets within the AS. Then a mutual redistribution between BGP and the IGP should be performed.
- D. The external routes should be redistributed into a Link-State IGP (like OSPF or IS-IS) with filtering to reduce the size of the Link-State database.

Answer: B

QUESTION 43:

What are the three main BGP scaling issues for service providers? (Choose three.)

- A. BGP policy scaling issues
- B. IBGP full mesh scaling issues
- C. BGP to IGP redistribution scaling issue
- D. BGP updates and BGP table size scaling issues

Answer: A, B, D

QUESTION 44:

Which four attributes are used by BGP to detect routing loops? (Choose four.)

- A. AS.Path
- B. Cluster ID
- C. Cluster List
- D. Originator ID
- E. Community ID

Answer: A, B, C, D

QUESTION 45:

What is the difference between a managed CE router VPN and a central services VPN?

- A. In a managed CE router VPN, only customer loopback addresses are marked to be imported into the network management VPN.
- B. In a managed CE router VPN, only default customer routes are marked to be imported into the network management VPN.
- C. In a managed CE router VPN, all customer routes are marked to be imported into the network management VPN.
- D. In a managed CE router VPN, only customer loopback addresses and default customer routes are marked to be imported into the network management VPN.

Answer: A

QUESTION 46:

DRAG DROP

You work as network administrator at Certkiller.com. Your Certkiller boss asks you to match the VPN types to the appropriate descriptions.

Select from these

- | | |
|---------------------|----------------------|
| overlapping VPN | Central services VPN |
| Managed network VPN | Simple VPN |

Place here

- | | |
|---|------------|
| Some sites can participate in more than one VPN | Place here |
| All sites can communicate with a common server but not with each other. | Place here |
| Sometimes it is called a "rainbow" VPN implementation. | Place here |
| Every site can communicate with every other site. | Place here |

Answer:

- | | |
|---|----------------------|
| Some sites can participate in more than one VPN | overlapping VPN |
| All sites can communicate with a common server but not with each other. | Central services VPN |
| Sometimes it is called a "rainbow" VPN implementation. | Managed network VPN |
| Every site can communicate with every other site. | Simple VPN |

QUESTION 47:

Which three statements are true about route reflectors? (Choose three.)

- A. If the route is learned from an EBGP peer by the route reflector, it is reflected to all IBGP and EBGP peers.
- B. If the route is learned from a non-client IBGP peer by the route reflector, it is reflected to all EBGP peers only.
- C. If the route is learned from a non-client IBGP peer by the route reflector, it is reflected to EBGP peers and clients only.
- D. If the route is learned from a client IBGP peer by the route reflector, it is reflected to all clients only, except the originating client.
- E. If the route is learned from a client IBGP peer by the route reflector, it is reflected to all EBGP peers, non-clients, and clients (except the originating client).

Answer: A, C, E

Explanation:

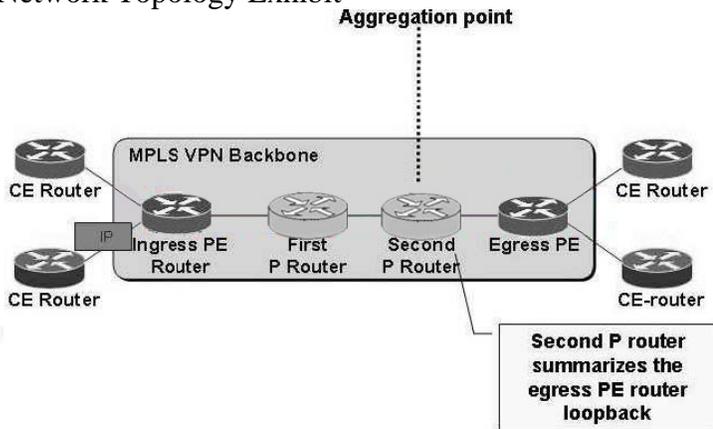
- 1. A routes update to EBGP peers and clients
- 2. C routes update to non-clients and clients
- 3. E route updates to all except if originator

Wrong

- 1. B route are sent to clients
- 2. D routes updates are sent to non clients and EBGP peers

QUESTION 48:

Network Topology Exhibit



What problem can be caused by the second P router summarizing the loop address of the egress PE router?

- A. The first P router will be faced with a VPN label which it does not understand.
- B. The second P router will be faced with a VPN label which it does not understand.
- C. The egress PE router will not be able to establish a label switch path (LSP) to the ingress PE router.
- D. A label switch path (LSP) will be established from the ingress PE router to the egress PE router, an event that is not desirable.
- E. The ingress PE router will not be able to receive the VPN label the egress PE router via MP-IBGP.

Answer: B

QUESTION 49:

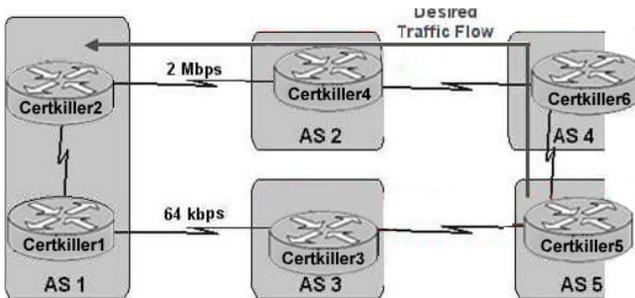
In cell-mode MPLS, where is the label imposed?

- A. Between Layer 1 and Layer 2 headers.
- B. Between Layer 2 and Layer 3 headers.
- C. In the VPI/VCI fields
- D. In the VPI field alone

Answer: C

QUESTION 50:

Network Topology Exhibit



AS-Path prepending is used AS1 in order to influence the return traffic path from AS 5 to AS 1 through the higher speed via AS 2. _____ needs to be configured for AS-Path prepending and a minimum of _____ of the AS number should be prepended.

- A. Certkiller 1; one copy
- B. Certkiller 2; one copy
- C. Certkiller 1; two copies
- D. Certkiller 2; two copies
- E. Certkiller 2; three copies

Answer: C

QUESTION 51:

Which one of the following is used to prevent routing loops in MPLS VNP networks with multihomed sites?

- A. RT
- B. SOO
- C. RD
- D. AS-Path
- E. MPLS Label
- F. MED

Answer: B

QUESTION 52:

In a multihomed environment with two ISP connections, which two statements are true? (Choose two.)

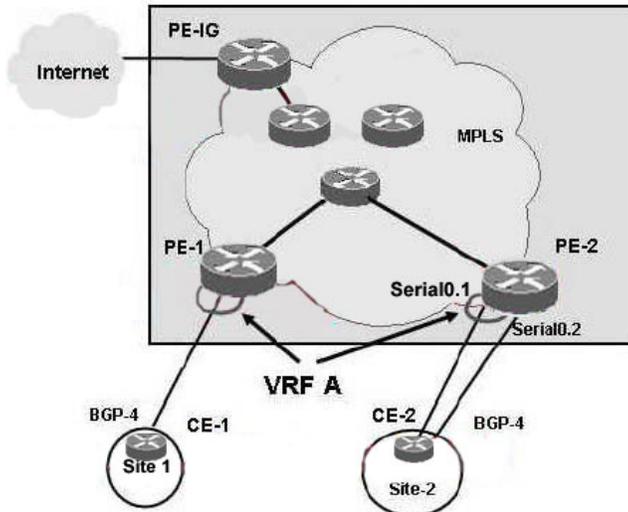
- A. The customer should not be configured to act as a transit AS between the two ISPs.
- B. It is recommend that the multi-homed customer use a registered (public) AS number.
- C. AS-Path prepending can be configured on the customer's edge router to influence the BGP path selection process for the outbound traffic (traffic from the customer to the ISPs).
- D. The customer can use Local Preference on the customer's edge routers to influence the BGP path selection process for the inbound traffic (traffic from the ISPs to the customer).

E. The advertisement of the customer's IP address space be conditioned by the customer's edge routers by using a static route to the null0 interface and by using the proper network statement under router bgp.

Answer: A, B

QUESTION 53:

Network topology exhibit; Internet Access through a Dedicated Subinterface



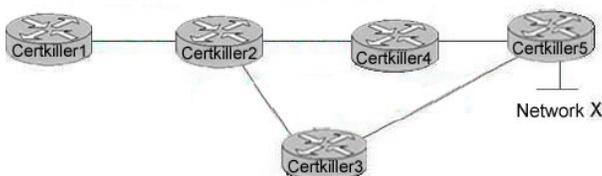
Internet access is through a dedicated subinterface implementation. Which of its routing tables will PE-2 use to forward packets from Site2 to Site 1?

- A. the global routing table
- B. the VPNv4 routing table
- C. the VRF A routing table
- D. the VRF A FIB
- E. the PE-IG routing table

Answer: A

QUESTION 54:

Network topology exhibit



Liberal label retention and unsolicited downstream distribution is being used in the frame-mode MPLS network shown in the exhibit. The best to network X as defined by the IGP for network is Certkiller 1- Certkiller 2- Certkiller 4- Certkiller 5.

What will Router Certkiller 3 and route

- A. Both routers will add it to their LIB.
- B. Both routers will ignore it.
- C. Router Certkiller 3 will add the label to its LIB. Router Certkiller 4 will ignore it.
- D. Router Certkiller 4 will add the label to its LIB. Router Certkiller 3 will ignore it.

Answer: A

QUESTION 55:

Exhibit

```
Certkiller1#show ip bgp neighbors 2.2.2.2
BGP neighbor is 2.2.2.2, remote AS 102, internal link
Index 1, Offset 0, Mask 0x2
BGP version 4, remote router ID 88.0.0.1
BGP state = Established, table version = 1, up for 00:14:52
Last read 00:00:52, said time is 180, keepalive interval is 60 seconds
Minimum time between advertisement runs is 5 seconds
Received 233 messages, 0 notifications, 0 in queue
Sent 206 messages, 0 notifications, 0 in queue
Prefix advertised 0, suppressed 0, withdrawn 0
Connection = established 17, dropped 16
Last reset 00:46:02 due to Use reset
18 accepted prefixes consume 576 bytes, maximum limit 20
Threshold for warning message 80%
0 history paths consume 0 bytes
```

Which two statements are true? (Choose two).

- A. Certkiller 1 has accepted 20 prefixes from the 2.2.2.2 IBGP neighbor.
- B. Certkiller 1 generated a warning message to the router's console after the 2.2.2.2 IBGP neighbor sent 16 prefixes to Certkiller 1.
- C. Certkiller 1 generated a warning message to the router's console after the 2.2.2.2 IBGP neighbor sent 15 prefixes to Certkiller 1.
- D. Certkiller 1 will drop its neighbor relationship to the 2.2.2.2 IBGP neighbor if 2.2.2.2 sends two more additional prefixes to Certkiller 1.
- E. Certkiller 1 will drop its neighbor relationship to the 2.2.2.2 IBGP neighbor if 2.2.2.2 sends three more additional prefixes to Certkiller 1.

Answer: B, E

QUESTION 56:

What is the main advantage of the BGP Route Refresh capability as compared to the "BGP soft-reconfiguration in" option?

- A. Route Refresh is not Cisco proprietary.
- B. Route Refresh does not require additional memory.
- C. Route Refresh does not require the BGP session to be taken down.
- D. Route Refresh allows BGP routers to automatically send out new BGP updates.

Answer: B

QUESTION 57:

With MPLS VPNs, which four elements will be included in an MP-BGP update exchange between the PE routers? (Choose four.)

- A. VPNv4 Address
- B. Route tag
- C. Label used for VPN packet forwarding
- D. Mandatory BGP attributes like As-Path and Next-Hop
- E. Route target
- F. Implicit null label used for PHP operation

Answer: A, C, D, E

QUESTION 58:

What is true of MPLS TE?

- A. Only the ingress LSR must see the entire topology of the network.
- B. Every LSR needs additional information about links in the network, available resources, and constraints.
- C. Every core router must be able to create an LSP tunnel on demand.
- D. Both RSVP and CR-LDP are used in conjunction to establish traffic engineering (TE) tunnels and to propagate the labels.

Answer: B

QUESTION 59:

Network Topology Exhibit

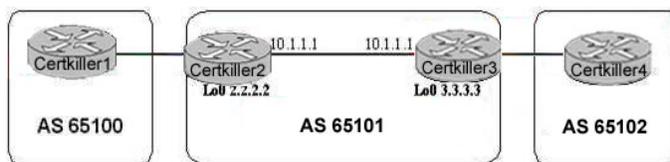


Exhibit #2: Certkiller 2 configuration

Given the following configuration for Certkiller2

```
! output omitted
|
hostname Certkiller2
|
interface loopback 0
ip address 2.2.2.2

interface e0
ip address 10.1.1.1 255.255.255.0
|
router bgp 65101
neighbor 172.16.1.1 remote-as 65100
neighbor 3.3.3.3 remote-as 65101
no sync

router eigrp 101
network 10.0.0.0
network 2.0.0.0
```

When Certkiller 2 sends the TCP SYN packet to Certkiller 3 to establish the IBGP session, what will be the source IP address of the TCP SYN packet from Certkiller 2 to Certkiller 3?

- A. 2.2.2.2
- B. 3.3.3.3
- C. 10.1.1.1
- D. 10.1.1.2

Answer: C

QUESTION 60:

Exhibit, Output

```
Certkiller12#sh ip route bgp
B 192.168.24.0/24 [20/0] via 172.31.1.1, 02:34:00
B 192.168.21.0/24 [20/0] via 172.31.1.2, 02:34:00
B 192.168.20.0/24 [20/0] via 172.31.1.2, 01:20:33
B 10.0.0.0/8 [20/0] via 172.31.1.1, 02:34:29
B 192.168.23.0/24 [20/0] via 172.31.1.1, 02:34:00
B 192.168.22.0/24 [20/0] via 172.31.1.2, 02:34:00
B 192.168.12.0/22 [20/0] via 172.31.1.1, 01:40:07
B 192.168.20.0/22 [20/0] via 172.31.1.1, 01:38:12
```

Which prefix-list can be used to permit only the 192.168.12.0/22 and the 192.168.20.0/22 prefixes?

- A. ip prefix-list test permit 192.168.0.0/19 le 22
- B. ip prefix-list test permit 192.168.0.0/16
- C. ip prefix-list test permit 192.168.0.0/22

- D. ip prefix-list test permit 192.0.0.0/8 ge 22
- E. ip prefix-list test permit 192.168.0.0/20 le 22

Answer: A

QUESTION 61:

What is the VPN type where sites of two different companies are connected together via Frame Relay virtual circuits

- A. overlay intranet VPN
- B. overlay extranet VPN
- C. peer-to-peer access VPN
- D. peer-to-peer Internet VPN
- E. MPLS simple VPN
- F. MPLS overlapping VPN

Answer: B

QUESTION 62:

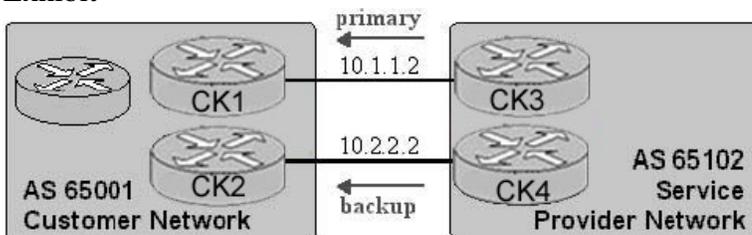
What is the range of values from which an ISP can assign a private AS number?

- A. 32768 to 65535
- B. 64512 to 65535
- C. 65101 to 65535
- D. 65001 to 65535

Answer: B

QUESTION 63:

Exhibit



What is the correct BGP configuration on CK1 and CK2 to have AS 65102 prefer the path to CK1 over the path to CK2 to reach AS 65001? (Influence the path selection decision by using the MED attribute)

- A. hostname CK1
- !
- router bgp 65001

```
neighbor 10.1.1.2 remote-as 65102
neighbor 10.1.1.2 route-map setmed
!
route-map setmed permit 10
set metric 100
hostname CK2
!
router bgp 65001
neighbor 10.2.2.2 remote-as 65102
neighbor 10.2.2.2 route-map setmed
!
route-map setmed permit 10
set metric 200
B. hostname CK1
!
router bgp 65001
neighbor 10.1.1.2 remote-as 65102
neighbor 10.1.1.2 route-map setmed out
!
route-map setmed permit 10
set metric 100
hostname CK2
!
router bgp 65001
neighbor 10.2.2.2 remote-as 65102
neighbor 10.2.2.2 route-map setmed out
!
route-map setmed permit 10
set metric 200
C. hostname CK1
!
router bgp 65001
neighbor 10.1.1.2 remote-as 65102
neighbor 10.1.1.2 route-map setmed in
!
route-map setmed permit 10
set metric 100
hostname CK2
!
router bgp 65001
neighbor 10.2.2.2remote-as 65102
neighbor 10.2.2.2 route-map setmed in
!
route-map setmed permit 10
set metric 200
D. hostname CK1
```

```
!  
router bgp 65001  
neighbor 10.1.1.2 remote-as 65102  
neighbor 10.1.1.2 route-map setmed in  
!  
route-map setmend permit 10  
set metric 200  
hostname CK2  
!  
router bgp 65001  
neighbor 10.2.2.2 remote-as 65102  
neighbor 10.2.2.2 route-map setmed in  
!  
route-map setmed permit 10  
set metric 100
```

Answer: B

QUESTION 64:

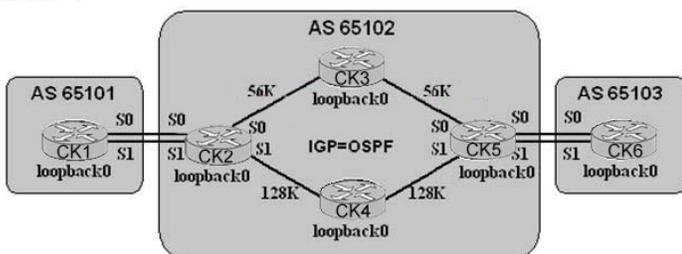
Which type of VPN implementation is simple and cost effective because all sites connect only to the PE router and as a result optimum routing between sites is enabled by default?

- A. peer-to-peer VPN
- B. overlapping VPN
- C. hub-and-spoke overlay VPN
- D. fully-meshed overlap VPN
- E. client-server VPN

Answer: A

QUESTION 65:

Exhibit



What can prevent the corresponding BGP session from being successfully established?

- A. CK2 and CK5 cannot establish the IBGP session if the BGP Hello Timer between the

two IBGP neighbors is different.

B. CK1 and CK2 cannot establish the EBGP session if the BGP Hello Timer between the two EBGP neighbors are different.

C. CK1 and CK2 cannot establish the EBGP session using the loopback0 interface if the EBGP-multihop value is set to 2.

D. CK2 and CK5 cannot establish the IBGP session because they are not using the loopback0 interface to establish the IBGP session.

E. CK2 and CK5 cannot establish the IBGP session if CK3 and CK4 have an access list permitting only TCP port 80 and IP protocol number 89 traffic.

Answer: E

QUESTION 66:

What is true about the label in frame-mode MPLS?

A. 32 bit label with an 8 bit label field.

B. 32 bit label with a 20 bit label field.

C. 20 bit label with a 3 bit Bottom of Stack field.

D. 20 bit label with a 1 bit TTL field.

Answer: B

QUESTION 67:

Which routers in an MPLS VPN service provider's network are used to connect to the customer's routers?

A. P routers

B. CE routers

C. PE routers

D. P and PE routers

Answer: C

QUESTION 68:

Which three statements are true about a Cluster-ID? (Choose three)

A. A Cluster-ID configuration is only required on the route reflectors.

B. The Cluster-ID is the router ID of the originator of a route within the local AS.

C. Cluster-List is used to track Cluster IDs similar to the way that the AS-PATH is used to track AS numbers.

D. When a cluster contains multiple route reflectors, all the route reflectors in the cluster need to be configured with the same Cluster-ID.

Answer: A, C, D

QUESTION 69:

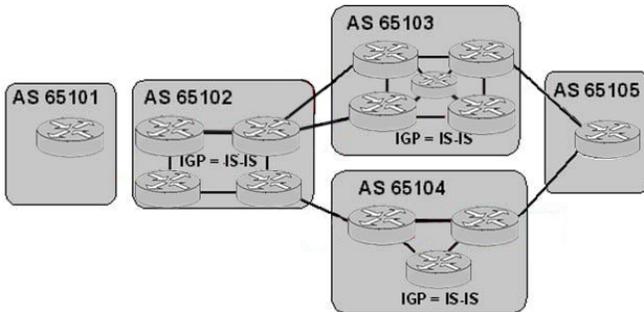
In reference to BGP route flap dampening feature, when will the penalty and the flap history of a route clear?

- A. When the penalty associated with a route reaches the reuse limit.
- B. When a route is in the history state for longer than the half-time limit.
- C. When the penalty associated with a route reaches the suppress limit.
- D. When a route in the suppressed state for longer than the half-time limit.
- E. When the penalty associated with a route drops below half the reuse limit.

Answer: E

QUESTION 70:

Exhibit



What is the limitation of BGP?

- A. AS 65101 cannot use BGP to connect to AS 65102.
- B. AS 65102 cannot use MED to influence the return traffic from AS 65103.
- C. AS 65102 cannot influence how AS 65103 will route traffic out to AS 65105.
- D. AS 65102 cannot load balance the traffic to AS 65105 via AS 65103 and AS 65104.
- E. Since AS 65102, 65103, and AS 65104 are Transit AS, they must enable synchronization.

Answer: C

QUESTION 71:

Which prefix-list statement can be used to deny all Class C Private IP addresses?

- A. ip prefix-list blkc seq 5 deny 192.168.0.0
- B. ip prefix-list blkc seq 5 deny 192.168.0.0/24
- C. ip prefix-list blkc seq 5 deny 192.168.0.0/32 le 32
- D. ip prefix-list blkc seq 5 deny 192.168.0.0/16 le 32

Answer: D

QUESTION 72:

Which MPLS VPN implementation allows selected sites in one simple VPN to communicate with selected sites of a second VPN?

- A. central services VPN
- B. managed CE router services VPN
- C. overlapping VPN
- D. managed PE router services VPN

Answer: C

QUESTION 73:

In a Transit AS, how do the internal routers within the Transit AS forward packets destined for the external networks using a scalable solution?

- A. Using the default route.
- B. Using the IGP routes where the external networks are redistributed into the IGP by the edge routers.
- C. Using the EBGp routes where the external networks are redistributed into the IBGP by the edge routers.
- D. Using the IBGP routes, then using recursive lookup based on IGP information to resolve the BGP next-hop.

Answer: D

QUESTION 74:

Which BGP dampening configuration parameter is not configurable by the user?

- A. half-time
- B. reuse-limit
- C. max-suppress
- D. suppress-limit
- E. per-flap-penalty

Answer: E

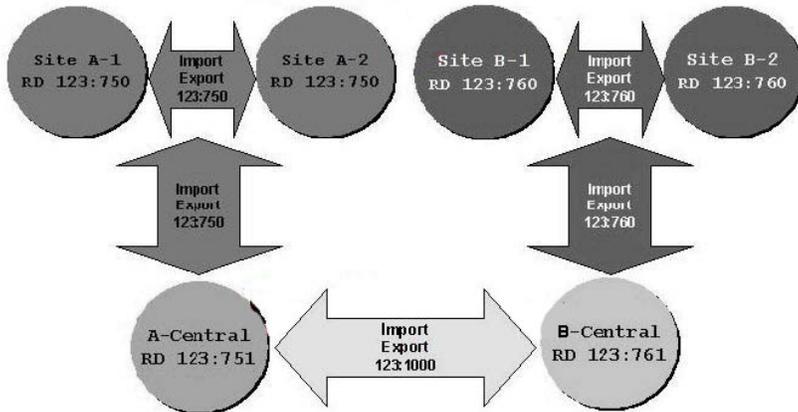
QUESTION 75:

In a central services topology, which routes do client VRFs contain?

- A. Routes from the client site, but not from the server site.
- B. Routes from the server site, but not from the client site.
- C. Routes from both the client site and the server site.
- D. Only EBGP routes from either the client site or the server site.

Answer: C

QUESTION 76:



In reference to the diagram depicting the flow of routing updates in an overlapping VPN, which statement is true?

- A. Sites A1 and A2 will import all networks with RTs 123:750 and 123:1000.
- B. Sites B1 and B2 will export all networks with RTs 123:760 and 123:1000.
- C. Sites A1 and A2 communicate with Sites B1 and B2 via A-Central and B-Central.
- D. Site A-Central exports RTs 123:750 and 123:1000.
- E. Site B-Central is exporting and importing RT 123:1000 so it can communicate with sites A1 and A2.

Answer: D

QUESTION 77:

Which prefix will be permitted by the following prefix-list?
ip prefix-list test permit 10.1.0.0/16 le 22

- A. 10.2.2.0 255.255.255.0
- B. 10.1.32.0 255.255.240.0
- C. 10.1.1.64 255.255.255.240
- D. 10.0.0.0 255.255.0.0
- E. 10.1.3.0 255.255.255.0

Answer: B

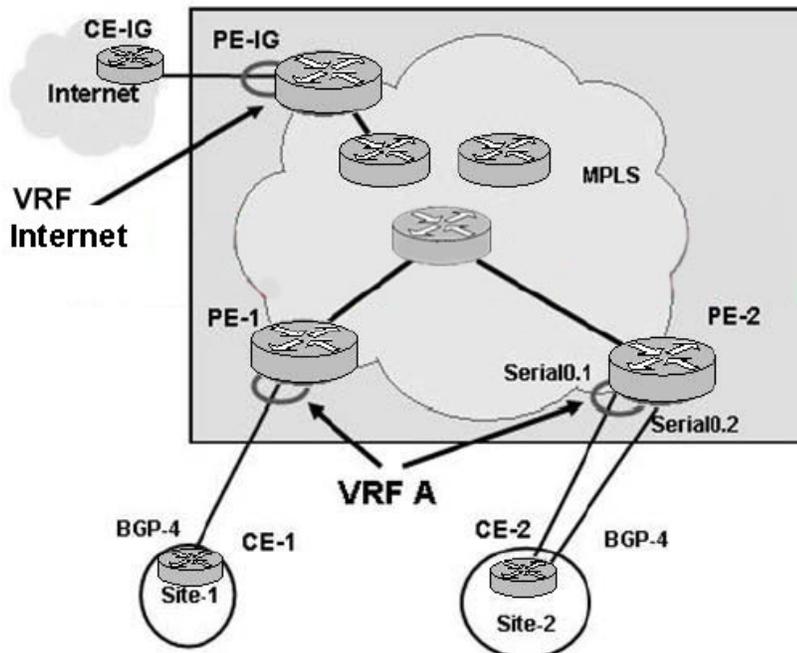
QUESTION 78:

Which component of MPLS architecture exchanges Layer 3 routing information and labels?

- A. control plane
- B. data plane
- C. forwarding plane
- D. routing plane

Answer: A

QUESTION 79:



In the diagram, the Internet is accessed through a dedicated Internet VPN implementation. Which routing table will PE-IG use to forward packets from the Internet to Site-2?

- A. Global routing table
- B. VPNv4 routing table
- C. VRF A routing table
- D. VRF Internet routing table

Answer: D

QUESTION 80:

What is the main reason for a multihomed customer to apply an outbound route

filter to filter the BGP updates from the customer router to the service provider router?

- A. To ensure that the return traffic into the customer network will be load balanced between the different service providers.
- B. To ensure that the outbound traffic from the customer network will be load balanced between the different service providers.
- C. To prevent the customer network from becoming a transit AS.
- D. To allow the customer network to become a transit AS.
- E. To reduce the size of the BGP table on the customer network internal (core) routers.
- F. To reduce the size of the BGP table on the service provider internal (core) routers.

Answer: C

QUESTION 81:

With MPLS VPNs, what is prepended to an IPv4 address to make it globally unique?

- A. route target (RT)
- B. route distinguisher (RD)
- C. VPNv4 header
- D. Site of origin (SOO) prefix
- E. Route tag

Answer: B

QUESTION 82:

By default, which BGP attribute is stripped in outgoing IBGP updates?

- A. next-hop
- B. as-path
- C. community
- D. local-preference
- E. origin

Answer: C

QUESTION 83:

Based on the following show ip bgp neighbor 10.1.1.1 output, what is the neighbor relationship?

```
CK1 #show ip bgp neighbor 10.1.1.1
BGP neighbor is 10.1.1.1, remote AS 65002, external link
Index 2, Offset 0, Mask 0x4
```

BGP version 4, remote router ID 12.1.2.3
Neighbor under common administration
BGP state = Established, table version = 5, up for 00:09:15
Last read 00:00:16, hold time is 180, keepalive interval is 60 seconds
Minimum time between advertisement runs is 30 seconds
Received 13 messages, 0 notifications, 0 in queue
Sent 13 messages, 0 notifications, 0 in queue
Prefix advertised 1, suppressed 0, withdrawn 0
Connections established 1; dropped 0
Last reset never
1 accepted prefixes consume 32 bytes
0 history paths consume 0 bytes
External BGP neighbor may be up to 255 hops away

- A. Neighbor 10.1.1.1 is a regular EBGP neighbor.
- B. Neighbor 10.1.1.1 is an Intra-Confederation IBGP neighbor.
- C. Neighbor 10.1.1.1 is an Intra-Confederation EBGP neighbor.
- D. Neighbor 10.1.1.1 is an IBGP neighbor that belongs to the same AS as CK1 .
- E. Neighbor 10.1.1.1 is an IBGP neighbor that belongs to the same Member As CK1 .

Answer: C

QUESTION 84:

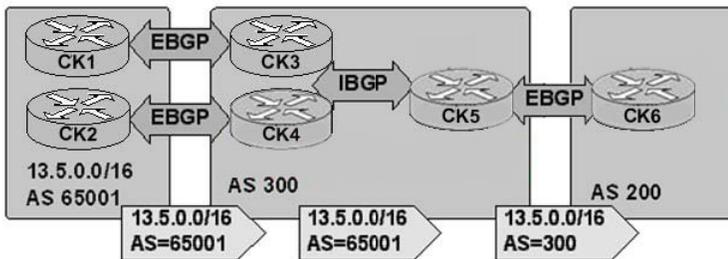
Which kind of link does a Layer 3 overlay VPN use?

- A. emulated point-to-point
- B. dedicated point-to-point
- C. point-to-multipoint
- D. permanent virtual circuits

Answer: A

QUESTION 85:

Exhibit



Complete the following BGP configuration on CK5 so that it will remove the 65001 private AS number before it sends the 13.5.0.0/16 update to AS 200.

Hostname CK5

```
!  
!Output omitted  
!  
router bgp 300  
neighbor 3.3.3.3 remote-as 300  
neighbor 4.4.4.4 remote-as 300  
neighbor 6.6.6.6 remote-as 300  
-----
```

- A. neighbor 6.6.6.6 allowas-out
- B. neighbor 6.6.6.6 local-as 300
- C. neighbor 6.6.6.6 remove-private-as
- D. neighbor 6.6.6.6 remove-private-out
- E. neighbor 6.6.6.6 as-prepend private-as

Answer: C

QUESTION 86:

Exhibit



You work as a network administrator at Certkiller .com. Study the exhibit carefully. Which two other BGP router commands are required on Certkiller 1 and Certkiller or on both so that Certkiller 2 can upload the input prefix-list named testing to Certkiller 1 to be used as an output filter? Select two.

- A. On Certkiller 1, use the neighbor 10.1.1.2 capability orf prefix-list receive command.
- B. On Certkiller 1, use the neighbor 10.1.1.2 capability orf prefix-list send command.
- C. On Certkiller 1, use the neighbor 10.1.1.2 prefix-list testing out command.
- D. On Certkiller 2, use the neighbor 10.1.1.2 capability orf prefix-list receive command.
- E. On Certkiller 2, use the neighbor 10.1.1.2 capability orf prefix-list send command.
- F. On Certkiller 3, use the neighbor 10.1.1.2 prefix-list testing out command.

Answer: A, E

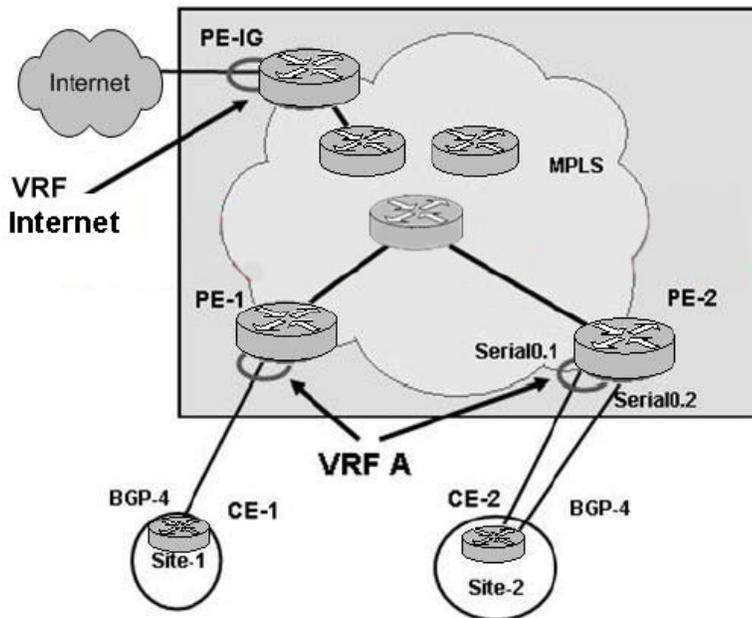
QUESTION 87:

Which two statements are true regarding the BGP network configuration command? (Choose two)

- A. When using the "network ip-prefix" BGP configuration command to advertise a major network into BGP, at least one of the subnets of the major-network must be present in the BGP table.
- B. When using the "network ip-prefix mask subnet-mask" BGP configuration command to advertise a classless prefix into BGP, the prefix must have an exact match in the IP routing table.
- C. When using the "network ip-prefix" BGP configuration command to advertise a major network into BGP, at least one of the subnets of the major-network must be present in the IP routing table.
- D. When using the "network ip-prefix mask subnet-mask" BGP configuration command to advertise a classless prefix into BGP, the prefix will be announced only if there is at least one network in the specified range in the IP routing table.

Answer: B, C

QUESTION 88:



In the case of Internet access through a dedicated sub-interface implementation, which routing table will PE-2 use to forward packets from Site-2 to the Internet?

- A. Global routing table
- B. VPNv4 routing table
- C. VRF A routing table
- D. CE-2 routing table
- E. PE-IG routing table

Answer: A

QUESTION 89:

Which BGP path attribute is only used within the local router and NOT shared between all the BGP routers within the AS?

- A. MED
- B. Origin
- C. Weight
- D. Local Preference

Answer: C

QUESTION 90:

Which two statements regarding route reflectors are true? (Choose two)

- A. A non-route reflector capable router cannot be a client.
- B. A hierarchical route reflector design is where a route reflector client is not directly connected to the route reflector.
- C. If a client has IBGP sessions to other clients in the same cluster, those clients will receive unnecessary duplicated BGP updates.
- D. If a client in one cluster has an IBGP session to a route reflector that belongs to different clusters, the clients in the other cluster will receive unnecessary duplicated BGP updates.

Answer: C, D

QUESTION 91:

What is wrong with the following BGP configuration?

```
!  
interface loopback0  
ip address 1.1.1.1 255.255.255.255  
!  
interface s0  
ip address 172.16.1.2 255.255.255.0  
!  
interface e1  
ip address 10.1.1.1 255.255.255.0  
!  
interface e2  
ip address 10.1.2.1 255.255.255.0  
!  
router bgp 65101  
neighbor 172.16.1.1 remote-as 65102  
neighbor 3.3.3.3 remote-as 65101
```

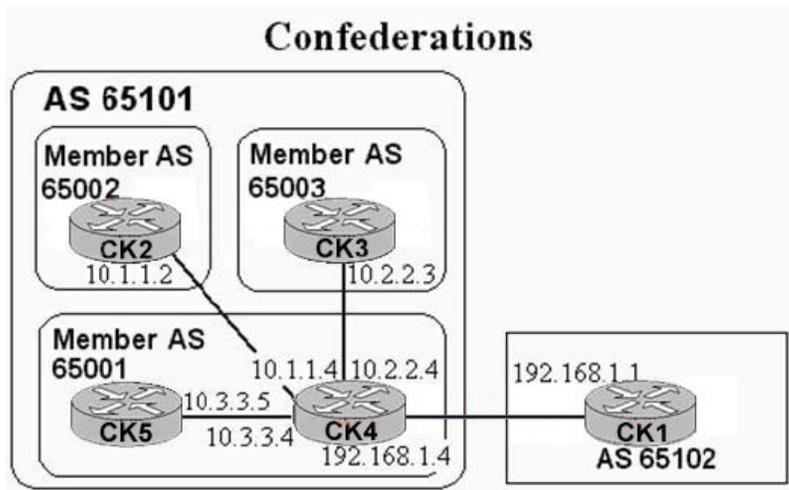
```
neighbor 3.3.3.3 update-source loopback0
network 10.1.0.0 mask 0.0.255.255
no auto-summary
!
router eigrp 100
network 10.0.0.0
network 1.0.0.0
```

- A. The network 10.1.0.0 mask 0.0.255.255 statement is wrong.
- B. The auto-summary configuration statement is missing under router bgp 65101.
- C. The neighbor 3.3.3.3 ebgp-multihop 2 configuration statement is missing under router bgp 65101.
- D. The neighbor 172.16.1.1 ebgp-multihop 2 configuration statement is missing under router bgp 65101.
- E. The neighbor 3.3.3.3 remote-as 65101 and the neighbor 3.3.3.3 update-source loopback0 configuration statements should be changed to neighbor 1.1.1.1 remote-as 65101 and neighbor 1.1.1.1 update-source loopback0.
- F. The loopback0 interface IP address should be 3.3.3.3 255.255.255.255.

Answer: A

QUESTION 92:

Exhibit



Given the following:

```
router bgp 65001
bgp confederation identifier 65101
-----
```

```
neighbor 10.3.3.5 remote-as 65001
neighbor 10.2.2.3 remote-as 65003
neighbor 10.1.1.2 remote-as 65002
neighbor 192.168.1.1 remote-as 65101
```

Based on the topology shown in the exhibit, what is the missing configuration

statement required on CK4 to complete the BGP confederation configuration?

- A. bgp confederation peers 65002 65003
- B. bgp confederation peers 10.2.2.3 10.1.1.2
- C. bgp confederation peers 65101 65002 65003
- D. bgp confederation peers 65001 65002 65003
- E. bgp confederation peers 10.3.3.5 10.2.2.3 10.1.1.2

Answer: A

QUESTION 93:

Which of the following statements about the BGP MED path attribute on Cisco routers is correct?

- A. If the MED is not attached to a BGP route, it defaults to 0 (best).
- B. By default, MED can be considered for routers coming from different autonomous systems.
- C. By default, MED can be considered for routers coming from different intrafederation autonomous systems.
- D. The bgp deterministic-med command is used to ensure that an accurate MED comparison is made across all routes received from different autonomous systems.
- E. The bgp deterministic-med command is used to allow the MED to determine the BGP path selection over the other BGP attributes.
- F. The bgp deterministic-med command is used to ensure that an accurate MED comparison is made across all routes received from the same autonomous system.

Answer: A

QUESTION 94:

Which of the following errors is the most common when configuring BGP routers within a transit AS?

- A. using BGP next-hop-self on the edge routers
- B. BGP to IGP redistribution errors
- C. IGP to BGP redistribution errors
- D. BGP synchronization enabled
- E. eBGP and iBGP administrative distance configuration errors

Answer: D

QUESTION 95:

Exhibit:

```
Certkiller 3(config-router)#neighbor 10.1.1.1 maximum-prefix 100 80 warning-only
```

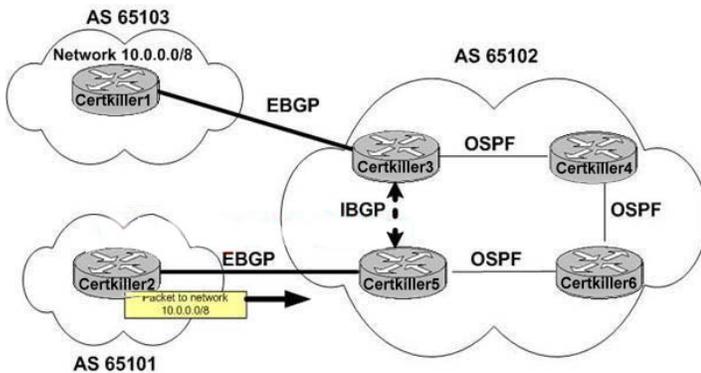
What does the command displayed in the exhibit accomplish?

- A. It prevents the 10.1.1.1 neighbor from exhausting the memory of the Certkiller 3 router.
- B. It limits the number of BGP Update messages the Certkiller 3 router will send to the 10.1.1.1 BGP peer.
- C. It logs a warning message on the Certkiller 3 router if the Certkiller 3 router BGP contains more than 100 prefixes.
- D. It logs a warning message on the Certkiller 3 router if the Certkiller 3 router BGP table is 80% full.
- E. It resets the BGP session to the 10.1.1.1 BGP peer of the Certkiller 3 router BGP table contains more than 100 prefixes.
- F. It resets the BGP session to the 10.1.1.1 BGP peer of the Certkiller 3 router BGP table is 80% full.

Answer: A

QUESTION 96:

Exhibit:



There is no direct physical connectivity between Certkiller 3 and Certkiller 5. The Transit AS 65102 is only running partial-meshed IBGP between Certkiller 3 and Certkiller 5 and is using OSPF as its IGP. Both Certkiller 3 and Certkiller 5 are using the next-hop-self options. The external BGP routes are not being redistributed into OSPF. BGP synchronization is disabled.

What happens to the packet from AS 65101 to AS 65103?

- A. Packets will be black-holed because it will be dropped by Certkiller 4.
- B. Packets will be black-holed because it will be dropped by Certkiller 5.
- C. Packets will be black-holed because it will be dropped by Certkiller 6.
- D. Packets can be transported via the Transit AS 65102 without any problems.

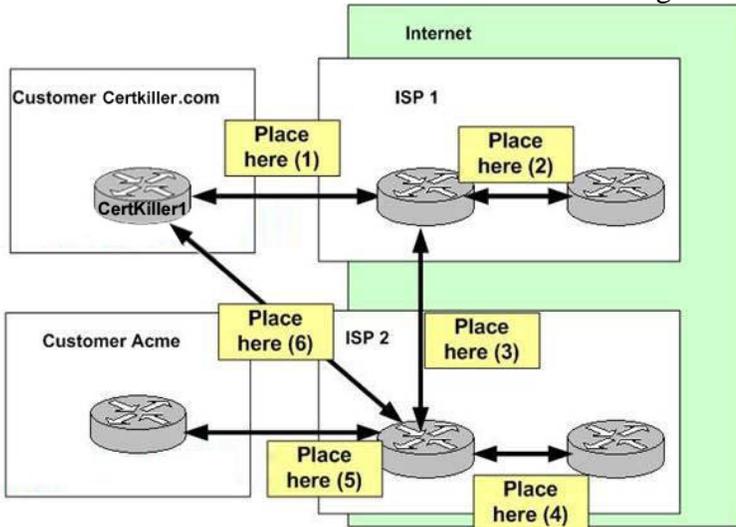
Answer: C

Explanation: Even with synchronization off testing 6 does not know how to route the packet because OSPF does not know about the route that BGP knows about. Synchronization being off assumes that there is a full mesh of BGP routers.

QUESTION 97:

DRAG DROP

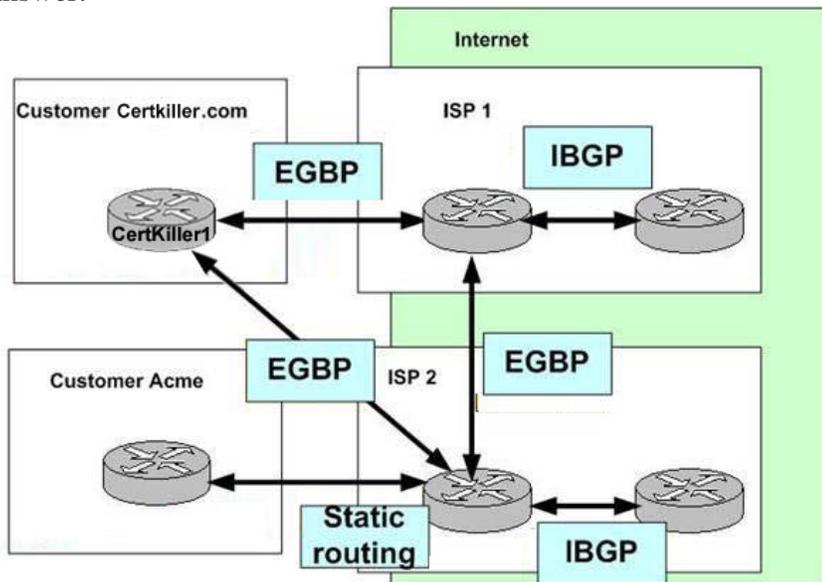
Identify the most appropriate method to connect the customers to the ISPs, to interconnect the two ISPs, and to interconnect the router within each ISP. Place the methods in their correct locations in the diagram.



Methods, Select from these

- EGBP
- IBGP
- Static routing

Answer:



Methods, Select from these

- EGBP
- IBGP
- Static routing

Explanation:

642-691

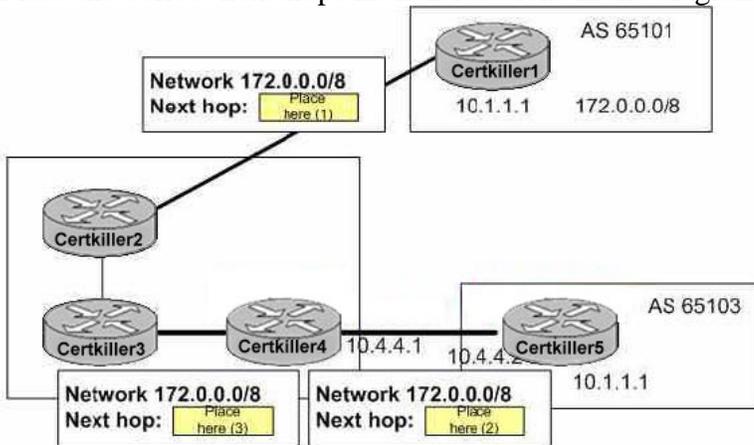
1. Box1 EBGP Certkiller .com 1 to ISP1. This is a multi-homed Internet connection. Most appropriate
2. Box 2IBGP Routers within ISP1
3. Box 3EBGP Between ISP1 and ISP2
4. Box 4 IBGP Routers within ISP 2
5. Box 5 Static Routing (Here is the catch). Because this is not a multi-homed connection the most appropriate method is static routes
6. Box 6 EBGP Certkiller .com 2 to ISP2. This is a multi-homed Internet connection. Most appropriate

QUESTION 98:

DRAG DROP

By default, NOT using the next-hop-self option, what is the next-hop for the 172.0.0.0/8 BGP updates shown in the diagram?

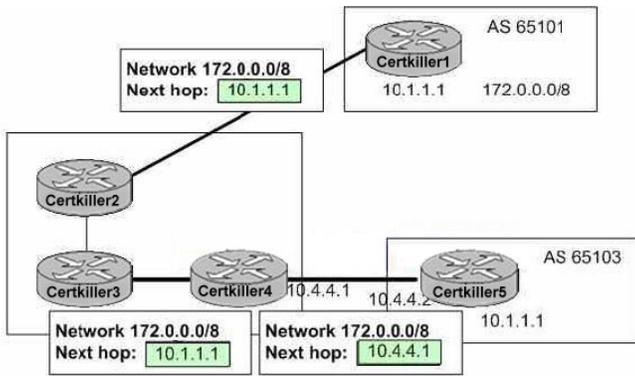
Place the correct next-hops in their locations in the diagram.



Next-hop, Select from these

- 10.1.1.1
- 10.1.1.2
- 10.2.2.1
- 10.2.2.2
- 10.3.3.1
- 10.3.3.2
- 10.4.4.1
- 10.4.4.2

Answer: :



Next-hop, Select from these

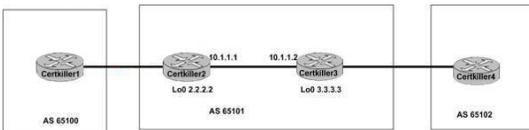
10.1.1.1 10.1.1.2 10.2.2.1 10.2.2.2
 10.3.3.1 10.3.3.2 10.4.4.1 10.4.4.2

Explanation:

1. Certkiller 1- Certkiller 2 Next Hop 10.1.1.1 This is an EBGP connection. Next hop is address of EBGP peer that advertised the route.
2. Certkiller 4- Certkiller 5 Next Hop 10.4.4.1 This is an EBGP connection. Next hop is address of EBGP peer that advertised the route.
3. Certkiller 3- Certkiller 4 Next Hop 10.1.1.1 This is an IBGP connection. Next hop is address of EBGP peer that advertised the route.

QUESTION 99:

Exhibit:'



Given the following configuration for Certkiller 2:

```
! output omitted
!
hostname Certkiller 2
!
interfaceloopback 0
ipaddress 2.2.2.2
!
interfacee0
ipaddress 10.1.1.1 255.255.255.0
!
routerbgp 65101
neighbor172.16.1.1 remote-as 65100
neighbor3.3.3.3 remote-as 65101
nosync
!
routereigrp 101
```

network 10.0.0.0

network 2.0.0.0

When Certkiller 2 sends the TCP SYN packet to Certkiller 3 to establish the IBGP session, what will be the source IP address of the TCP SYN packet from Certkiller 2 to Certkiller 3?

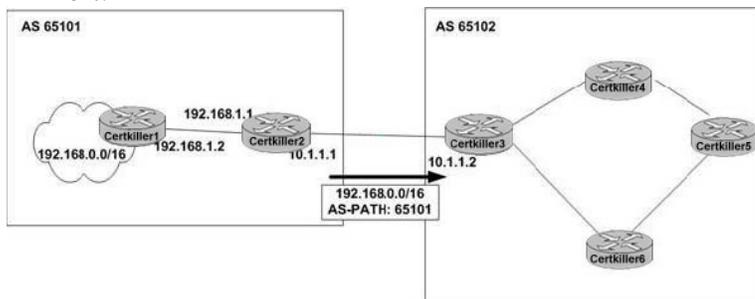
- A. 2.2.2.2
- B. 3.3.3.3
- C. 10.1.1.1
- D. 10.1.1.2

Answer: C

Explanation: Answer is 10.1.1.1 since neighbor 3.3.3.3 update-source lo0 was NOT entered.

QUESTION 100:

Exhibit:



What is the required configuration on Certkiller 2 to allow Certkiller 2 to announce the 192.168.0.0/16 prefix to Certkiller 3 via BGP?

- A.

```
router bgp 65101
neighbor 10.1.1.2 remote-as 65102
neighbor 192.168.1.2 remote-as 65101
network 192.168.1.0
auto-summary
```
- B.

```
router bgp 65101
neighbor 10.1.1.2 remote-as 65102
neighbor 192.168.1.2 remote-as 65101
network 192.168.0.0
!
```
- C.

```
router bgp 65101
neighbor 10.1.1.2 remote-as 65102
neighbor 192.168.1.2 remote-as 65101
network 192.168.0.0 mask 255.255.0.0
auto-summary
```

```
!  
iproute 192.0.0.0 255.0.0.0 192.168.1.2  
D. router bgp 65101  
neighbor10.1.1.2 remote-as 65102  
neighbor192.168.1.2 remote-as 65101  
network192.168.0.0 mask 255.255.0.0  
!  
iproute 192.168.0.0 255.255.0.0 192.168.1.2
```

Answer: D

Explanation: Without knowing the IGP, we must ensure that an exact route into the routing table for the network advertised.

Wrong Answer

1. A Auto summary would only summarize 192.168.1.0/24 and not /16
2. B Will only advertise 198.168.0.0/24 and not /16
3. C Need an exact match on route in Routing table and network command

QUESTION 101:

Based on the following show ip bgp neighbor 10.1.1.1 output:

```
R1#show ip bgp neighbors 10.1.1.1  
BGP neighbor is 10.1.1.1, remote AS 65103, external link  
BGP version 4, remote under ID 0.0.0.0  
BGP state = Active  
Last read 00:00:04, hold time is 180, keepalive interval is  
60 seconds  
Received 44 messages, 0 notifications, 0 in queue  
Sent 45 messages, 6 notifications, 0 in queue  
Route refresh request: received 0, sent 0  
Default minimum time between advertisement runs in 30  
seconds  
For address family: IPv4 Unicast  
BGP table version 1, neighbor version 0  
Index 2, Offset 0, Mask 0x4  
0 accepted prefixes consume 0 bytes  
Prefix advertised 0, suppressed 0, withdraw 0  
Number of NLRI in the update sent: max 0, min 0  
Connections established 7; dropped 7  
Last reset 00:05:18, due to BGP Notification received, peer  
in wrong AS  
External BGP neighbor may be up to 2 hops away.  
No active TCP connection  
What is causing the BGP session to the 10.1.1.1 neighbor to toggle between the Idle and  
Active state?
```

- A. There is an AS number configuration error.
- B. The BGP neighbor 10.1.1.1 is not reachable.
- C. The EBGP-multihop value for neighbor 10.1.1.1 is set to the default value.
- D. The BGP session is using the loopback interface but the update source is not set to specify the loopback interface.

Answer: A

Explanation: The key is looking at the last section of the output. We specified the neighbor in AS 65103 but received an error (Notification) that it is not in that AS.

QUESTION 102:

Which forward mechanism pre-builds the complete IP forwarding table (FIB) based on the IP routing table?

- A. CEF Switching
- B. Fast Switching
- C. Process Switching
- D. Optimum Switching

Answer: A

Explanation: CEF is the most advanced form and build a separate FIB. See [cef.htm#5757](#)

QUESTION 103:

Exhibit:

```
R1#show ip bgp
BGP table version is 1, local router ID is 10.0.0.1
Status codes: s suppressed, d damped, h history, *valid,>best,iinternal
Origin codes: i - IGP, e - EGP,? - incomplete
Network Next Hop Metric LocPrf Wight Path
*i133.3.0.0 3.3.3.3 0 100 0 65101 i
*i172.0.0.0/8 20.1.1.5 0 100 0 65102 i
R1#show ip route
Codes: C - connected, S - static, I - IGRP, R -RIP, M -mobile, B
- BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter
area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate
default
```

U - per-user static route, o - ODR, P - periodic downloaded static route

T - traffic engineered route

Gateway of last resort is not set

C 30.0.0.0/8 is directly connected, FastEthernet0/0

C 10.0.0.0/8 is directly connected, Loopback111

C 40.0.0.0/8 is directly connected, Serial0/0

Based on the exhibits show ip bgp and show ip route output, what is the most likely problem that causes the two IBGP routes NOT to be selected as the best route in the BGP table?

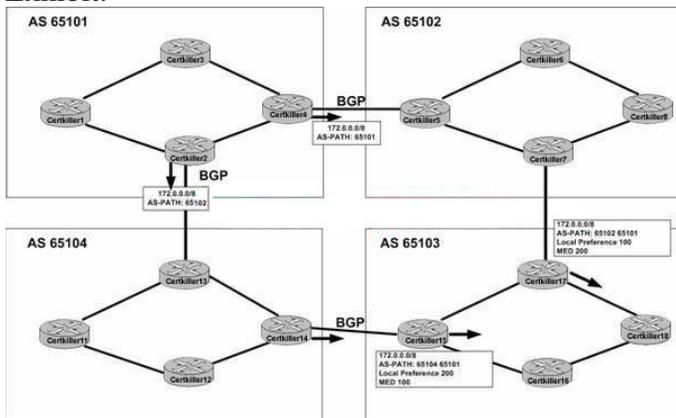
- A. The Weight is 0.
- B. The origin code is "i,".
- C. The Metric (MED) is 0.
- D. BGP synchronization is disabled.
- E. The BGP next-hop is not reachable.

Answer: E

Explanation: One of the first criteria in BGP to accept a route is that the next hop must be reachable.

QUESTION 104:

Exhibit:



The best path from Certkiller 18 (AS 65103) to the 172.0.0.0/8 prefix (AS 65101) is _____.

- A. Via Certkiller 15 to AS 65104 then to AS 65101 because it has the best MED.
- B. Via Certkiller 17 to AS 65102 then to AS 65101 because it has the best MED.
- C. Load balanced between Certkiller 15 and Certkiller 17 due to the equal AS-Path length.
- D. Via Certkiller 15 to AS 65104 then to AS 65101 because it has the best Local Preference.
- E. Via Certkiller 17 to AS 65102 then to AS 65101 because it has the best Local Preference.

Answer: D

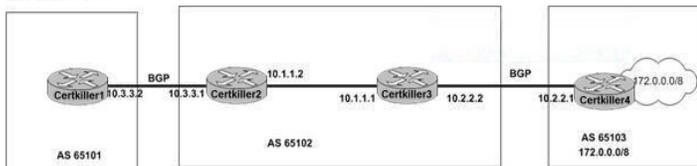
Explanation: BGP selects routes based on Higher Local Preference and then lower MED. LP is use to route to a designation. MED is used for path back.

Wrong Answers

1. A. MED is not used for path there
2. B. MED is not used for path there
3. C Not load balanced because Local Preference takes precedence. BGP only hand over best path to IP by default. multi-path would have to be specified for the chance to load balance
4. E Best local preference is the one with the highest value.

QUESTION 105:

Exhibit:



What is the proper BGP configuration on Certkiller 3 to have Certkiller 3 announce the 172.0.0.0/8 prefix from Certkiller 4 to Certkiller 2 via BGP with a next hop of 10.1.1.1?

- A. router bgp 65102
neighbor 10.2.2.1 remote-as 65103
neighbor 10.1.1.2 remote-as 65102
- B. router bgp 65102
neighbor 10.2.2.1 remote-as 65103
neighbor 10.1.1.2 remote-as 65102
neighbor 10.2.2.1 next-hop-self
- C. router bgp 65102
neighbor 10.2.2.1 remote-as 65103
neighbor 10.1.1.2 remote-as 65102
neighbor 10.1.1.2 next-hop-self
- D. router bgp 65102
neighbor 10.2.2.1 remote-as 65103
neighbor 10.1.1.2 remote-as 65102
neighbor 10.1.1.2 update-source-self

Answer: C

Explanation: Causes next-hop to be outgoing interface to neighbor 10.1.1.2

Wrong Answer

1. A is missing Next-hop-self command
2. B This is not done at the EBGp level

3. D update-source-self invalid parameter update-source address is the proper neighbor parameter

QUESTION 106:

What are two methods used to configure the local-preference BGP path attribute on Cisco routers? (Choose two)

- A. Using a route-map with the set local-preference value command.
- B. Using the bgp default local-preference value BGP configuration command.
- C. Using the default-metric local-preference value BGP configuration command.
- D. Using the neighbor ip-address local-preference value BGP configuration command.

Answer: A, B

Explanation:

- 1. A Valid command , See BGP command reference guide
- 2. B Valid command , See BGP command reference guide

Wrong Answer

- 1. C Invalid command
 - 2. D Invalid command
-

QUESTION 107:

BGP table

Address	Prefix	AS-Path	Next hop		Other attr.
				Communities	
10.0.0.0	/8	65100	1.1.1.1	65101:111	
		65101			
...

IP routing table

Protocol	Address	Prefix	Next-hop	Outgoing interface
BGP	10.0.0.0	/8	1.1.1.1	---

OSPF	1.1.1.1	/32	172.16.1.2	Ethernet 0
conn.	172.16.1.0	/24	---	Ethernet 0

To forward a packet to 10.0.0.0/8, the router perform the following steps:

Step 1. Search the ip routing table for a route to reach the 10.0.0.0/8 network.

Step 2. _____

Step 3. Find the connected outgoing interface to reach 172.16.1.2.

Step 4. Arp for the 172.16.1.2 MAC address if it is not already in the ARP cache.

Step 5. Store the 172.16.1.2 MAC address in the Fast Switching cache for successive packets to network 10.0.0.0.

What is step 2?

- A. Search the BGP table for an IGP route to reach the BGP next-hop 1.1.1.1.
- B. Search the BGP table for an IBGP route to reach the BGP next-hop 1.1.1.1.
- C. Search the IP routing table for an IGBP route to reach the BGP next-hop 1.1.1.1.
- D. Search the IP routing table for an IGP route to reach the BGP next-hop of 1.1.1.1.

Answer: D

Explanation: Routing table specifies to get to 10.0.0.0 goto 1.1.1.1. To get to 1.1.1.1 do a lookup in the routing table on the next hop.

QUESTION 108:

What type of BGP path attributes are propagated to BGP peers with the partial bit set even if they are not recognized?

- A. Optional Transitive
- B. Well Known Mandatory
- C. Optional Non-Transitive
- D. Well Known Discretionary

Answer: A

Explanation: Optional transitive. Attributes are passed on even if not recognized.

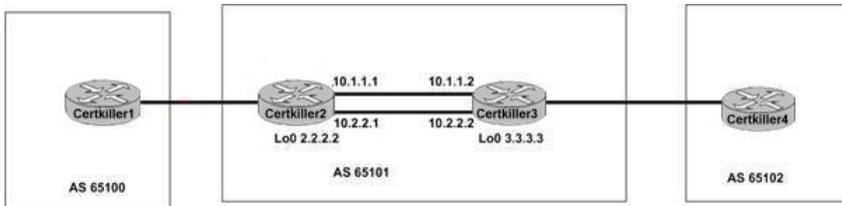
Wrong Answer

1. B. Well known mandatory must be recognized
2. C. Non-transitive would not be propagated
3. D. Well known discretionary must be recognized

QUESTION 109:

Exhibit:

642-691



Given the following configurations, Certkiller 2 and Certkiller 3 are not able to successfully establish the IBGP session using the loopback 0 interface.

What could be the cause of this problem?

!

! output omitted

!

```
hostname Certkiller 2
```

```
interface lo 0
```

```
ip address 2.2.2.2
```

!

```
interface e0
```

```
ip address 10.1.1.1 255.255.255.0
```

```
noshut
```

!

```
interface e1
```

```
ip address 10.2.2.1 255.255.255.0
```

```
noshut
```

!

```
router bgp 65101
```

```
neighbor 172.16.1.1 remote-as 65100
```

```
neighbor 3.3.3.3 remote-as 65101
```

```
neighbor 3.3.3.3 update-source loopback0
```

!

```
router eigrp 101
```

```
network 10.0.0.0
```

!

! output omitted

!

```
hostname Certkiller 3
```

!

```
interface lo 0
```

```
ip address 3.3.3.3
```

!

```
interface e0
```

```
ip address 10.1.1.2 255.255.255.0
```

```
noshut
```

!

```
interface e1
```

```
ip address 10.2.2.2 255.255.255.0
```

```
noshut
```

!

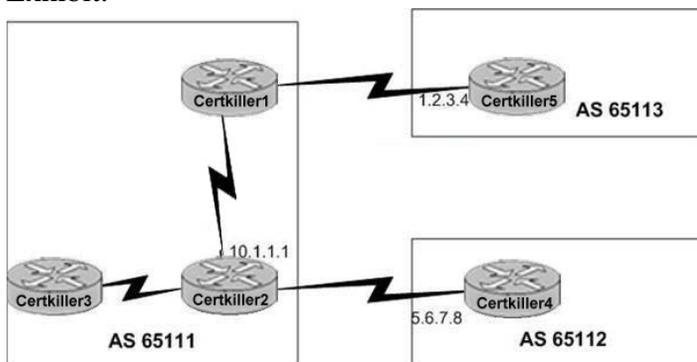
```
routerbgp 65101
neighbor192.168.1.1 remote-as 65102
neighbor2.2.2.2 remote-as 65101
neighbor2.2.2.2 update-source loopback0
!
routereigrp 101
network10.0.0.0
!
```

- A. The "No Sync" BGP configuration command is missing.
- B. Certkiller 2 and Certkiller 3 are using the loopback0 IP address as the source address for the BGP messages to each other.
- C. The "network 2.0.0.0" EIGRP configuration command is missing on Certkiller 2 and the "network 3.0.0.0" EIGRP configuration command is missing on Certkiller 3.
- D. The "Neighbor 2.2.2.2 ibgp-multihop 2" BGP configuration command is missing on Certkiller 3 and the "neighbor 3.3.3.3 ibgp-multihop 2" BGP configuration command is missing on Certkiller 2.

Answer: B

QUESTION 110:

Exhibit:



In the exhibit, how does the following configuration on Certkiller 1 affect the AS 65111 traffic flow?

```
hostname Certkiller 1
!
!output omitted
!
routerbgp 65111
neighbor10.1.1.1. remote-as 65111
neighbor10.2.2.2. remote-as 65111
neighbor1.2.3.4 remote-as 65113
neighbor1.2.3.4 route-map L2M in
!
route-mapL2M permit 10
setlocal-preference 101
```

- A. Inbound traffic from AS 65111 will prefer the path via AS 65113 over the path via AS 65112.
- B. Inbound traffic from AS 65111 will prefer the path via AS 65112 over the path via AS 65113.
- C. Outbound traffic from AS 65111 will prefer the path via AS 65113 over the path via AS 65112.
- D. Outbound traffic from AS 65111 will prefer the path via AS 65112 over the path via AS 65113.

Answer: C

Explanation: Default local preference is 100. Highest local preference is preferred. Routes advertised from AS 65113 are given a higher local preference. Local Preference applies to outbound traffic only.

QUESTION 111:

Given the following:

```
routerbgp 65102
neighbor 10.1.1.1 remote-as 65103
neighbor 10.2.2.2 remote-as 65104
neighbor 172.16.1.4 remote-as 65102
neighbor 10.1.1.1 route-map setlp in
!
route-mapsetlp permit 10
match community 1
set local-preference 150
!
route-mapsetlp permit 20
!
ipcommunity-list 1 permit 65103:150
What will the configuration on R1 accomplish?
```

- A. Routes from AS 65103 marked with a Community of 65103: 150 will have the default Local Preference; all other routes will have a Local Preference of 150.
- B. Routes from AS 65103 marked with a Community of 65103: 150 will have a Local Preference of 150; all other routes will have the default Local Preference.
- C. Routes from AS 65103 and AS 65104 marked with a Community of 65103: 150 will have the default Local Preference; all other routes will have a Local Preference of 150.
- D. Routes from AS 65103 and AS 65104 marked with a Community of 65103: 150 will have a Local Preference of 150; all other routes will have the default Local Preference.

Answer: B

Explanation: Route Map applies to only BGP routes from AS 65103. Local

Preference is increased if the community string is set. Route Map is only applied to 10.1.1.1 neighbor (AS 65103).

QUESTION 112:

Complete the BGP configuration on Certkiller 1 so that the show ip bgp output reflects the following path selection?

```
Certkiller 1#show ip bgp
BGP table version is 9, local router ID is 172.16.10.1
Status codes: s suppressed, d damped, h history, *valid, >
best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
Network Next Hop Metric LocPrf Weight Path
*> 172.20.0.0 10.10.20.1 150 65123 65122 i
* 10.10.10.2 0 120 65122 i
*> 172.25.0.0 10.10.21.1 150 65123 65124 i
* 10.10.10.2 120 65122 65124 i
*> 172.30.0.0 10.10.21.1 0 150 65123 i
* 10.10.10.2 120 65122 65123 I
hostname Certkiller 1
!
router bgp 65121
neighbor 10.10.10.2 remote-as 65122
neighbor 10.10.20.1 remote-as 65123
```

-
- A. neighbor 10.10.20.1 metric 0
 - B. neighbor 10.10.10.2 metric 0
 - C. neighbor 10.10.10.2 weight 120
neighbor 10.10.20.1 weight 150
 - D. neighbor 10.10.10.2 local-preference 0
neighbor 10.10.20.1 local-preference 100

Answer: C

Explanation: As per output the weight value is modified.

Wrong Answer

1. A. does not match output
 2. B does not match output
 3. D no local preference appears to have been used
-

QUESTION 113:

What is the correct BGP configuration to have AS 65111 prefer the path to AS 65221 over the path to AS 65237 except for the traffic which originates in AS 65237?

```
A. router bgp 65111
neighbor 10.10.10.10 remote-as 65221
neighbor 10.20.20.20 remote-as 65237
neighbor 10.10.10.10 weight 160
neighbor 10.20.20.20 weight 150
neighbor 10.20.20.20 filter-list 1 weight 170
!
ip as-path access-list 1 permit _65237$

B. router bgp 65111
neighbor 10.10.10.10 remote-as 65221
neighbor 10.20.20.20 remote-as 65237
neighbor 10.10.10.10 weight 160
neighbor 10.20.20.20 weight 150
neighbor 10.20.20.20 access-group 1 weight 170
!
access-list 1 permit _65237$

C. router bgp 65111
neighbor 10.10.10.10 remote-as 65221
neighbor 10.20.20.20 remote-as 65237
neighbor 10.10.10.10 weight 160
neighbor 10.20.20.20 weight 150

D. router bgp 65111
neighbor 10.10.10.10 remote-as 65221
neighbor 10.20.20.20 remote-as 65237
neighbor 10.10.10.10 weight 160
neighbor 10.20.20.20 weight 150
neighbor 10.20.20.20 route-map test
!
route-map test permit 10
match as-path _65221$
```

Answer: A

Explanation: Command requires the use of as path _65327\$ mean originating from 65237

Wrong Answer

1. B Needs the filter-list parameter to be use
2. C Will not prefer traffic originating from AS 65327 to go to AS 65327
3. D You are supposed to prefer traffic originating from AS 65237 and not AS 65221

QUESTION 114:

The purpose of the bgp always-compare-med command is to compare the MED _____.

- A. Even if the routes originate in different ASs.
- B. Even if the routes originate from the same AS.

- C. Even if the routes have different AS-Path length.
- D. Last, after all other path attributes are compared.
- E. First, before any other path attributes are compared.

Answer: A

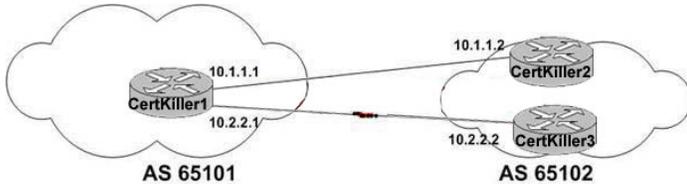
Explanation: See command reference guide. Command allow the comparison of MED for paths from neighbors in different AS.

Wrong Answer

- 1. B not the purpose of always-compare-med.
- 2. C,D,E is not relevant here

QUESTION 115:

Exhibit:



Complete the following Certkiller 1 BGP configuration to load balance the traffic from AS 65101 to AS 65102 if two identical EBGP routes are learned from AS 65102.

```
hostname Certkiller 1
!  
routerbgp 65101  
neighbor10.1.1.2 remote-as 65102  
neighbor10.2.2.2 remote-as 65102
```

- A. variance 2
- B. maximum-paths 2
- C. neighbor 10.1.1.2 weight 100
neighbor10.2.2.2 weight 100
- D. neighbor 10.1.1.2 ebgp-multihop 2
neighbor10.2.2.2 ebgp-multihop 2
- E. neighbor 10.1.1.2 local-preference 100
neighbor10.2.2.2 local-preference 100

Answer: B

Explanation: Maximum-paths 2 is required to load balance traffic between the two AS's.

Wrong

- 1. A Variance 2 is an EIGRP command to balance between over two BGP sessions
- 2. C Similar weight is required to load balance
- 3. D Multihop is not required
- 4. E Local preference needs to be similar to allow load balancing

QUESTION 116:

Which two statements are true about route reflector configurations? (Choose two)

- A. Clients must have IBGP sessions to the other clients and to the route reflector.
- B. The neighbor ip-address route-reflector-client configuration command is required on the route reflectors.
- C. When a cluster contains multiple route reflectors, all the route reflectors in the cluster need to be configured with the same Cluster-ID.
- D. When a cluster contains multiple route reflectors, the bgp cluster-id cluster-id configuration command is required on the clients on the route reflectors.

Answer: B, C

Explanation:

1. B. Route reflectors are configured by deciding who the clients are
2. C For redundancy Multiple Route reflectors can be in a cluster. They need to know what there cluster ID is to make this the routers know there peers.

Wrong Answer

1. A Clients do not have sessions to other clients. This defeats the purpose of the route reflector
2. D Clients do not know that they are clients

QUESTION 117:

Based on the following show ip bgp neighbors output from Certkiller 2, which two statements are true? (Choose two)

```
Certkiller 2#show ip bgp neighbors
BGP neighbor is 10.1.1.5, remote AS 65101, internal link
Index 1, Offset 0, Mask 0x2
Route-Reflector Client
BGP version 4, remote router ID 5.5.5.5
BGP state = Established, table version = 4, up for
00:01:29
Last read 00:00:29, hold time is 180, keepalive interval
is 60 seconds
Minimum time between advertisement runs is 5 seconds
Received 13556 messages, 0 notifications, 0 in queue
Sent 13454 messages, 0 notifications, 0 in queue
Prefix advertised 0, suppressed 0, withdraw 0
Default weight 900
Inbound path policy configured
Route map for incoming advertisement is test
Connections established 31; dropped 30
Last reset 00:01:40, due to User reset
```

3 accepted prefixes consume 96 bytes
0 history paths consume 0 bytes
0 history paths consume 0 bytes

- A. Certkiller 2 is a route reflector client of the 10.1.1.5 neighbor.
- B. Certkiller 2 is a route reflector and the 10.1.1.5 neighbor is a client of Certkiller 2.
- C. Certkiller 2 has successfully established the IBGP session with the 10.1.1.5 neighbor.
- D. All the BGP updates from Certkiller 2 to the 10.1.1.5 neighbor must go through the route-map called "test".

Answer: B, C

Explanation:

- 1. B The neighbor does not know that it is a route reflector client.
- 2. C Certkiller 2 established the IBGP session with 10.1.1.5 . Must be IBGP because we are using route reflectors.

Wrong Answer

- 1. A. Certkiller 2 is a reflector and not a client
- 2. D. BGP updates from 10.1.1.5 not Certkiller 2 must go through route map test. Key this route map is used for incoming advertisement.

QUESTION 118:

Within a Transit AS, when is it necessary to redistribute the connected interface into IGP at the edge routers?

- A. When synchronization is enabled.
- B. When the edge router is a route reflector.
- C. When the edge routers are redistributed the BGP into IGP.
- D. When the edge routers are not using the next-hop-self option.
- E. When the edge routers are not using loopback interface to establish the IBGP sessions.

Answer: D

Explanation: When not using the next-hop-self option the next hop is the address of the EBGP peer. The routing table must have a route to the next-hop or the route is not considered. By redistributing connected into the IGP insures that you can get to the next hop neighbor.

Wrong Answer

- 1. A Synchronization ensures that the routes of the EGP match the IGP.
- 2. B this has no bearing
- 3. C Generally this is insane
- 4. E You still need to be sure that you have a route back to the loopback interface

QUESTION 119:

Which optional configuration parameter is required on a static route so a route-map can be used to match the static route?

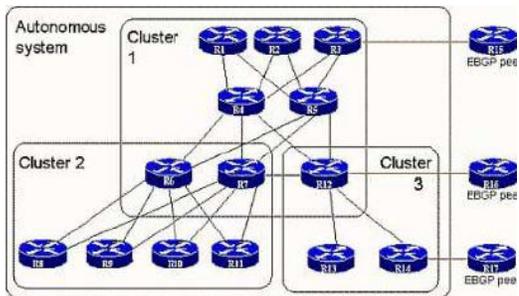
- A. Tag
- B. Distance
- C. Community
- D. Permanent

Answer: A

Explanation: Command is ip route prefix mask { address | interface } [distance] [tag tag][permanent]. Where tag can be used a a match value in route-maps

QUESTION 120:

Exhibit:



Based on the network topology shown in the exhibit, which routers should be configured as route reflectors but also act as clients to other router reflectors to provide the most scalable solution?

- A. R1, R2, and R3
- B. R6, R7, and R12
- C. R3, R12, and R14
- D. R1, R2, R3, R4, and R5
- E. R4, R5, R6, R7, and R12

Answer: B

Explanation:

Key is which routers are route reflectors and clients at the same time

1. B R6, R7, and R12 are both clients and route reflectors

Wrong Answer

1. A R1, R2, R3 would not route reflectors

2. C makes no sense, they are all the EBGP directly connect peers.

3. D no sense

4. E R4 and R5 are route reflectors but not clients as well. The best way to look at this is to look at each cluster individually to determine which should be route reflectors

QUESTION 121:

Which two statements about route reflectors are true? (Choose two)

- A. A non-route reflector capable router can be a client.
- B. Clients do not need IBGP sessions to all the route reflectors in the same cluster.
- C. A hierarchical route reflector design is where a route reflector is also client of another route reflector.
- D. Clients should have IBGP sessions to route reflectors that belong to different clusters to provide redundancy.

Answer: A, C

Explanation:

1. A Route reflector clients never know that they are clients
2. C Route reflectors can be clients to other route reflectors

Wrong Answers

1. B Clients need to peer with all route reflectors in a cluster
 2. D Clients should only connect to route reflectors in the cluster they belong to
-

QUESTION 122:

Besides BGP policy scaling and IBGP mesh scaling, service providers need to concentrate on controlling _____.

- A. The AS-Path length by using AS-PATH prepending.
- B. The BGP query process by using route summarization.
- C. The number of routes being redistributed from BGP into the IGP using route-maps.
- D. The number of routes being redistributed from the IGP into the BGP using route-maps.
- E. The number of BGP updates and the size of the BGP forwarding tables using route aggregation.

Answer: E

QUESTION 123:

Which Cisco IOS command is used to view the amount of CPU resources consumed (utilization) due to running BGP processes?

- A. show ip bgp process
- B. show memory | include BGP
- C. show process cpu | include BGP

- D. show ip protocols | include BGP
- E. show cpu utilization | include BGP

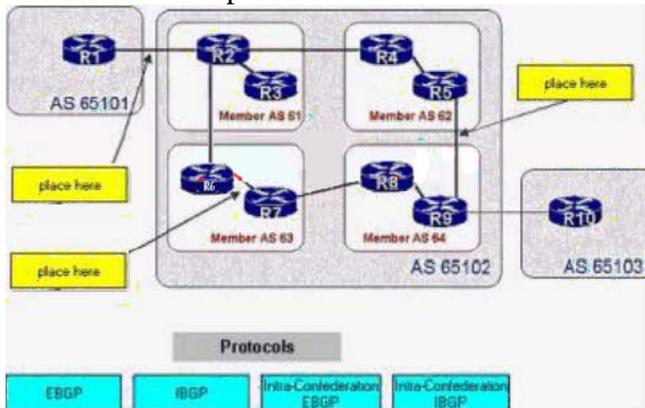
Answer: C

Explanation: show process cpu is the command used to show cpu resources of all processes. Piping to include show only BGP's usage.

QUESTION 124:

DRAG DROP

Place the correct protocols in their locations on the diagram.



Answer:

Explanation:

From TOP Left going clockwise

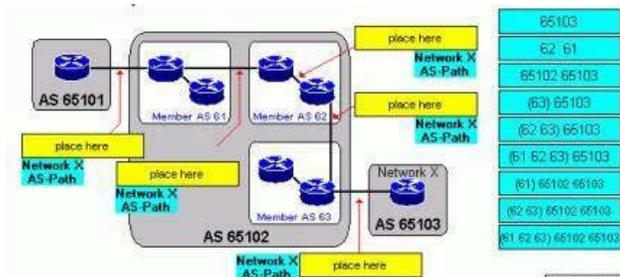
1. Box 1 EBGP R1-R2
2. Box 2 Intra confederation EBGP
3. Box 3 IBGP

QUESTION 125:

DRAG DROP

Network X originates in AS 65103. Select the correct AS Path information for the Network X update and drag it into the appropriate box.

Some of the AS-Paths listed will not be used.



Answer:

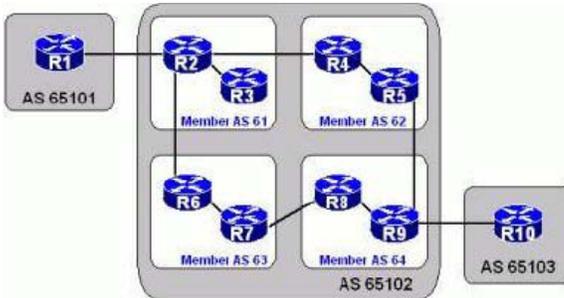
Explanation:

From Top left going clockwise

1. Box 1 65102 65103 R1-R2
2. Box 2 (62 63) 65103 AS 61-AS 62
3. Box 3 (63) 65103 Confed AS 62
4. Box 4 65103

QUESTION 126:

Exhibit:



How many and what kind of neighbors are configured on R2?

- A. 1 EBGP neighbor, 3 IBGP neighbors
- B. 3 EBGP neighbors, 1 IBGP neighbor
- C. 3 EBGP neighbors, 1 Intra-Confederations IBGP neighbor
- D. 1 EBGP neighbor, 3 Intra-Confederations EBGP neighbors
- E. 1 EBGP neighbor, 1 IBGP neighbor, 2 Intra-Confederations EBGP neighbors

Answer: E

Explanation: R2 - R1 EBGP, R2 - R6 Confed EBGP, R2 - R4 Confed EBGP, R3 IBGP. = 1 EBGP, 2 confed EBGP, 1 IBGP

QUESTION 127:

Which three statements about an EBGP session or an IBGP session are true? (Choose three)

- A. EBGP uses AS-Path to detect routing information loops.
- B. BGP route selection rules favor IBGP routes over EBGP routes.
- C. No BGP attributes are changed in IBGP updates except for the next-hop attribute if next-hop-self is configured.
- D. Routes learned from an IBGP peer are not advertised to another IBGP peer to prevent routing information loops within the AS.
- E. EBGP uses split horizon to prevent routing information loops; routes learned from an EBGP peer are not advertised to another EBGP peer.

Answer: A, C, D

Explanation:

1. A EBGp uses AS-Path to ensure that it is not passing its own route back to itself
2. C No BGP attributes are changed
3. D Router learned via IBGP are not advertised to prevent loops

Wrong

1. B IBGP routes have admin distance of 200 and EBGp of 20. 20 is preferred
2. E routes learned from EBGp peers are advertised to other EBGp peers.

QUESTION 128:

Given the following:

```
routerbgp 65123
neighborcustomers peer-group
neighbor10.1.1.2 remote-as 65213
neighbor10.1.1.2 peer-group customers
neighbor10.2.2.2 remote-as 65314
neighbor10.2.2.2 peer-group customers
neighbor10.2.2.2 route-map test2 in
```

Based on the BGP configuration for a router in AS 65123, which two statements are true?
(Choose two)

- A. The peer group is named "customers".
- B. The two EBGp peer belongs to the same peer-group.
- C. The route-map named "test2" will be applied both EBGp peers.
- D. The peer group configuration is used to reduce the full mesh BGP requirements.
- E. AS 65123, AS 65314 and AS 65213 belong to the same BGP confederation named "customers".

Answer: A, B

Explanation:

1. A The peer group name is "customers" (configuration statement #2 creates the peer group)
2. B neighbors 10.1.1.2 and 10.2.2.2 belong to the peer group customers (config statements #4 & #6 define who belongs to peer group)

Wrong Answers

1. C test2 is not applied to the peer group but only 10.2.2.2
2. D nothing to do with full mesh. Configuration is used to simplify policy
3. E We are talking about peer groups and not confederations

QUESTION 129:

A Transit AS contains 10 routers. To enable full mesh IBGP within that Transit AS, how many IBGP sessions are required?

- A. 10
- B. 28
- C. 45
- D. 50
- E. 100

Answer: C

Explanation: full mesh = $(n)(n-1)/2 = 10*9/2 = 45$

QUESTION 130:

What is the purpose of the following BGP configuration command?

neighbor 2.2.2.2 maximum-prefix 1000

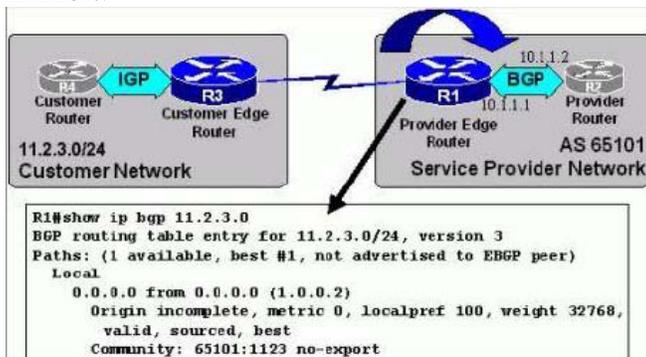
- A. It limits the number of prefixes that can be sent to neighbor 2.2.2.2 to 1000.
- B. It limits the number of prefixes that can be sent to and received from the 2.2.2.2 neighbor to 1000.
- C. If the number of prefixes sent to the 2.2.2.2 neighbor exceeds 650, the router starts to generate a warning message.
- D. If the number of prefixes received from the 2.2.2.2 neighbor exceeds 650, the router starts to generate a warning message.
- E. If the number of prefixes received from the 2.2.2.2 neighbor exceeds 1000, the neighbor relationship to 2.2.2.2 will be dropped.

Answer: E

Explanation: Maximum-prefix control how many prefixes can be received from a neighbor. Default warning is 75%, peering will terminated if the maximum is exceeded by default.

QUESTION 131:

Exhibit:



Based on the show ip bgp 11.2.3.0 output shown in the exhibit, which two statements are true? (Choose two)

- A. The route will not be sent to any EBGP neighbors.
- B. The route is using the default Weight and Local Preference values.
- C. The route will not be propagated to any IBGP or EBGP peers because it has the "no-export" community.
- D. The origin code is "incomplete" because the network command was used to advertise the route via BGP.

Answer: A, B

Explanation:

1. A Route will not be sent to any EBGP neighbor, no-export ensures this

2. B Default local preference is 100

Wrong

1. C no-export means do not advertise to EBGP peer. no-advertise means do not advertise to any peer.

2. D incomplete is if IGP is redistributed into BGP.

QUESTION 132:

Which statement is true about BGP peer group configurations?

- A. The use of peer group will increase the CPU load of the BGP process.
- B. Peer group is used to reduce the full mesh IBGP requirements for Transit AS.
- C. Peer group is used to sub-divide an AS into multiple sub-ASs to increase scalability.
- D. The use of peer groups allows the router to build BGP update only once for the entire peer group.
- E. IBGP and EBGP neighbors can be assigned to the same peer groups as long as all peer group members have a common outbound policy.

Answer: D

Explanation: One update table is created for all members of a peer group.

Wrong Answer

1. A Peer groups reduce CPU, because less tables are maintained

2. B Peer groups do not change full mesh requirement

3. C Confederations are used to subdivide an AS

QUESTION 133:

What are the two methods for configuring BGP route dampening? (Choose two)

- A. Using the bgp dampening command.
- B. Using a route-map with the set dampening command.
- C. Using the neighbor ip-address maximum prefix command.
- D. Using the neighbor ip-address unsuppress-map command.
- E. Using the neighbor ip-address soft-configuration command.

Answer: A, B

Explanation:

1. A bgp dampening [half-life reuse suppress max-suppress-time][route-map map]

2. B bgp dampening [half-life reuse suppress max-suppress-time

Wrong

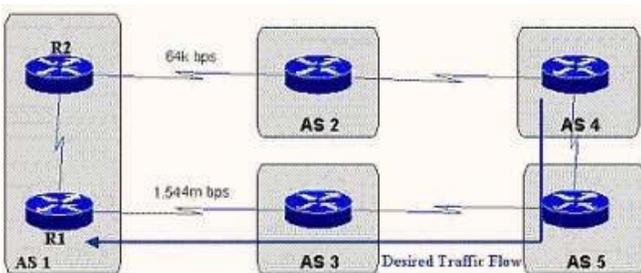
1. C used for max prefixes

2. D not a command

3. E control refresh

QUESTION 134:

Exhibit:



AS-Path prepending is used in AS1 in order to influence the return traffic from AS 4 to AS 1 through the higher speed path via AS 3. _____ needs to be configured for AS-Path prepending and a minimum of _____ of the AS number should be prepended.

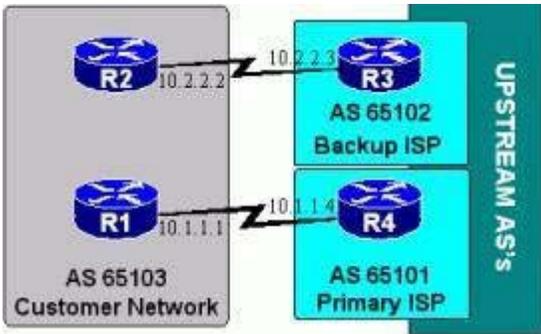
- A. R1; one copy
- B. R2; one copy
- C. R1; two copies
- D. R2; two copies
- E. R2; three copies

Answer: D

Explanation: For AS 4 to prefer AS 3 path to AS 1, the path needs to be shorter than through AS 2. AS 4 to AS 1 through AS 2 path is "2 1". AS4 through AS 3 to AS 1 is "5 3 1". Need to prepend AS twice on R2 to make it less desirable.

QUESTION 135:

Exhibit:



R2 is configured to prepend AS number 65102 in updates to R3.
Which statement is true?

- A. The BGP updates from R2 to R3 will be rejected by R3 due to loop prevention.
- B. The EBGP session between R2 and R3 will be dropped due to an AS-Path prepend configuration error.
- C. The AS-Path length via AS 65102 will increase to influence the return traffic path selected by the remote ASs.
- D. The AS-Path length via AS 65102 will decrease to influence the return traffic path selected by the remote ASs.
- E. The AS-Path prepend configuration command using the remote AS number (65102) will be rejected by IOS at R2.

Answer: A

Explanation: R2 is prepending the AS to which it is peering with. BGP prevents loops by not accepting routes from its own AS.

QUESTION 136:

Every time a flap occurs on a route, the route receives _____.

- A. 750 per-flap penalty points which are user configurable.
- B. 1500 per-flap penalty points which are user configurable.
- C. 2000 per-flap penalty points which are user configurable.
- D. 1000 per-flap penalty points which are not user configurable.
- E. 2000 per-flap penalty points which are not user configurable.

Answer: D

Explanation: Everytime a flap occurs a penalty of 1000 is made and is not configurable. See sections on route dampening in Internet routing architectures guide.

QUESTION 137:

DRAG DROP

642-691

Complete the following BGP configuration so that all prefixes announced to neighbor 10.1.1.1 will have a MED of 100. Drag the required configuration statements into the box.

The possible configuration statements are listed in order. Follow the same order when placing your selections into the box. Some of the statements will not be used.

<pre>! output omitted ! router 65123 neighbor 10.1.1.1 remote-as 65111</pre>	<pre>neighbor 10.1.1.1 route-map 1 in neighbor 10.1.1.1 route-map 1 out neighbor 10.1.1.1 filter-list 1 in neighbor 10.1.1.1 filter-list 1 out neighbor 10.1.1.1 metric 100 route-map 1 permit 10 match any set metric 100</pre>
--	--

Answer:

Explanation:

1. neighbor 10.1.1.1 route-map 1 out
2. route-map 1 permit 10
3. set metric 100

Wrong

1. match any (not required)
2. neighbor 10.1.1.1 metric 100 (no such command)
3. neighbor 10.1.1.1 route-map 1 in (wrong direction)
4. filter list can not set MED

QUESTION 138:

AS 65101 is a dual home customer. When starting the EBGP sessions with both ISPs announcing its networks, which two statements are true about the default behavior of BGP? (Choose two)

- A. The default BGP route selection prefers the shortest AS-Path.
- B. If the AS-Paths are equal in length, BGP prefers the route from the peer with the highest router ID.
- C. The bandwidth available to reach the ISPs is not taken into consideration in the BGP path selection process.
- D. Using the default BGP "maximum-paths" of 4, BGP can load balance the outbound traffic from the multihomed customer between the two ISPs.

Answer: A, C

Explanation: Route selection is > W, > LP, local, shortest AS, < origin cod i<e<?, < MED, ext over interior, closest neighbor, IP address, BGP router ID

1. A default route selection is shortest AS-path
2. C bandwidth is not taken into consideration at all

Wrong

1. B Lowest RID
2. D default is one

QUESTION 139:

Which regular expression can be used to match all routes originating in the local AS?

- A. *\$
- B. .\$
- C. ^\$
- D. .*
- E. ^*

Answer: C

Explanation: ^\$ ^ begin path, \$ end path. with nothing in between means now AS's pass through it, hence, local.

QUESTION 140:

Exhibit:

Certkiller router#show ip route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2, EEGP

i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, *

-candidate default

U - per-user static route, o - ODR

T - traffic engineered route

Gateway of last resort is not set

172.16.0.0/24 is subnetted, 2 subnets

B 172.16.10.0 [20/0] via 10.1.1.100, 00:00:24

B 172.16.11.0 [20/0] via 10.1.1.100, 00:00:24

172.26.0.0/28 is subnetted, 3 subnets

B 172.26.1.48 [200/0] via 192.168.1.50, 00:00:31

B 172.26.1.32 [200/0] via 192.168.1.50, 00:00:31

B 172.26.1.16 [200/0] via 192.168.1.50, 00:00:31

10.0.0.0/8 is variable subnetted, 2 subnets, 2 masks

B 10.0.0/8 [20/0] via 10.1.1.100, 00:00:24

C 10.1.1.0/24 is directly connected, Serial3

192.168.1.0/28 is subnetted, 3 subnets

C 192.168.1.32 is directly connected, Serial1

C 192.168.1.48 is directly connected, Serial2
C 192.168.1.16 is directly connected, Serial0
192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
B 192.168.2.64/28 [20/0] via 10.1.1.100, 00:00:26
Based on the show ip route output in the exhibit, how can you tell if a BGP route is learned via IBGP or EBGP?

- A. By the Metric
- B. By the Next Hop
- C. By the Admin Distance
- D. By the code "B" or "B Ex"

Answer: C

Explanation: admin distance. Distance of EBGP is 20, Distance of IBGP is 200.

QUESTION 141:

The Outbound Route Filter (ORF) capability is negotiated between BGP neighbors during the BGP ____ process via the ____ message.

- A. route propagation; Update
- B. session establishment; Open
- C. route propagation; Notification
- D. route propagation; Route Refresh
- E. session establishment; Route Refresh

Answer: B

Explanation: Capabilities are negotiated during session establishment in the open message. Appendix C P456 Internet routing Architectures.

Wrong

1. A needs to be decided before update
 2. C notification is an error message
 3. D not a message type
 4. E not a message type
-

QUESTION 142:

Which condition must be met before Route Refresh messages can be sent to a BGP peer?

- A. The soft reconfiguration feature must be enabled.
- B. The BGP peer must advertise Route Refresh capability.
- C. The Outbound Route Filter (ORF) feature must be enabled.
- D. The BGP peer must have enough memory to store a copy of the original update.

Answer: B

Explanation: The peer must announce it is capable of refresh. Without this capability soft config or refresh could not be enabled.

Wrong

1. A Not necessary for outbound. Most important to have capability
2. C ORF optimizes the performance of refresh
3. D Copy is stored local, ORF will cause update to be stored on peer

QUESTION 143:

How would you complete the following BGP configuration to allow the BGP router in AS 65111 to only accept BGP updates from AS 65101 having 172.16 in the first 16 bits and a subnet mask of 24 bits or less? (Any other updates from AS 65101 are denied.)

```
routerbgp 65111
neighbor10.1.1.1 remote-as 65101
```

- A. neighbor 10.1.1.1 access-list 1 in
access-list1 permit 172.16.0.0/24
- B. neighbor 10.1.1.1 access-list 1 in
access-list1 permit 172.16.0.0 0.0.255.255
- C. neighbor 10.1.1.1 access-list 1 in
access-list1 permit 172.16.0.0 255.255.0.0
- D. neighbor 10.1.1.1 prefix-list test in
ipprefix-list test seq 5 permit 172.16.0.0&17 le 24
- E. neighbor 10.1.1.1 prefix-list test in
ipprefix-list test seq 5 permit 172.16.0.0/16 ge 24

Answer: No Right Answer

Explanation: Answer should be ip prefix-list test 5 permit 172.16.0.0/16 le 24.
172.16.0.0 /16 is the correct network length statement. le 24 is the mask to permit
This is a cross between D and E

QUESTION 144:

DRAG DROP

What prefixes will be matched by the prefix-list statements?
Place each prefix next to the corresponding prefix-list statement.

Place here	Prefix-list Statement	Prefix
Place here	ip prefix-list certkiller1 seq 5 permit 0.0.0.0/0 ge 32	all routers
Place here	ip prefix-list certkiller2 seq 5 permit 0.0.0.0/0 ie 32	all host routes
Place here	ip prefix-list certkiller3 seq 5 permit 0.0.0.0/0	just the default route
Place here	ip prefix-list certkiller4 seq 5 permit 0.0.0.0/1 le 24	any prefix in class-A address space covering at least 256 addresses

Answer:

Place here	Prefix-list Statement	Prefix
all host routes	ip prefix-list certkiller1 seq 5 permit 0.0.0.0/0 ge 32	
all routers	ip prefix-list certkiller2 seq 5 permit 0.0.0.0/0 ie 32	
just the default route	ip prefix-list certkiller3 seq 5 permit 0.0.0.0/0	
any prefix in class-A address space covering at least 256 addresses	ip prefix-list certkiller4 seq 5 permit 0.0.0.0/1 le 24	

QUESTION 145:

What does the route map named test in the following BGP configuration accomplish?

```
routerbgp 65111
neighbor10.1.1.1 remote-as 65112
neighbor10.1.1.1 route-map test out
!
ipas-path access-list 1 permit_65222$
ipas-path access-list 2 permit.*
!
route-maptest permit 10
matchas-path 1
setmetric 200
!
route-maptest permit 20
matchas-path 2
```

- A. Marks all prefixes advertised to the 10.1.1.1 neighbor with a MED of 200.
- B. Marks all prefixes received from the 10.1.1.1 neighbor with a MED of 200.
- C. Marks all prefixes originating in autonomous system 65222 advertised to the 10.1.1.1 neighbor with a MED of 200.
- D. Marks all prefixes originating in autonomous system 65222 received from the 10.1.1.1 neighbor with a MED of 200.

Answer: C

Explanation: _65222\$. _ space, \$ end of path. The first route map statement does this. Route map is applied out bound to 10.1.1.1

QUESTION 146:

Given the following:

The customer (AS 65100) is advertising the 172.0.0.0/8 prefix over the upper link without using AS-Path prepending.

The customer (AS 65100) is advertising the 172.0.0.0/8 prefix with a prepended AS-Path of AS 65100 65100 65100 65100 over the lower link.

Which two statements are true about the incoming traffic to 172.0.0.0/8? (Choose two)

- A. The lower link will be preferred over the upper link.
- B. The upper link will be preferred over the lower link.
- C. The upper link will act as a backup for the lower link.
- D. The lower link will act as a backup for the upper link.
- E. The upper and lower links will be used simultaneously for load balancing

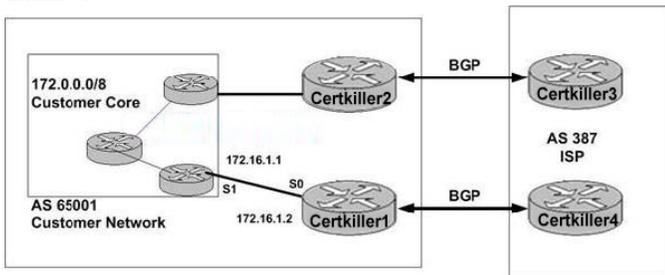
Answer: B, D

Explanation:

1. B The upper link is preferred over the lower link, since, the shortest AS-PATH is preferred
2. D The lower link will act as a backup. If the upper link fails then the lower link will be the preferred path

QUESTION 147:

Exhibit:



Which two configuration commands will complete the BGP configuration on Certkiller 1 so it will conditionally announce the 172.0.0.0/8 to Certkiller 4 via BGP? (Choose two) hostname Certkiller 1

```
!  
!output omitted  
!  
1. _____  
!  
routerbgp 65001  
neighbor 172.16.1.1 remote-as 65001  
neighbor 2.2.2.2 remote-as 65001  
neighbor 4.4.4.4 remote-as 387  
!
```

2. _____
!

- A. 2.network 172.16.0.0
auto-summary
- B. 2. network 172.0.0.0 mask 255.0.0.0
- C. 1.ip route 172.0.0.0 255.0.0.0 null0
- D. 1.ip route 172.0.0.0 255.0.0.0 null0 255
- E. 1.ip route 172.0.0.0 255.0.0.0 172.16.1.1
- F. 2.aggregate-address 172.0.0.9 mask 255.0.0.0

Answer: B, E

Explanation:

- 1. 1 E, conditional static route disappears when the link to 172.16.1.1 disappears
- 2. 2B, correct network command, correct mask, will not advertise when link goes down.

Wrong

- 1. A wrong mask need /8 this defaults to /16
- 2. C This is unconditional
- 3. D This is an unconditional floating static route
- 4. F should be 172.0.0.0 not 172.0.0.9

QUESTION 148:

Which two statements are true about multihomed customers? (Choose two)

- A. The customer usually uses provider assigned (PA) address space.
- B. The customer usually uses provider independent (PI) address space.
- C. Static routing is always adequate between the customer and service provider.
- D. BGP is only needed if physical link failures on the primary link cannot be detected.
- E. The customer runs BGP with the service provider using either a private or a public AS number.

Answer: B, E?

Explanation:

- 1. B use PI when possible so route aggregation of ISP is not prevented
- 2. E Can be either and private or public AS number. However, the private number is not common practice. See ISP essentials.

Wrong

- 1. A Provider assigned is not common for multi-home, common for Single connected
- 2. C Static routing is not appropriate when multi-homing to multiple ISP's
- 3. D multi-homing can be used for load balancing.

QUESTION 149:

In a multihomed environment with two ISP connections, which two statements are true?
(Choose two)

- A. The customer should not be configured to act as a transit AS between the two ISPs.
- B. It is recommended that the multi-homed customer use a registered (public) AS number.
- C. AS-Path prepending can be configured on the customer's edge router to influence the BGP path selection process for the outbound traffic (traffic from the customer to the ISPs).
- D. The customer can use Local Preference on the customer's edge routers to influence the BGP path selection process for the inbound traffic (traffic from the ISPs to the customer).
- E. The advertisement of customer's IP address space can be conditioned by the customer's edge routers by using a static route to the null0 interface and by using the proper network statement under router bgp.

Answer: A, E

Explanation:

1. A Customers do not want to become a transit for ISP's
2. E This a valid way to advertise routes to the ISP

Wrong

1. B ISP essentials says most popular, does not say recommend
2. C affects inbound traffic
3. D local preference is for outbound traffic

QUESTION 150:

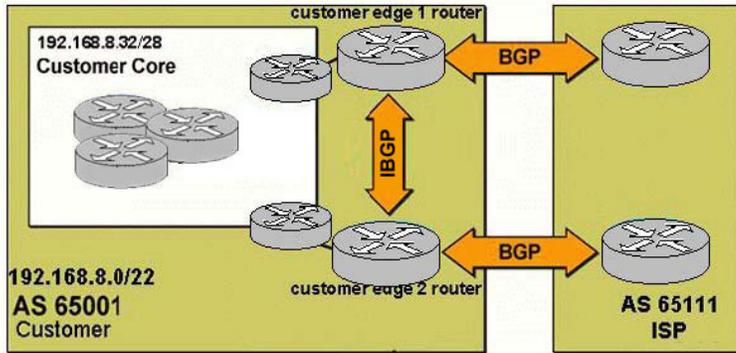
Within a transit AS, which router(s) are typically configured with the BGP next-hop-self option?

- A. All edge routers running both eBGP and iBGP sessions.
- B. All internal routers running iBGP sessions.
- C. The internal routers configured as route reflectors and running iBGP sessions to the client routers.
- D. The internal routers configured as route reflector clients and running iBGP sessions to the route reflector.
- E. All BGP routers within the transit AS.

Answer: A

QUESTION 151:

Exhibit



Based on the network diagram shown in the exhibit, what is the correct configuration on the customer edge router used to conditionally announce the customer networks to the ISP?

A. router bgp 65001

! neighbor commands not shown
network 192.168.0 mask 255.255.252.0

!

ip route 192.168.8.0 255.255.252.0 192.168.8.33

B. router bgp 65001

! neighbor commands not shown
aggregate-address 192.168.8.0 255.255.252.0 summary-only

!

ip route 192.168.8.0 255.255.252.0 192.168.8.23

C. route bgp 65001

! neighbor commands not shown

network 192.168.8.0

network 192.168.9.0

network 192.168.10.0

network 192.168.11.0

!

ip route 192.168.8.0 255.255.255.0 null0

ip route 192.168.9.0 255.255.255.0 null0

ip route 192.168.10.0 255.255.255.0 null0

ip route 192.168.11.0 255.255.255.0 null0

D. router bgp 65001

! neighbor commands not shown

aggregate-address 192.168.8.0 255.255.252.0 summary-only

!

router ospf 1

network 192.168.8.0 0.0.3.255 area 0

E. router bgp 65001

! neighbor commands not shown

aggregate-address 192.168.8.0 255.255.252.0

!

ip route 192.168.0 255.255.252.0 null0

Answer: A

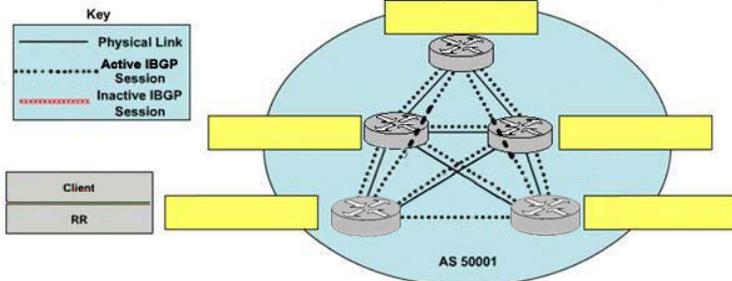
QUESTION 152:

DRAG DROP

Based on the physical topology of AS 50001, to reduce the IBGP full-mesh requirements, only two redundant route reflectors (RR) will be used in a single cluster. In this case, which two routers should be the RRs and which three routers should be the clients based on the recommended route reflector design rule?

Task: 1) Drag the matching text labels given on the left side of the topology to the correct targets on the right to identify these routers.

2) Also, identify all the unnecessary IBGP sessions once the redundant RRs have been configured by clicking on the IBGP session links to disable them. You can re-enable the IBGP sessions link by clicking on it again.



Answer:

QUESTION 153:

What is the effect on existing BGP connections when Path MTU (PMTU) discovery is enabled?

- A. They dynamically renegotiate the new MTU size.
- B. After timing out, they will inherit the configured MTU size.
- C. The connections are reset and negotiate MTU size.
- D. There is no effect on existing connections.

Answer: D

QUESTION 154:

Which show command can be used to display the originator ID and cluster-list?

- A. show ip bgp
- B. show ip bgp sum
- C. show ip route bgp
- D. show ip route {prefix}
- E. show ip bgp {prefix}

F. show ip bgp neighbors {ip address}

Answer: A

QUESTION 155:

CK4 >show ip bgp 172.0.0.0

BGP routing table entry for 172.0.0.0/8, version 4

Paths: (1 available, best #1, table Default-IP-Routing-Table)

Advertised to non peer-group peers:

192.168.1.6

50111

2.2.2.2 (metric 2323456) from 3.3.3.3 (3.3.3.3)

Origin IGP, metric 0, localpref 100, valid, internal, best

Originator: 2.2.2.2, Cluster list: 3.3.3.3

Based on the shown ip bgp 172.0.0.0 output shown, which statement is true about the 172.0.0.0/8 prefix?

- A. It is received from an EBGP neighbor.
- B. It is received from an intra-confederation EBGP neighbor.
- C. It is received from an intra-confederation IBGP neighbor.
- D. It is a prefix locally originated by CK4 .
- E. It is received from a route reflector.

Answer: E

QUESTION 156:

What does the following command accomplish?

Certkiller #clear ip bgp 10.1.1.1 in prefix-filter

- A. The Certkiller router will perform an outbound soft reconfig to the 10.1.1.1 neighbor.
- B. The Certkiller router will send out the ORF prefix-list so that a new route refresh will be received from the 10.1.1.1 neighbor.
- C. The 10.1.1.1 router will perform an inbound soft reconfig on the updates from the Certkiller neighbor.
- D. The 10.1.1.1 router will send out the ORF prefix-list so that a new route refresh will be received from the Certkiller neighbor.
- E. The bgp session between the Certkiller and the 10.1.1.1 router will be reset so that all the new bgp updates from the 10.1.1.1 router can be processed by the inbound prefix-list at the Certkiller router.
- F. The bgp session between the Certkiller and the 10.1.1.1 router will be reset so that all the new bgp updates from the Certkiller router can be processed by the inbound prefix-list at the 10.1.1.1 router.

Answer: B

QUESTION 157:

When BGP is configured for aggregate addressing, which attribute is set by default unless the as-set keyword is specified?

- A. summary-only
- B. atomic aggregate
- C. AS-path
- D. Aggregate route

Answer: B

QUESTION 158:

What is one method an ISP can use to reduce the BGP requirements within the ISP network core?

- A. Use MPLS within the ISP network core.
- B. Use MPLS to connect to multihomed customers.
- C. Carry the external routes using an IGP like IS-IS.
- D. Use recursive routing to reach the BGP next hop.
- E. Redistribute the external routes into MPLS.

Answer: A

QUESTION 159:

Examine the following route-map configuration:

```
route-map test permit 10
match as-path 10
set community no-export additive
```

Which of the statements below is correct?

- A. Any route matched by as-path access-list 10 will not be advertised to eBGP peers.
- B. Any route matched by as-path access-list 10 will have its BGP community attribute stripped off.
- C. Any route matched by as-path access-list 10 will have its BGP community attribute stripped off when the route is propagated to eBGP peers.
- D. Any route matched by as-path access-list 10 will have its BGP community overwritten with the no-export community.

Answer: A

QUESTION 160:

Within a transit AS, which routes should the IGP carry if the edge routers are using next-hop-self?

- A. The core subnets and loopbacks within the transit AS.
- B. The external subnets used in the EBGP sessions with the neighboring ASs.
- C. The external subnets used in the EBGP sessions with the neighboring ASs and the core subnets and loopbacks within the transit AS.
- D. The external routes to remote ASs.
- E. The external routes to remote ASs and the core subnets and loopbacks within the transit AS.

Answer: A

QUESTION 161:

Which two command scenarios will alter the BGP local-preference attribute to 120? (Choose two)

- A. router(config-route-map)#set local-preference 120
- B. router(config-router)#neighbor ip-address local-preference 120
- C. router(config-router)#bgp default local-preference 120
- D. router(config-router)#neighbor ip-address prefix-list test in local-preference 120

Answer: A, C

QUESTION 162:

```
PIR1#show route-map test
route-map test, permit, sequence 10
  Match classes:
    ip address prefix lists.test
    as-path (as-path filter): 1 2
    community (community-list filter): 1 2
  Set clauses:
    weight 150
Policy routing matches: 0 packets, 0 bytes
```

Based on the show route-map test output shown in the exhibit, which matching routes will be set to a weight of 150?

- A. Routes that match the prefix-list test AND also match either the as-path filter 1 OR 2 AND also match either the community-list 1 OR 2
- B. Routes that match the prefix-list named test OR match the as-path filter 1 AND 2 OR match the community-list 1 AND 2
- C. Routes that match the prefix-list named test AND match the as-path filter 1 AND 2

AND match the community-list 1 AND 2
D. Routes that match the prefix-list named test
OR match the as-path filter 1 OR 2
OR match the community-list 1 OR 2

Answer: A

QUESTION 163:

How do BGP confederations reduce the requirements for a full-mesh iBGP configuration?

- A. By splitting the autonomous system into a number of smaller autonomous systems and using iBGP sessions among them.
- B. By modifying the iBGP split horizon rule.
- C. By modifying the eBGP split horizon rule.
- D. The combining multiple iBGP sessions into a single iBGP session over loopback interface.
- E. By synchronizing the iBGP updates to the IGP updates.

Answer: A

QUESTION 164:

Which one of the following regarding the BGP route refresh feature is correct?

- A. Negotiates feature capability during the BGP UPDATE message exchange process.
- B. Stores a copy of all the inbound BGP routing information in router memory.
- C. Requires preconfiguration using the neighbor soft-reconfiguration command on the router and the BGP peer.
- D. Allows a router to request the retransmission of BGP routes from a BGP neighbor.

Answer: D

QUESTION 165:

To establish a BGP session between two BGP speakers, what will be the first BGP message sent by each BGP speaker after the transport protocol connection is established?

- A. hello
- B. query
- C. open
- D. sync
- E. update
- F. notification

Answer: C

QUESTION 166:

Which two commands scenarios will alter the BGP weight attribute to 120? (Choose two)

- A. router(config-route-map)# set weight 120
- B. router(config-router)# neighbor ip-address weight 120
- C. router(config-router)# bgp default weight 120
- D. router(config-router)# neighbor ip-address prefix-list test in weight 120

Answer: A, B

QUESTION 167:

Which two BGP prefix-update optimization strategies can be used to minimize the impact of BGP routing policy changes? (Choose two)

- A. route flap dampening
- B. soft reconfiguration
- C. BGP scanner
- D. No synchronization
- E. Route refresh
- F. Synchronization

Answer: B, E

QUESTION 168:

Which one of the following as-path access-list will match routes which last came from BGP confederation member-AS 65111?

- A. router(config)#ip as-path access-list 1 permit ^(65111_
- B. router(config)#ip as-path access-list 1 permit ^65111_
- C. router(config)#ip as-path access-list 1 permit _65111_
- D. router(config)#ip as-path access-list 1 permit ^(65111_
- E. router(config)#ip as-path access-list 1 permit (65111_

Answer: A

QUESTION 169:

When you are verifying a BGP neighbor relationship using the show ip bgp summary command, which one of the following from the State/PfxRcd field of the

show output will confirm that the neighbor relationship has been established?

- A. connect
- B. active
- C. OpenConfirm
- D. In integer
- E. Established

Answer: D

QUESTION 170:

Which configuration task requires configuring the bgp cluster-id {cluster-id} command?

- A. Configuring the member ASs within a BGP confederation.
- B. Configuring the BGP confederation ID.
- C. Configuring hierarchical BGP confederations.
- D. Configuring redundant BGP confederations.
- E. Configuring hierarchical route reflectors.
- F. Configuring redundant route reflectors.

Answer: F

QUESTION 171:

DRAG DROP

Place the BGP attributes in the correct order as they are examined by the BGP route selection process.

Weight	1
MED	2
Local <i>Preference</i>	3
Origin Code	4
As Path	5

Answer:

Weight
Local <i>Preference</i>
As Path
Origin Code
MED

QUESTION 172:

What do the following commands accomplish?

```
router(config-router)#neighbor {ip addressII} route-map {name} out
router(config)#route-map {name} permit 10
router(config-route-map)#set as-path prepend {AS-Number} {AS-Number}
```

- A. They influence outbound-traffic BGP path selection.
- B. They influence inbound-traffic BGP path selection.
- C. They influence both inbound and outbound-traffic BGP path selection.
- D. They influence BGP path selection using AS-path filtering.
- E. They limit the size of the BGP table using AS-path filtering.
- F. They set the BGP confederation sub-ASs numbers.

Answer: B

QUESTION 173:

```
sanfran#show ip bgp 172.31.1.0
BGP routing table entry for 172.31.1.0/24, version 15
Paths: (2 available, best #2, table Default-IF-Routing-Table)
Advertised to non-peer-group peers:
10.1.0.2 10.200.200.3 10.200.200.4 172.31.11.1
65100
172.31.1.1 from 10.1.0.2 (10.200.200.2)
Origin IGP, metric 0, localpref 100, valid, internal
65100
172.31.1.1 from 172.31.1.1 (172.31.1.1)
Origin IGP, metric 0, localpref 100, valid, external, best
Based on the output shown in the exhibit, which two statements are correct?
(Choose two)
```

- A. The eBGP route to 172.31.1.0/24 is preferred over the iBGP route to 172.31.1.0/24.
- B. 172.31.1.1 is an iBGP peer.
- C. 10.1.0.2 is an iBGP peer.
- D. 10.200.200.3 is an eBGP peer.
- E. The sanfran router will load balance the traffic to 172.31.1.0/24 via two equal cost paths.

Answer: A, C

QUESTION 174:

Which two statements are true regarding the BGP network configuration command? (Choose two)

- A. When using the "network ip-prefix" BGP configuration command to advertise a major network into BGP, at least one of the subnets of the major-network must be present in the BGP table.
- B. When using the "network ip-prefix mask subnet-mask" BGP configuration command to advertise a classless prefix into BGP, the prefix must have an exact match in the IP routing table.
- C. When using the "network ip-prefix" BGP configuration command to advertise a major network into BGP, at least one of the subnets of the major-network must be present in the IP routing table.
- D. When using the "network ip-prefix mask subnet-mask" BGP configuration command to advertise a classless prefix into BGP, the prefix will be announced only if there is at least one network in the specified range in the IP routing table.

Answer: B, C

QUESTION 175:

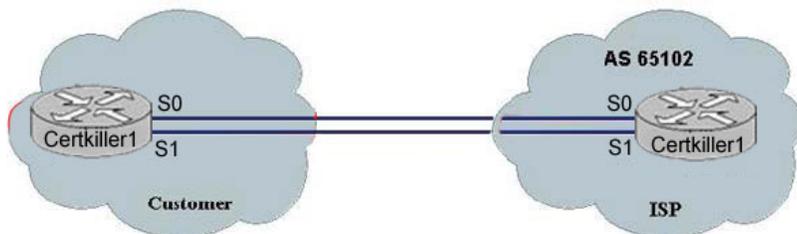
What type of access control list (ACL) is configured using a regular expression?

- A. prefix-list ACL
- B. AS-path ACL
- C. Named ACL
- D. Extended ACL
- E. BGP table-map ACL
- F. BGP community ACL

Answer: B

QUESTION 176:

Exhibit



Based on the topology diagram shown in the exhibit, when should BGP be used as the routing protocol between the customer and the ISP?

- A. If physical link failures can not be detected by the link-level procedures.
- B. If the customer wants to affect how the ISP will route the customer's traffic out to the rest of the Internet.
- C. If the ISP does not support static routing with the customer.

- D. If the customer is using provider-assigned (PA) address inside the customer's network.
- E. If the customer is using private addresses inside the customer's network.

Answer: A

QUESTION 177:

Which BGP router configuration mode command is best used to summarize prefixed learned from other BGP peers?

- A. network
- B. aggregate-address
- C. auto-summary
- D. synchronization
- E. summary-address
- F. redistribute bgp

Answer: A

QUESTION 178:

As the penalty for a flapping route decreases and falls below a certain limit, the route is unsuppressed.

What is the name of that limit?

- A. half-life limit
- B. suppress limit
- C. max-suppress-time limit
- D. reuse limit
- E. unsuppress limit
- F. penalty limit

Answer: D

QUESTION 179:

What is the proper configuration to enable route reflector redundancy?

- A. To avoid a single point of failure, IBGP full mesh can be used between all the clients in a cluster.
- B. To avoid a single point of failure, a route reflector client can have IBGP sessions to more than one route reflector within the cluster.
- C. To avoid a single point of failure, a route reflector client can have IBGP sessions to other clients in another cluster.
- D. To avoid a single point of failure, a route reflector client have IBGP sessions to other route reflectors in another cluster.

E. To avoid a single point of failure, a route reflector client can have IBGP sessions to other route reflectors and other clients in another cluster.

Answer: B

QUESTION 180:

In the establishment of a BGP session between two BGP speakers, what will be the first BGP message sent by each BGP speaker after the transport protocol connection is established?

- A. hello
- B. query
- C. open
- D. synch
- E. update
- F. notification

Answer: C

QUESTION 181:

Which three statements about a transit AS are correct? (Choose three)

- A. A transit AS has eBGP connection(s) to only one external AS.
- B. Routes between ASs are always exchanged via eBGP.
- C. A transit AS uses an IGP like OSPF or ISIS to propagate the external networks within the transit AS.
- D. Core routers within a transit AS normally use default routing to reach the external networks.
- E. iBGP sessions can be established between non directly connected routers.
- F. Edge routers within a transit AS usually announce themselves as the next hop in the iBGP updates.

Answer: B, E, F

QUESTION 182:

When the border routers in a transit autonomous system are to be configured, which configuration command should be used if the internal routers within the transit autonomous system have no IGP route to reach the router interface of the EBGP neighbor?

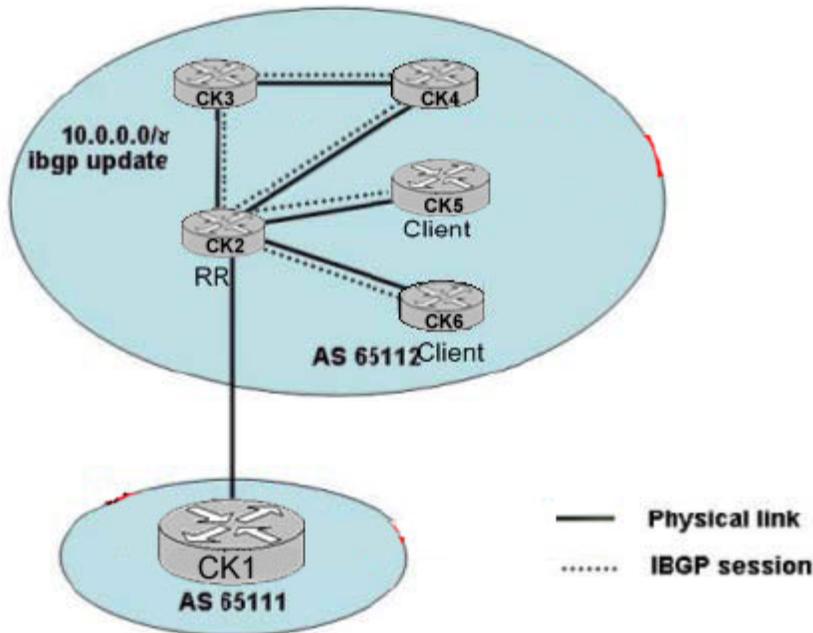
- A. neighbor {ip-address} ebgp-multihop {hop-count}
- B. neighbor {ip-address} ibgp-multihop {hop-count}
- C. neighbor {ip-address} update-source {interface}

- D. neighbor {ip-address} next-hop-self
- E. sync
- F. maximum-paths {paths}

Answer: D

QUESTION 183:

Exhibit:



Based on the network diagram shown in the exhibit, both CK5 and CK6 are clients of the CK2 RR. When the 10.0.0.0/8 iBGP update from CK3 is received by the CK2 RR, which router(s) will CK2 reflect the update to?

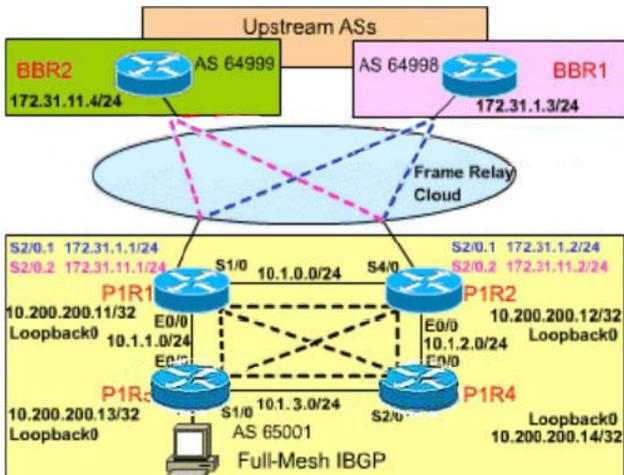
- A. CK1 only
- B. CK5 and CK6
- C. CK5 , CK6 and CK1
- D. CK4 , CK5 and CK6
- E. CK4 , CK5 , CK6 , and CK1
- F. To no other router

Answer: C

QUESTION 184:

SIMULATION

Exhibit:



You work as a network engineer at Certkiller .com. In the Transit AS 65001, router P1R3 has no BGP routes in its routing table and as a result has a problem reaching any networks external to AS 65001.

Unlike router P1R3, the other routers in AS 65002 (P1R1, P1R2 and P1R4) are not experiencing any network connectivity issue.

All BGP routes are properly inserted into the P1R4 routing table as shown below.

P1R4#sh route bpg

```
B 192.168.12.0/24 [200/0] via 10.200.200.12, 00:00:26
B 192.168.13.0/24 [200/0] via 10.200.200.12, 00:00:26
B 192.168.14.0/24 [200/0] via 10.200.200.12, 00:00:26
B 192.168.24.0/24 [200/0] via 10.200.200.12, 00:00:22
B 192.168.11.0/24 [200/0] via 10.200.200.12, 00:00:26
B 192.168.21.0/24 [200/0] via 10.200.200.12, 00:00:21
10.0.0.0/8 is variable subnetted, 9 subnets, 3 masks
B 10.0.0.0/8 [200/0] via 10.200.200.12, 00:00:26
B 192.168.23.0/24 [200/0] via 10.200.200.12, 00:00:22
B 192.168.22.0/24 [200/0] via 10.200.200.12, 00:00:21
```

Task:

1. Correct the configuration on Router P1R3 so its routing table shows the various BGP routes.
2. On router P1R3, adjust the default weight on all of the IBGP updates from router P1R2 in order to cause P1R3's BGP path selection process to choose the path through router P1R2 to exit AS 65001. Do this without initiating a hard BGP neighbor reset.

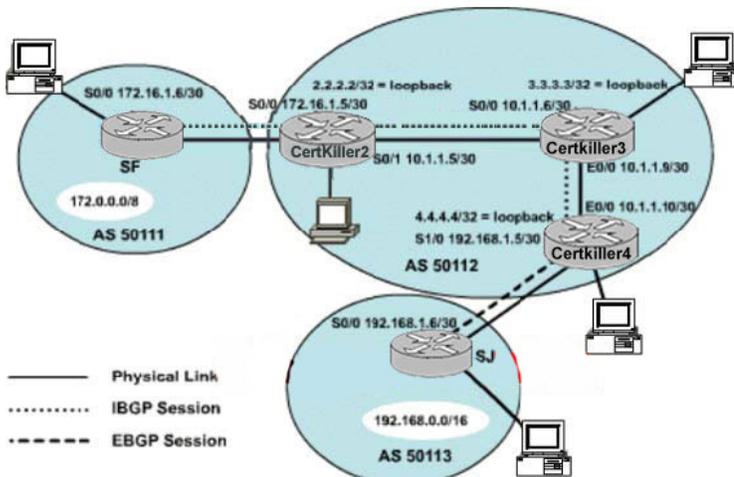
Answer:

1. no synchronization
2. neighbor 10.200.200.12 weight 100
clear ip bgp * soft in

QUESTION 185:

SIMULATION

Exhibit:



You work as a network engineer at Certkiller .com. AS 50112 is a Transit AS with two EBGP connections, one to AS 50111 and one to AS 50113.

Currently there are some connectivity problems with the network:

Updates from AS 50111 are not being propagated via AS 50112 to AS 50113

Updates from AS 50113 are not being propagated via AS 50112 to AS 50111

The current EBGP and IBGP sessions setup are shown in the diagram.

You have access to all router consoles to issue show commands, display their running configurations, BGP tables, and IP routing tables to troubleshoot the problem.

Your goal is to fix the configuration on the Certkiller 3 router to solve the problem.

Enable secret password is Certkiller .

Answer:

Explanation:

The answer could be to setup an IBGP session between Certkiller 4 and Certkiller 2 as follows:

Certkiller 4:

```
router bgp 50112
```

```
neighbor 2.2.2.2 remote-as 50112
```

```
neighbor 2.2.2.2 update source Loopback
```

```
neighbor 2.2.2.2 next-hop-self
```

and the same on the Certkiller 2:

```
router bgp 50112
```

```
neighbor 4.4.4.4 remote-as 50112
```

```
neighbor 4.4.4.4 update source Loopback
```

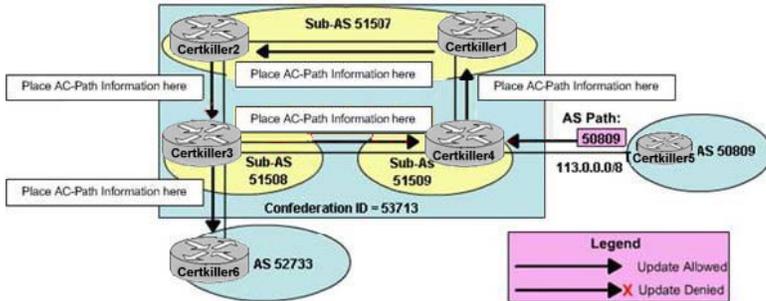
```
neighbor 4.4.4.4 next-hop-self
```

This should cause the Certkiller 2 and Certkiller 4 to forward the incoming updates to each other which would cause to forward the updates coming from the IGP to the EBGP neighbors.

or Certkiller 3 could be setup as route reflector like:

```
neighbor 2.2.2.2 route-reflector-client
```

```
neighbor 4.4.4.4 route-reflector-client
```

QUESTION 186:**SIMULATION****Exhibit, Simulation**

Select from these AC-Path informations
 *** MISSING ***

You work as a network engineer at Certkiller .com. An EBGP update about the 113.0.0.0/8 prefix is received by the router in the 51509 Sub-AS from AS 50809.

Tasks:

1. Identify the AS Path information as the 113.0.0.0/8 prefix propagated within the Confederation and to AS 52733 by selecting the correct AS-Path information from the available options.
2. Also identify if the update will be denied by any of the routers. Click on the update path to place an "X" on the respective router(s) that will deny the update.

Answer:

QUESTION 187:

Which statement is correct about transit AS and iBGP?

- A. The iBGP full mesh is a physical full mesh of all the routers within a transit AS.
- B. The iBGP full mesh is a logical mesh of TCP sessions.
- C. Transit AS running full-mesh iBGP requires synchronization to be enabled.
- D. Transit AS running full-mesh iBGP requires next-hop-self to be enabled on all the core routers.
- E. iBGP neighbors within a transit AS must be directly connected to each other.
- F. To prevent routing loops within a transit AS, eBGP updates are not propagated to iBGP peers.

Answer: B

QUESTION 188:

Exhibit:

ip tcp path-mtu-discovery age-timer infinite

What is the result of the command in the exhibit?

- A. sets the age time to zero
- B. turns off the age timer
- C. allows an infinite MTU size for the default age time of 10 minutes
- D. sets the age time to its default maximum value of 30 minutes.

Answer: B

QUESTION 189:

Which method could be used to improve the scalability of IS-IS used as the IBGP of an ISP?

- A. Use stub or totally stubby nonbackbone IS-IS areas.
- B. Use IS-IS route leaking into the nonbackbone IS-IS areas.
- C. Use IS-IS CSNP on Level-1-2 circuits.
- D. Use IS-IS pseudo-node on Level-1-2 circuits.
- E. Use proper interarea route summarization.
- F. Use proper external route summarization when redistributing BGP routes into IS-IS.

Answer: E

QUESTION 190:

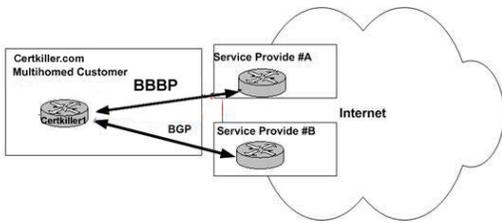
Which of the following statements about the BGP MED path attribute on Cisco routers is correct?

- A. If the MED is not attached to a BGP route, it defaults to 0 (best).
- B. By default, MED can be considered for routes coming from different autonomous systems.
- C. By default, MED can be considered for routes coming from different intraconfederation autonomous systems.
- D. The `bgp deterministic-med` command is used to ensure that an accurate MED comparison is made across all routes received from different autonomous systems.
- E. The `bgp deterministic-med` command is used to allow the MED to determine the BGP path selection over the other BGP attributes.
- F. The `bgp always-compare-med` command is used to ensure that an accurate MED comparison is made across all routes received from the same autonomous system.

Answer: A

QUESTION 191:

Exhibit



You work as a network technician at Certkiller .com. Study the exhibit carefully. What is the best way for Certkiller .com to make sure that it does not become a transit autonomous system?

- A. Use static routes to reach its service providers.
- B. Synchronize routes between its service providers.
- C. Place access list on its border routers to block transit traffic.
- D. Filter out all external routes from the updates it sends to its service providers.

Answer: D

QUESTION 192:

Which command is used to configure the external, confederation-wide AS number?

- A. Certkiller 3(config)#router bgp {as-number}
- B. Certkiller 3(config-router)#bgp confederation peers {as-number}
- C. Certkiller 3(config-router)#bgp confederation identifier{as-number}
- D. Certkiller 3(config-router)#bgp cluster-id{as-number}
- E. Certkiller 3(config-router)#neighbor {ip address} remote-as {as-number}

Answer: C

QUESTION 193:

Which BGP path attribute is used to inform the BGP peer that the routing information that was presented in the original routing updates may have been lost when the updates were summarized?

- A. AGGREGATOR
- B. AS_PATH
- C. ATOMIC_AGGREGATE
- D. CLUSTER_LIST
- E. COMMUNITY
- F. MULTI_EXIT_DISC

Answer: C

QUESTION 194:

You work as a network engineer at Certkiller .com. You are troubleshooting a BGP

neighbor relationship in which the neighbor relationship has gone down. You issue the following command:

```
show ip bgp summary
```

You see the string "PfxCt" and Idle appear in the State/PfxRcd field.

What could be the cause of this issue?

- A. The neighbor has been manually shutdown.
- B. The maximum prefix limit had been set and exceeded.
- C. The interface on the neighbor router is down.
- D. There is a problem with the service provider.
- E. The neighbor has no prefix to send.

Answer: B

QUESTION 195:

Exhibit:

(output omitted)

!

```
router bgp 65500
```

```
neighbor 172.25.25.25 remote-as 65501
```

```
neighbor 172.22.22.25 remote-as 65502
```

```
neighbor 172.23.23.22 remote-as 65503
```

!

(output omitted)

You work as a network technician at Certkiller .com. Study the exhibit carefully.

Which of the following is the correct set of commands to establish a minimum routing update interval of 60 seconds for all neighbors?

- A. Certkiller 3(config-router)#neighbor advertisement-interval 60
- B. Certkiller 3(config-router)#neighbor 172.25.25.25 advertisement-interval 60
Certkiller 3(config-router)#neighbor 172.25.22.22 advertisement-interval 60
Certkiller 3(config-router)#neighbor 172.23.23.23 advertisement-interval 60
- C. Certkiller 3(config-router)#bgp neighbor advertisement-interval 60
- D. Certkiller 3(config-router)#timers bgp 60 60 60

Answer: B

QUESTION 196:

DRAG DROP

As a technician at Certkiller .com you are required to match the correct show command to its usage.

Usage	Show commands, place here
Display the BGP Table	Place here
Display BGP attributes attached to a prefix like the community value.	Place here
Verify the BGP neighbors state.	Place here
Display the BGP routes in the routing table.	Place here
Verify AS-path access-list configuration	Place here

Show commands, Select from these

show ip route	show ip bgp
show ip bgp {network number}	show ip bgp summary
show ip bgp regexp {regular-expression}	

Answer:

Usage	Show commands, place here
Display the BGP	show ip bgp
Display BGP attributes attached to a prefix like the community value.	show ip bgp {network number}
Verify the BGP neighbors state.	show ip bgp summary
Display the BGP routes in the routing table.	show ip route
Verify AS-path access-list configuration	show ip bgp regexp {regular-expression}

Show commands, Select from these

QUESTION 197:

When configuring BGP on routers within a transit AS, under what circumstances should BGP synchronization be enabled?

- A. when a full-mesh iBGP is used
- B. when route reflectors are used
- C. when BGP confederation is used
- D. when redistributing the BGP routes into IGP
- E. when redistributing the IGP routes into BGP
- F. when the edge routers are using the BGP next-hop self option

Answer: D

QUESTION 198:

In the user of hierarchical route reflector design, what is the maximum number of route reflectors that can be configured?

- A. 1
- B. 2
- C. 3
- D. 4
- E. 5

- F. 6
- G. 10
- H. no limit

Answer: H

QUESTION 199:

When the border routers in a transit autonomous system are to be configured, which configuration command should be used if the internal routers within the transit autonomous system have no IGP route to reach the router interface of the EBGP neighbor?

- A. neighbor {ip-address} ebgp-multihop {hop-count}
- B. neighbor {ip-address} ibgp-multihop {hop-count}
- C. neighbor {ip-address} update-source {interface}
- D. neighbor {ip-address} next-hop-self
- E. sync
- F. maximum-paths {paths}

Answer: D

QUESTION 200:

Which configuration will set the local preference for the 172.16.10.0/24 prefix to 200 and allow all others prefixes from the 10.1.1.1 BGP neighbor to use the default local preference?

- A. router bgp 50001
neighbor 10.1.1.1 remote-as 50002
neighbor 10.1.1.1 route-map Certkiller in
!
ip prefix-list Certkiller seq 5 permit 172.16.0.0/8 ge 17
!
route-map Certkiller permit 10
match ip address prefix-list Certkiller
set local-preference 200
!
- B. router bgp 50001
neighbor 10.1.1.1 remote-as 50002
neighbor 10.1.1.1 route-map Certkiller in
!
ip prefix-list Certkiller seq 5 permit 172.16.0.0/24
!
route-map Certkiller permit 10
match ip address prefix-list Certkiller

```
set local-preference 200
!  
C. router bgp 50001  
neighbor 10.1.1.1 remote-as 50002  
neighbor 10.1.1.1 route-map Certkiller t in  
!  
ip prefix-list Certkiller seq 5 permit 172.16.0.0/16 ge 16  
!  
route-map Certkiller permit 10  
match ip address prefix-list Certkiller  
set local-preference 200  
!  
route-map Certkiller permit 20  
D. router bgp 50001  
neighbor 10.1.1.1 remote-as 50002  
neighbor 10.1.1.1 route-map Certkiller in  
!  
ip prefix-list Certkiller seq 5 permit 172.16.0.0/24  
!  
route-map Certkiller permit 10  
match ip address prefix-list Certkiller  
set local-preference 200  
!  
route-map Certkiller permit 20
```

Answer: D

QUESTION 201:

Exhibit

```
Certkiller 3(config-router)# neighbor 10.1.1.1 maximum-prefix 100 80 warning-only  
What effect does the command in the exhibit have?
```

- A. It prevents the 10.1.1.1 neighbor from exhausting the memory of the Certkiller 3 router.
- B. It limits the number of BGP Update messages the Certkiller 3 router will send to the 10.1.1.1 BGP peer.
- C. It logs a warning message on the Certkiller 3 router if the Certkiller 3 router BGP table contains more than 100 prefixes.
- D. It logs a warning message on the Certkiller 3 router if the Certkiller 3 router BGP table is 80% full.
- E. It resets the BGP sessions to the 10.1.1.1 BGP peer if the Certkiller 3 router BGP table contains more than 100 prefixes.
- F. It resets the BGP sessions to the 10.1.1.1 BGP peer if the Certkiller 3 router BGP table is 80% full.

Answer: A

QUESTION 202:

What can cause a single sourced iBGP route not to be selected as the best route?

- A. The BGP MED is 0.
- B. The BGP next-hop is unreachable.
- C. The BGP origin is incomplete.
- D. The BGP weight is 0.
- E. The BGP local preference is 0.
- F. BGP synchronization is disabled.

Answer: B

QUESTION 203:

Because of large BGP tables being received from your ISP Certkiller , you must filter any possible addresses that are within the private IP (RFC 1918) address space.

Which prefix lists would filter any of these private addresses?

- A. ip prefix-list Certkiller seq 5 deny 192.168.0.0/16 le 32
ip prefix-list Certkiller seq 10 deny 172.16.0.0/12 le 32
ip prefix-list Certkiller seq 15 deny 10.0.0.0/8 le 32
- B. ip prefix-list Certkiller seq 5 deny 192.168.0.0/16 ge 16
ip prefix-list Certkiller seq 10 deny 172.16.0.0/12 ge 24
ip prefix-list Certkiller seq 15 deny 10.0.0.0/8 ge 8
- C. ip prefix-list Certkiller seq 5 deny 192.168.0.0/16
ip prefix-list Certkiller seq 10 deny 172.16.0.0/12
ip prefix-list Certkiller seq 15 deny 10.0.0.0/16
- D. ip prefix-list Certkiller seq 5 deny 192.168.0.0/16
ip prefix-list Certkiller seq 10 deny 172.16.0.0/12
ip prefix-list Certkiller seq 15 deny 10.0.0.0/32

Answer: A

QUESTION 204:

What are the two ways for a router to inject its local routes into the BGP table for advertising to other BGP routes? Select two.

- A. Use the network command in BGP router configuration mode.
- B. Use the redistribute command in BGP router configuration mode.
- C. Use the neighbor command in BGP router configuration mode.
- D. Use the synchronization command in BGP router configuration mode.
- E. Use the route-map command in BGP router configuration mode.

Answer: A, B

QUESTION 205:

Exhibit

```
Certkiller# show ip route bgp
B 192.168.12.0/24 [20/0] via 172.31.1.3, 00:53:00
B 192.168.13.0/24 [20/0] via 172.31.1.3, 00:53:00
B 192.168.14.0/24 [20/0] via 172.31.1.3, 00:53:00
B 192.168.15.0/24 [20/0] via 172.31.1.3, 00:53:00
B 192.168.20.0/24[20/0] via 172.31.11.4 00:52:57
B 192.168.21.0/24[20/0] via 172.31.11.4, 00:52:57
B 192.168.22.0/24 [20/0] via 172.31.11.4, 00:52:57
B 192.168.23.0/24 [20/0] via 172.31.11.4, 00:52:57
```

```
Certkiller# show run
!
! partial show run output
```

```
router bgp 65101
 aggregate-address 192.168.12.0 255.255.252.0 summary-only
   *** MISSING ***
 neighbor 10.1.1.1 remote-as 65101
 neighbor 172.31.1.3 remote-as 65102
 neighbor 172.31.11.4 remote-as 65103
```

Based on the output of the show ip route bgp command and the BGP configuration shown in the exhibit, which BGP prefixes will be advertised by Certkiller ? Select four.

- A. 192.168.12.0/22
- B. 192.168.12.0/24
- C. 192.168.15.0/24
- D. 192.168.20.0/22
- E. 192.168.20.0/24
- F. 192.168.23.0/24

Answer: A, D, E, F

QUESTION 206:

Exhibit

```
route-map certkiller permit 10
match ip address prefix-list def
match as-path 1
set weight 111
!
route-map certkiller permit 20
match ip address prefix-list def
set weight 110
!
route-map test permit 30
set weight 109
!IP as-path access list 1 permit_64111$
!
ip prefix-list def seq 10 permit 0.0.0.0/0
```

What does the route map in the exhibit accomplish? Select three.

- A. The weight of the default route that originates in AS65111 will be set to 111.
- B. The weight of the default route that originates in AS65222 will be set to 110.
- C. The weight of all routes that originates in AS65111 will be set to 111.
- D. The weight of all routes that originates in any AS will be set to 109.
- E. The weight of all routes that originates in any AS will be set to 110.
- F. The weight of all routes other than a default route that originates in any AS will be set to 109.

Answer: A, B, F

QUESTION 207:

What are two purposes of the BGP scan-time command? (Choose two.)

- A. to tune the BGP process which walks the BGP table and confirms the reachability of next hops
- B. to allow faster detection of downed BGP peers
- C. to improve BGP convergence time
- D. to tune the BGP update interval
- E. to decrease the effects of unstable routes by increasing the route suppression time

Answer: A, C

QUESTION 208:

When creating iBGP multipaths which three criteria must be met by multiple paths to the same destination? (Choose three.)

- A. Router IDs must be the same on all routers.

- B. Each destination must have a different next-hop address.
- C. The destination AS-number must be different for each destination.
- D. Multi-exit discriminator attributes must be the same on all paths.
- E. Interior Gateway Protocol distance must be identical on each path.

Answer: B, D, E

QUESTION 209:

DRAG DROP

List the BGP route selection steps in the correct order.

Select from these	Place here
Prefer highest local preference (global within AS)	First
Prefer lowest MED	Second
Prefer routes that the router originated	Third
Prefer lowest origin code	Fourth
Prefer shorter AS paths (only length is compared)	Fifth
Prefer highest weight (local to router)	Sixth

Answer:

Place here

Prefer highest weight (local to router)
Prefer highest local preference (global within AS)
Prefer routes that the router originated
Prefer Shorter AS paths (only length is compared)
Prefer lowest origin code
Prefer lowest MED

QUESTION 210:

Which two of these statements about hierarchical route reflectors are correct?
(Choose two.)

- A. A route reflector can be a client of another route reflector.
- B. Each cluster within the hierarchy can only contain one route reflector.
- C. The hierarchy can be as deep as needed.
- D. A route reflector can have clients in different clusters.

E. Hierarchical route reflectors are set up using three levels (access, distribution, and core layers).

Answer: A, C

QUESTION 211:

DRAG DROP

Drag the method used to influence BGP path selection on the left to the traffic flow it influences on the right.

Select from these

- Setting the weight
- Setting the MED
- Setting the local preference
- Setting the admin distance
- Using recursive route lookup
- Using AS-path prepending

Used to influence the outbound traffic flow

Used to influence the inbound traffic flow

Answer:

Select from these

- Setting the admin distance
- Using recursive route lookup

Used to influence the outbound traffic flow

Setting the weight

Setting the local preference

Used to influence the inbound traffic flow

Setting the MED

Using AS-path prepending

QUESTION 212:

Which two of the following are true regarding the BGP Prefix-Based outbound route filtering feature? (Choose two.)

- A. IP multicast routes are not supported.
- B. Outbound route filtering is configured only on a per-address family basis.
- C. Outbound route filtering can be configured for either iBGP or eBGP sessions.
- D. The outbound route filter can be defined in a Prefix list, Distribute list or Access lists.
- E. Outbound route filtering is more effective when a distance vector IGP is used.

Answer: A, B

QUESTION 213:

Which configuration will enable the CK1 router in the AS51003 sub-AS (member-AS) as a route reflector with neighbors 10.1.1.1 and 10.2.2.2 as its route-reflector clients?

A. !

```
CK1 router bgp 51003
  bgp confederation identifier 55111
  bgp confederation peers 51001 51002
  neighbor 10.1.1.1 remote-as 51003
  neighbor 10.2.2.2 remote-as 51003
  neighbor 10.1.1.1 route-reflector-client
  neighbor 10.2.2.2 route-reflector-client
```

B. !

```
CK1 router bgp 51003
  bgp confederation identifier 55111
  bgp confederation peers 51001 51002
  neighbor 10.1.1.1 remote-as 51001
  neighbor 10.2.2.2 remote-as 51002
  neighbor 10.1.1.1 route-reflector-client
  neighbor 10.2.2.2 route-reflector-client
```

C. !

```
CK1 router bgp 55111
  bgp confederation identifier 51003
  neighbor 10.1.1.1 remote-as 51003
  neighbor 10.2.2.2 remote-as 51003
  neighbor 10.1.1.1 route-reflector-client
  neighbor 10.2.2.2 route-reflector-client
```

D. !

```
CK1 router bgp 55111
  bgp confederation identifier 51003
  neighbor 10.1.1.1 remote-as 55111
  neighbor 10.2.2.2 remote-as 55111
  neighbor 10.1.1.1 route-reflector-client
  neighbor 10.2.2.2 route-reflector-client
```

Answer: A

QUESTION 214:

Which BGP configuration option is designed to reduce router processing load caused by unstable routes?

- A. neighbor {ip-address} maximum-prefix {number}
- B. bgp dampening
- C. no sync
- D. bgp deterministic-med
- E. sync
- F. bgp scan-time

Answer: B

QUESTION 215:

How can you prevent multihomed customers with connections to two service providers from acting as a transit AS?

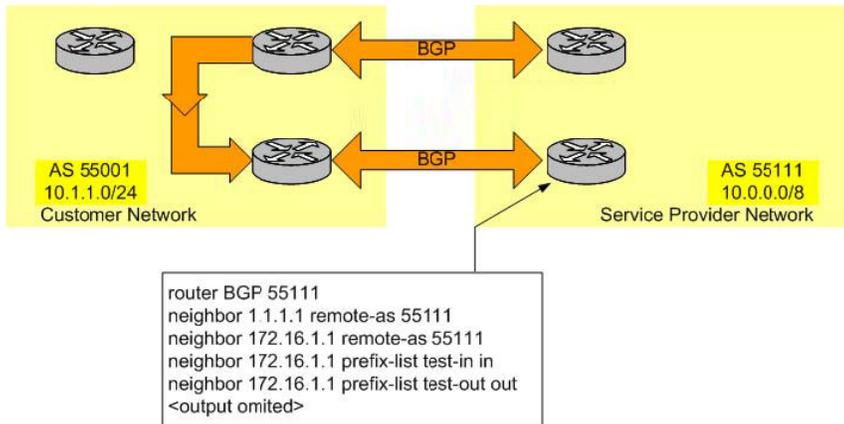
- A. Enable BGP synchronization on all the customer routers.
- B. Use MED to influence the inbound traffic from the ISPs.
- C. Use static routing to the ISPs.
- D. Use an AS-path access-list to filter the BGP updates to the ISPs.
- E. Use conditional advertisements when sending BGP updates to the ISPs.

Answer: D

QUESTION 216:

In the diagram, the customer is using BGP to connect to a single ISP over two permanent links. In this scenario, which input and output prefix-list filtering is typically enabled on the ISP routers? (Choose two.)

Exhibit:



- A. ip prefix-list test-in permit 10.1.1.0/24 le 32
- B. ip prefix-list test-in permit 10.0.0.0/8 le 32
- C. ip prefix-list test-in permit 0.0.0.0/0
- D. ip prefix-list test-out permit 10.1.1.0/24 le 32
- E. ip prefix-list test-out permit 10.0.0.0/8 le 32
- F. ip prefix-list test-out permit 0.0.0.0/0

Answer: A, F

QUESTION 217:

Based on the R1 router BGP configuration shown, which three statements are correct? (Choose three.)

Exhibit:

```
hostname R1
!
router bgp 50001
  bgp confederation identifier 50101
  bgp confederation peers 50002 50003
  neighbor 10.1.1.1 remote-as 50001
  neighbor 10.2.2.2 remote-as 50001
  neighbor 10.3.3.3 remote-as 50001
  neighbor 10.1.1.1 route-reflector-client
  neighbor 10.2.2.2 route-reflector-client
  neighbor 10.3.3.3 route-reflector-client
  neighbor 10.4.4.4 remote-as 50002
  neighbor 10.5.5.5 remote-as 50003
  neighbor 192.168.100.1 remote-as 50102
  neighbor 192.168.100.1 route-map setlp in
  neighbor 192.168.100.1 route-map setmed out
no sync
!
```

- A. R1 is in AS 50101 according to the 192.168.100.1 neighbor.
- B. R1 is in AS 50101 according to the 10.1.1.1 neighbor.
- C. The 192.168.100.1 neighbor must be directly connected to R1.
- D. R1 is a route-reflector client.
- E. The 10.4.4.4 neighbor is an EBGP neighbor.
- F. BGP updates coming in from the 192.168.100.1 neighbor must be processed by the setlp route-map.

Answer: A, C, F

QUESTION 218:

Which of these situations best describe when to use the AS number translation feature?

- A. All single-homed customers are using public AS numbers.
- B. All single-homed customers are using private AS numbers.
- C. All multihomed customers are using public AS numbers.
- D. All multihomed customers are assigned different AS numbers from different ISPs.

Answer: D

QUESTION 219:

The neighbor {ip-address} maximum-prefix {maximum number} command prevents which router condition?

- A. frequent BGP session resets
- B. routing instability
- C. asymmetric routing
- D. CPU and memory exhaustion
- E. route flaps

Answer: D

QUESTION 220:

Refer to the exhibit. What effect will the route-map PEER-FILTER have on the route 24.11.62.0/24 with a community of 10:100 injected by the peer router in AS632?

Exhibit:

```
router bgp 4224
neighbor 1.2.3.4 remote-as 632
neighbor 1.2.3.4 route-map PEER-FILTER in
!
route-map PEER-FILTER permit 10
 match ip address 1
 set weight 150
 continue 4
!
route-map PEER-FILTER permit 20
 match ip address 2
 set weight 100
!
route-map PEER-FILTER permit 30
 set weight 100
 continue
!
route-map PEER-FILTER permit 40
 match community 10:100
 set local-preference 105
!
access-list 1 permit 67.4.0.0 0.0.255.255
access-list 2 permit 24.11.0.0 0.0.255.255
```

- A. weight will be set to 100
- B. weight will be set to 150
- C. local preference will be set to 105
- D. weight will be set to 100, local preference will be set to 105
- E. weight will be set to 150, local preference will be set to 105

Answer: A

QUESTION 221:

Which two statements about a transit AS are correct? (Choose two.)

- A. A transit AS has eBGP connection(s) to only one external AS.

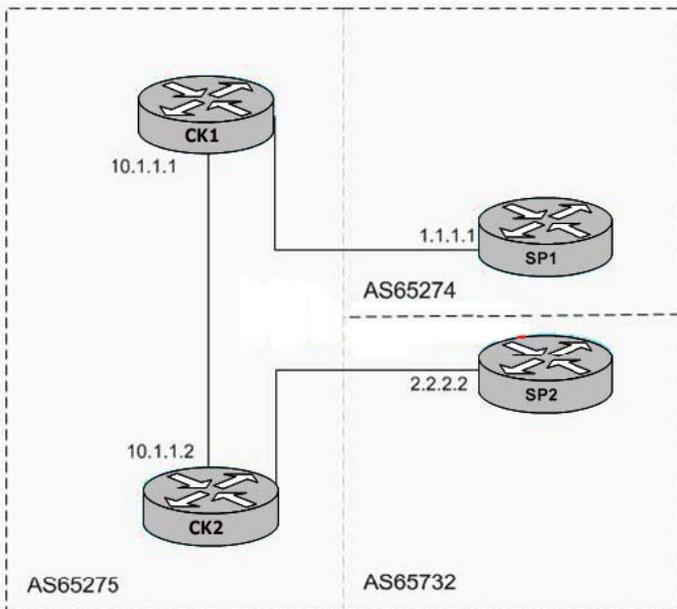
- B. Routes between ASs are always exchanged via eBGP.
- C. A transit AS uses an IGP like OSPF or ISIS to propagate the external networks within the transit AS.
- D. Core routers within a transit AS normally use default routing to reach the external networks.
- E. iBGP sessions can be established between non directly connected routers.

Answer: B, E

QUESTION 222:

Refer to the partial topology diagram shown. Service Provider 1 (SP1) assigned the customer an AS number of 65275. Service Provider 2 (SP2) assigned an AS number of 65745 to Certkiller. Certkiller decides to use AS 65275 internally. Which of the following is the correct partial router configuration to cause updates from CK1 to SP1 to report a source AS of 65275, while updates from CK2 to SP2 report the source AS of 65745 in addition to AS 65275?

Exhibit:



A. !

```
CK1 router bgp 65275
neighbor 1.1.1.1 remote-as 65274
neighbor 10.1.1.2 remote-as 65275
```

!

```
CK2 router bgp 65275
neighbor 2.2.2.2 remote-as 65732
neighbor 2.2.2.2 local-as 65745
neighbor 10.1.1.1 remote-as 65275
```

B. !

```
CK1 router bgp 65275
neighbor 1.1.1.1 remote-as 65274
```

```
neighbor 10.1.1.2 remote-as 65275
!  
CK2 router bgp 65745  
neighbor 2.2.2.2 remote-as 65732  
neighbor 2.2.2.2 local-as 65745  
neighbor 10.1.1.1 remote-as 65275  
C. !  
CK1 router bgp 65275  
neighbor 1.1.1.1 remote-as 65274  
neighbor 1.1.1.1 local-as 65745  
neighbor 10.1.1.2 remote-as 65275  
!  
CK2 router bgp 65275  
neighbor 2.2.2.2 remote-as 65732  
neighbor 2.2.2.2 local-as 65745  
neighbor 10.1.1.1 remote-as 65275  
D. !  
CK1 router bgp 65275  
neighbor 1.1.1.1 remote-as 65274  
neighbor 10.1.1.2 remote-as 65275  
!  
CK2 router bgp 65745  
neighbor 2.2.2.2 remote-as 65732  
neighbor 2.2.2.2 local-as 65275  
neighbor 10.1.1.1 remote-as 65275
```

Answer: A

QUESTION 223:

Which three actions are prerequisites to configuring the BGP link bandwidth feature? (Choose three.)

- A. enable BGP synchronization
- B. configure BGP multipath first
- C. configure eBGP multihop on eBGP peers
- D. enable Cisco Express Forwarding on the router
- E. enable BGP community propagation between link bandwidth advertising routers

Answer: B, D, E

QUESTION 224:

Which three conditions can make a transit autonomous system unable to forward packets from one neighboring autonomous system to another neighboring autonomous system? (Choose three.)

- A. sync enabled
- B. bgp next-hop not reachable
- C. no full-mesh IBGP and not using route reflectors or confederations
- D. non-directly connected IBGP neighbors
- E. non-directly connected EBGP neighbors
- F. EBGP to IBGP route redistribution configuration errors

Answer: A, B, C

QUESTION 225:

What are three characteristics of the local preference BGP attribute? (Choose three.)

- A. It is stripped in outgoing EBGP updates except with confederation peers.
- B. It is used to influence the incoming traffic.
- C. Its default local preference is 32768.
- D. It is used to select routes with equal weight.
- E. It can be used to ensure an AS-wide route selection policy.
- F. It is a Cisco-only BGP path attribute.

Answer: A, D, E

QUESTION 226:

Why can using the ip tcp path-mtu-discovery command improve BGP convergence?

- A. Single packet sizes in TCP sessions are limited.
- B. Smaller MSS sizes may reduce BGP convergence times.
- C. BGP is allowed to use a larger TCP window size.
- D. BGP is enabled to fragment its large update packets.
- E. The BGP memory requirements on routers are reduced.

Answer: B

QUESTION 227:

What state will a BGP session move to immediately after the router sends a BGP Open message to its neighbor?

- A. active
- B. OpenConfirm
- C. OpenSent
- D. established
- E. idle

Answer: C

QUESTION 228:

What is a key benefit of implementing the `bgp dmzlink-bw` command?

- A. establishes a secure BGP link between eBGP and iBGP peers
- B. enables the unequal cost path load balancing feature for external BGP links
- C. allows BGP to use link bandwidth as a metric when computing best route selection
- D. increases BGP security by dynamically allocating community strings to eBGP routes

Answer: B

QUESTION 229:

Which configuration command ensures BGP will not advertise a route until all the routers within the AS have learned about the route via an IGP?

- A. `router BGP {as-number} redistribute {IGP}`
- B. `router {IGP} redistribute BGP {as-number}`
- C. `synchronization`
- D. `bgp deterministic-med`
- E. `network {network number} mask {subnet mask}`
- F. `bgp dampening`

Answer: C

QUESTION 230:

What is the correct command to set the BGP scanner interval to two minutes?

- A. `bgp scan-time 2`
- B. `bgp scan-time 120`
- C. `bgp scan-time 2 60`
- D. The maximum scanning interval cannot exceed one minute.

Answer: D

QUESTION 231:

Refer to the configuration shown. What could cause the BGP prefixes from the 10.1.1.1 BGP peer to be absent from the routing table?

Exhibit:

```
!  
interface loopback0  
 ip address 10.10.10.10 255.255.255.255  
!  
router bgp 51001  
 synchronization  
  bgp log-neighbor-changes  
  bgp confederation identifier 51000  
  bgp confederation peers 51002 51003  
  network 10.0.0.0  
  neighbor 192.168.1.14 remote-as 51021  
  neighbor 192.168.1.18 remote-as 51022  
  neighbor 10.2.2.1 remote-as 51002  
  neighbor 10.1.1.1 next- Hcp-self  
  neighbor 10.1.1.1 update-source Loopback0  
  neighbor 10.2.2.1 remote-as 51002  
  neighbor 10.2.2.1 ebgp-multihop 255  
  neighbor 10.2.2.1 update-source Loopback0  
  neighbor 10.3.3.1 remote-as 51003  
  neighbor 10.3.3.1 ebgp-multihop 255  
  neighbor 10.3.3.1 update-source Loopback0  
 auto-summary  
!  
router rip  
 network 10.0.0.0  
!
```

- A. intraconfederation EBGp configurations error
- B. autosummarization
- C. ebgp-multihop issue
- D. BGP synchronization
- E. EBGp configurations error

Answer: D

QUESTION 232:

Which statement is correct when utilizing the BGP TTL security check feature?

- A. It is configured locally for each eBGP peering session.
- B. It compares the value of the received packet TTL against a BGP interface Access-List.
- C. It is not functional on multihop BGP peering sessions.
- D. It generates a syslog error if a "forged" packet is received.
- E. It eliminates the need for BGP keepalive packets.

Answer: A

QUESTION 233:

Which BGP router configuration command will enable an ISP edge router to advertise a default route to its BGP neighbor even if a default route is not present in its BGP table?

- A. router(config-router)#neighbor {ip-address} default-originate
- B. router(config-router)#network 0.0.0.0

- C. router(config-router)#network 0.0.0.0 mask 0.0.0.0
- D. router(config-router)#default network 0.0.0.0

Answer: A

QUESTION 234:

What is a key benefit of BGP dynamic update peer groups?

- A. Routing updates to the same destination are grouped to increase BGP efficiency.
- B. Newly configured BGP neighbors have peer group template configurations dynamically applied.
- C. Dynamic update groups use iBGP neighbor information to automatically calculate route reflector cluster configurations.
- D. Neighbors in a peer group are no longer required to share the same outbound routing policies.
- E. BGP configurations are automatically optimized by routers which dynamically create BGP peer groups.

Answer: D

QUESTION 235:

How does the extended community cost feature influence the BGP best path selection?

- A. alters the BGP AS exit path selection by adding the link cost to the local preference
- B. acts as a best path "tie breaker" when multiple IGP equal cost paths occur
- C. reflects the bandwidth of links entering the local AS from eBGP neighbors (in the MED attribute)
- D. selects the BGP route with the highest attached extended community cost value
- E. inserts the cost attribute after the MED attribute comparison, forcing best path route selection if all other preferred route selection criteria are equal

Answer: B

QUESTION 236:

DRAG DROP

Drag the correct description on the left to match the appropriate BGP feature on the right.

Select from these

- Increase stability by reducing BGP route flaps
- Reduce IBGP sessions by using special EBGP sessions
- Can use a hierarchical approach
- All routers within the AS must support it
- Override the BGP synchronization rule
- Start the BGP routing process using the cluster ID
- Limit the prefix being received from a BGP peer
- Modify the BGP split horizon rule

Confederations

-
-

Route Refelctors

-
-

Answer:

Select from these

- Increase stability by reducing BGP route flaps
- Override the BGP synchronization rule
- Start the BGP routing process using the cluster ID
- Limit the prefix being received from a BGP peer

Confederations

- All routers within the AS must support it
- Reduce IBGP sessions by using special EBGP sessions

Route Refelctors

- Modify the BGP split horizon rule
- Can use a hierarchical approach

QUESTION 237:

What is the maximum number of parallel routes that can be injected into the routing table when using the iBGP multipath load sharing feature?

- A. iBGP multipath can support up to 4 parallel paths.
- B. iBGP multipath can support up to 6 parallel paths.
- C. None, multipath is supported on eBGP only when used in conjunction with the update-source command.
- D. The number of parallel paths injected into the routing table will depend on the IGP currently in use within the iBGP domain.
- E. The maximum number of parallel paths injected into the routing table is specified by BGP MED and ranges from 2 to 256 parallel paths.

Answer: B

QUESTION 238:

What is a key benefit of implementing conditional route injection in BGP networks?

- A. allows the injection of prefix summaries without using the aggregate-address command
- B. eliminates the need for static routes to inject a default route
- C. provides for link redundancy by conditionally injecting active and standby routes
- D. allows the injection of more specific routes based on administrative policies
- E. eliminates the complexity of route aggregation when using the network command

Answer: D

QUESTION 239:

Which command is used to advertise a summary route while suppressing only a subset of the more specific routes?

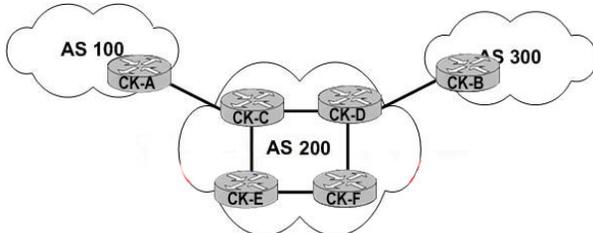
- A. network {ip prefix} {mask}
- B. network {ip prefix} {mask} unsuppress-map {route-map-name}
- C. aggregate-address {ip prefix} {mask}
- D. aggregate-address {ip prefix} {mask} summary-only
- E. aggregate-address {ip prefix} {mask} suppress-map {route-map-name}
- F. aggregate-address {ip prefix} {mask} as-set summary-only

Answer: E

QUESTION 240:

Refer to the exhibit. All routers are running BGP. Which of these statements is correct?

Exhibit:



- A. EBGP is being run between CK-A and CK-C and between CK-B and CK-D. All connections between routers inside AS 200 are IBGP.
- B. All connections between all routers are IBGP.
- C. All connections between all routers are EBGP.
- D. EBGP is being run between CK-A and CK-C, between CK-B and CK-D, and between CK-C and CK-D. IBGP is being run between all other routers.
- E. EBGP is being run between CK-A and CK-C, between CK-B and CK-D, and between CK-C and CK-D. CK-C and CK-D are route reflectors for CK-E and CK-F.

Answer: A

QUESTION 241:

When configuring BGP route dampening on Cisco routers, what is the correct default half-life-time?

- A. 15 minutes
- B. 45 minutes
- C. 60 minutes
- D. 120 minutes

Answer: A

QUESTION 242:

The core routers within a transit AS are running both IBGP and IGP. The edge routers within the transit AS are using the next-hop-self option to establish the IBGP sessions. What can be implemented to improve the routing performance to all external prefixes?

- A. enable route redistribution from BGP into IGP
- B. disable BGP synchronization on all the core routers
- C. enable CEF on all the core and edge routers
- D. enable route redistribution from IGP into BGP
- E. use route reflectors within the core

Answer: C

QUESTION 243:

Which command allows IBGP to scale within an AS?

- A. neighbor {ip-address} update-source loopback0
- B. neighbor {ip-address} route-reflector-client
- C. neighbor {ip-address} ebgp-multihop 2
- D. neighbor {ip-address} next-hop-self
- E. neighbor {ip-address} remote-as {confederation-id}
- F. no sync

Answer: B

QUESTION 244:

DRAG DROP

Drag the BGP MED function on the left to the command that enables it on the right.

Not all apply.

Exhibit:

Select from these

- Enable MED comparison among paths learned from confederation peers.
- Enable the MED to override the local preference in the route selection process.
- Allow the comparison of MED for paths from neighbors in different autonomous systems.
- Enable the a path with a higher MED value to be preferred over a path with a lower MED value.
- Enforce the comparison of MED between all paths received from the same autonomuos system.

Place here

- bgp deterministic med
- bgp always-compare-med
- bgp bestpath med-confed

Answer:

- Enable the MED to override the local preference in the route selection process.
- Enable the a path with a higher MED value to be preferred over a path with a lower MED value.

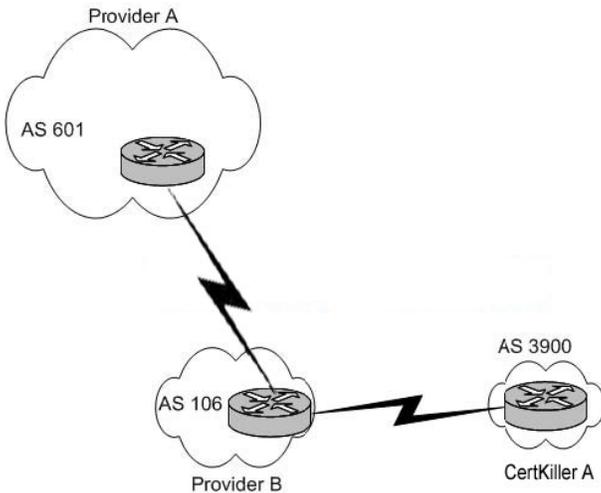
Place here

- Enforce the comparison of MED between all paths received from the same autonomuos system.
- Allow the comparison of MED for paths from neighbors in different autonomous systems.
- Enable MED comparison among paths learned from confederation peers.

QUESTION 245:

Refer to the Exhibit. How can Provider B move to AS 601 without forcing Certkiller A to immediately reconfigure?

Exhibit:



- A. Configure a second BGP AS on the Provider B router. Use the local AS feature with the dual-as keyword.
- B. Originate a default route from the Provider A network into Certkiller 's network. This will draw routed traffic into the provider network.
- C. Configure a confederation between Certkiller A and the provider networks. Use the BGP translate AS feature at the edge of the confederation.
- D. Certkiller A must update its BGP neighbor statements or the BGP sessions will no longer match.
- E. Add a wide area link between Provider A and Certkiller A configured for eBGP peering. Configure the customer BGP to prevent it from becoming a transit network.

Answer: A

QUESTION 246:

DRAG DROP

Drag the description related to route reflector on the left to the term that it describes.

Select from these	Place here
Created or updated on a route reflector only during route reflection and is removed when a route reflector advertises a route to an external peer.	Cluster ID
A group of redundant route reflectors and their clients.	Route reflector cluster
A BGP attribute configured on the route reflector and each route reflector cluster must have a unique value.	Originator ID
Occurs when the same packets are received by the router that sent them.	Routing loop
The route reflector set this attribute to the router ID of the peer from which it received the route.	Cluster list

Answer:

Select from these

Place here

The route reflector set this attribute to the router ID of the peer from which it received the route.
A group of redundant route reflectors and their clients.
Created or updated on a route reflector only during route reflection and is removed when a route reflector advertises a route to an external peer.
Occurs when the same packets are received by the router that sent them.
A BGP attribute configured on the route reflector and each route reflector cluster must have a unique value.

QUESTION 247:

External BGP peers must normally reside on a directly connected network. Sometimes it is useful to relax this restriction to enable load balancing. Which neighbor command option is used to permit this?

- A. soft-reconfiguration
- B. remote-as
- C. ebgp-multihop
- D. next-hop-self

- E. send-community
- F. no sync

Answer: C

QUESTION 248:

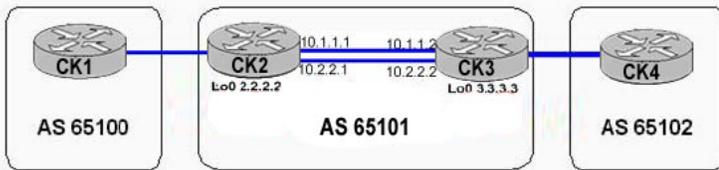
Given the following configurations, CK2 and CK3 are not able to successfully establish the IBGP session using the loopback 0 interfaces. What could be the cause of this problem?

```
!!
output omitted
!
hostname CK2
!
interface loopback 0
ip address 2.2.2.2
!
interface e0
ip address 10.1.1.1 255.255.255.0
no shut
!
interface e1
ip address 10.2.2.1 255.255.255.0
no shut
!
router bgp 65101
neighbor 172.16.1.1 remote-as 65100
neighbor 3.3.3.3 remote-as 65101
!
router eigrp 101
network 10.0.0.0
network 2.0.0.0
!!!
output omitted
!
hostname CK3
!
interface loopback 0
ip address 3.3.3.3
!
interface e0
ip address 10.1.1.2 255.255.255.0
no shut
!
interface e1
```

642-691

```
ip address 10.2.2.2 255.255.255.0
no shut
!
router bgp 65101
neighbor 192.168.1.1 remote-as 65102
neighbor 2.2.2.2 remote-as 65101
!
router eigrp 101
network 10.0.0.0
network 3.0.0.0
!
```

Exhibit:



- A. The "No Sync" BGP configuration command is missing.
- B. CK2 and CK3 are not using the loopback0 IP address as the source address for the BGP messages to each other.
- C. The "network 2.0.0.0" BGP configuration command is missing on CK2 and the "network 3.0.0.0" BGP configuration command is missing on CK3 .
- D. The "neighbor 2.2.2.2 ibgp-multihop 2" BGP configuration command is missing on CK3 and the "neighbor 3.3.3.3 ibgp-multihop 2" BGP configuration command is missing on CK2 .

Answer: B

QUESTION 249:

Refer to the outputs shown in the exhibit. What could be preventing the R1 router from receiving any prefixes from the R2 BGP neighbor?

Exhibit:

```

R1#show ip bgp summary
BGP router identifier 199.199.199.199, local AS number 20
BGP table version is 45, main routing table version 45
44 network entries using 4444 bytes of memory
81 path entries using 3888 bytes of memory
13 BGP path attribute entries using 780 bytes of memory
11 BGP AS-PATH entries using 264 bytes of memory
4 BGP route-map cache entries using 64 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 9440 total bytes of memory
BGP activity 88/44 prefixes, 191/110 paths, scan interval 5 secs

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.1.17  4     1   1628   2693     45    0    0 00:42:22    31
192.168.20.22 4    22     70     73     45    0    0 00:42:26    31
192.168.31.1  4 65002    172    274     0     0    0 00:00:13 Idle

R1#telnet 192.168.31.1
Trying 192.168.31.1 ... Open

User Access Verification

Password: cisco

R2#sh run | begin bgp
router bgp 65002
  bgp confederation identifier 1
  bgp confederation peers 65001
  network 10.0.0.0
  neighbor 192.168.31.2 remote-as 20

```

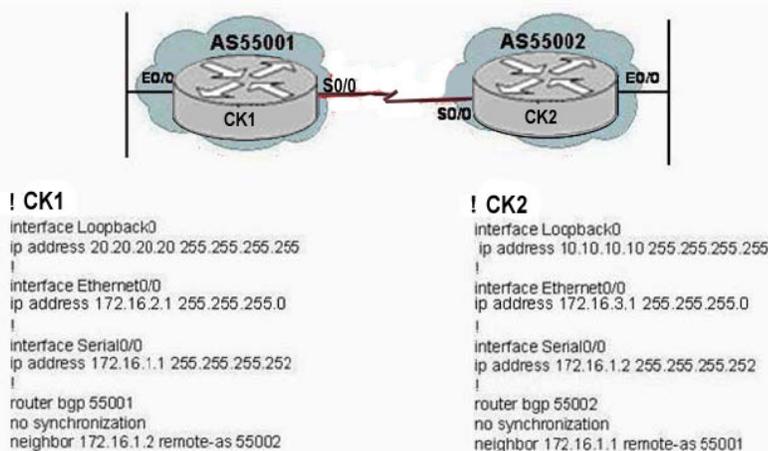
- A. There is a TCP session establishment problem between R1 and R2.
- B. The no sync command is missing on R2.
- C. The no sync command is missing on R1.
- D. R2 is using the wrong AS number in its neighbor 192.168.31.2 remote-as statement.
- E. R1 is using the wrong AS number in its neighbor 192.168.31.1 remote-as statement.
- F. Both R1 and R2 are not using a loopback address to source their BGP packets.

Answer: E

QUESTION 250:

Examine the topology and the configuration output. If the show ip bgp summary command was issued from R1, how many prefixes would it have in the State/PfxRcd field?

Exhibit:



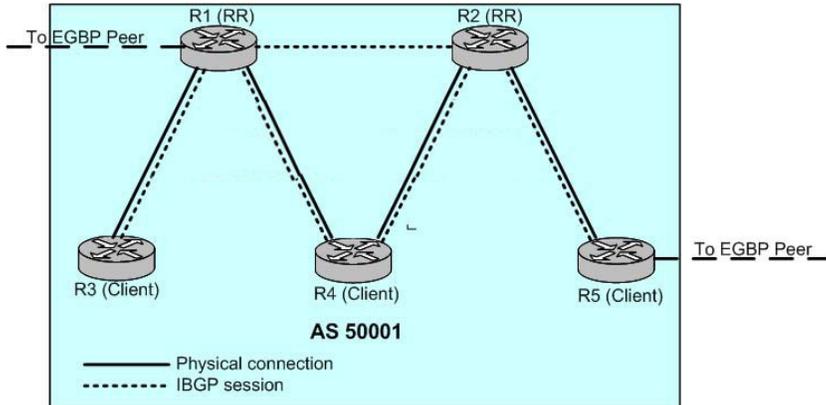
- A. 0
- B. 1
- C. 2
- D. 3

Answer: A

QUESTION 251:

Refer to the diagram. What should be changed within AS 50001 to improve the route reflector design?

Exhibit:



- A. Add a physical link between R1 and R2.
- B. Add a physical link between the clients (R3 and R4, and between R4 and R5).
- C. Remove the IBGP session between the two redundant RRs (R1 and R2).
- D. Add an IBGP session between each pair of clients (between R3 and R4, R4 and R5).
- E. Make R4 the RR and R1 and R2 its clients. R3 and R5 should be a non-RR/non-client.

Answer: A

QUESTION 252:

When verifying the BGP neighbor relationships on your router, you issue the show ip bgp summary command and there were no results. Which of the following could be the problem?

- A. The TCP session to the BGP neighbor can't be established.
- B. There are no BGP neighbors configured.
- C. All BGP updates from the BGP neighbor were filtered out.
- D. The neighbor link is down.

Answer: B

QUESTION 253:

You are a customer who is multihomed to two different ISPs—one for primary and another for backup. Each ISP assigned you a different AS number. How should you implement your AS number?

- A. Use the AS number assigned by the primary ISP.
- B. Use the AS number assigned by the backup ISP.

- C. Use both AS numbers.
- D. Use the AS number assigned by one of the ISPs, then use AS number translation when connecting to the other ISP.
- E. Use the AS number assigned by one of the ISPs, then use AS-path prepending to prepend the other AS number when connecting to the other ISP.

Answer: D

QUESTION 254:

A customer multihomed to two ISPs is assigned a different AS number from each ISP. In this case, which BGP configuration option can the customer use to resolve this issue?

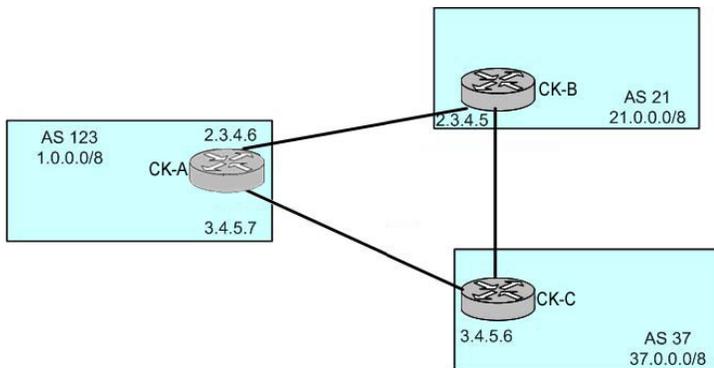
- A. set as-path prepend {as-number...}
- B. neighbor {ip-address} filter-list {as-path acl number} out
- C. neighbor {ip-address} filter-list {as-path acl number} in
- D. neighbor {ip-address} local-as {private-as}
- E. neighbor {ip-address} ebgp-multihop
- F. neighbor {ip-address} remove-private-as

Answer: D

QUESTION 255:

Refer to the exhibit. Which of these best describe why 21.0.0.0 via the next hop of 2.3.4.5 has been selected as the best path to the network?

Exhibit:



```

BGP: 2.3.4.5 rcv UPDATE about 37.0.0.0 255.0.0.0, next hop 2.3.4.5, path 21 37 metric 0
BGP: 2.3.4.5 rcv UPDATE about 1.0.0.0 255.0.0.0 - denied
BGP: 2.3.4.5 rcv UPDATE about 21.0.0.0 255.0.0.0, next hop 2.3.4.5, path 21 metric 0
  
```

```

CK-A #show ip bgp
BGP table version is 9, local router ID is 1.2.3.4
Status codes: s suppressed, d denied, h history, valid : belt , i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
  Network      Next Hop      Metric LocPrf  Weight  Path
  *> 1.0.0.0    0.0.0.0       0         32768   i
  * 21.0.0.0    3.4.5.6       0         100     37 21 i
  *>          2.3.4.5       500        100     21 i
  * 37.0.0.0    3.4.5.6       0          0        37 i
  *>          2.3.4.5       0          0         21 37 i
  
```

- A. Because it has the higher MED.
- B. Because it has the shortest AS path.
- C. Because the BGP table has not converged.
- D. Because it is the oldest learned path.

Answer: B

QUESTION 256:

Which three of these statements describe a transit AS? (Choose three.)

- A. It uses recursive routing lookup when forwarding packets toward external destinations.
- B. Core IGP stability can be improved using BGP route dampening.
- C. All core routers run IBGP and IGP.
- D. Core IGP can be scaled using route reflectors.
- E. It uses IGP routing information to reach the BGP next-hop.
- F. BGP split horizon and synchronization are disabled.

Answer: A, C, E

QUESTION 257:

Who should use private AS numbers?

- A. multihomed customers who configured their AS as a transit AS
- B. customers multihomed to two different ISPs
- C. customers multihomed to a single ISP using multiple permanent links
- D. customers connected to a single ISP using a single permanent link

Answer: C

QUESTION 258:

Which of these statements regarding the aggregate-address command is correct?

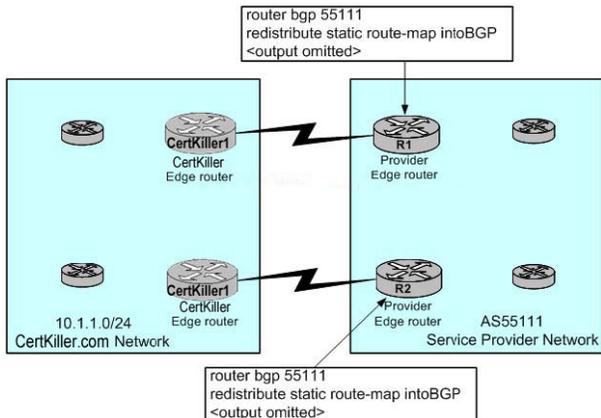
- A. There is no difference between the aggregate-address command and the network command.
- B. Using the aggregate-address command with no command options specified will only advertise the aggregate address and suppress all the more specific routes.
- C. The local BGP table must contain a more specific route if the aggregate-address command is used.
- D. Using the aggregate-address command with no command options specified will generate AS set path information.

Answer: C

QUESTION 259:

In the diagram, the customer is using static routing to connect to the ISP. Which configuration on the ISP edge routers will enable load balancing and backup of the traffic to the customer?

Exhibit:



A. ! R1

```
ip route 10.1.1.0 255.255.255.128 serial 0
```

! R2

```
ip route 10.1.1.128 255.255.255.128 serial 0
```

B. ! R1

```
ip route 10.1.1.0 255.255.255.0 serial 0
```

! R2

```
ip route 10.1.1.128 255.255.255.0 serial 0
```

C. !R1

```
ip route 10.1.1.0 255.255.255.128 serial 0
```

```
ip route 10.1.1.128 255.255.255.128 serial 0
```

!R2

```
ip route 10.1.1.128 255.255.255.128
```

```
serial 0 ip route 10.1.1.0 255.255.255.128 serial 0
```

D. ! R1

```
ip route 10.1.1.0 255.255.255.128 serial 0
```

```
ip route 10.1.1.0 255.255.255.0 serial 0
```

! R2

```
ip route 10.1.1.128 255.255.255.128 serial 0
```

```
ip route 10.1.1.0 255.255.255.0 serial 0
```

Answer: D

QUESTION 260:

DRAG DROP

Drag the correct description on the left to the EBGP, IBGP or IGP category on the right.

Select from these

- Has the lowest default admin distance
- Has the highest default distance
- Not affected by route flaps external to the AS
- Propagates the internal subnets no BSP next-hop
- Propagates routing updates to an external AS
- Propagates the external routes within the AS

EBGP

IBGP

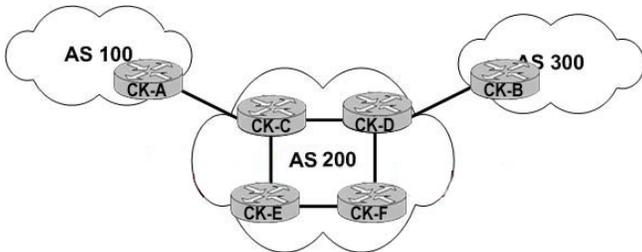
IGP

Answer:

QUESTION 261:

Refer to the exhibit. Which two statements are correct? (Choose two.)

Exhibit:



- A. The routers within AS 200 are IBGP peers; routes learned from an IBGP peer are not generally advertised to other IBGP peers.
- B. The routers within AS 200 are EBGP peers because this is not a full mesh implementation.
- C. The AS path is not manipulated when advertising a route to an IBGP peer; the local AS is added to the AS path only when advertising a route to an EBGP peer.
- D. Routers CK-C and CK-D will redistribute all the BGP prefixes into an IGP to make routing more scalable within AS 200.
- E. When router CK-D announces a prefix learned from AS 100 to router CK-B, the next-hop will be changed to the IP address of router CK-C (the interface that connects to router CK-D).

Answer: A, C

QUESTION 262:

What is a difference between using the network command vs. the aggregate-address command to perform route aggregation and summarization?

- A. The network command has more configuration options.

- B. The network command always requires the configuration of a static route pointing to the null interface.
- C. The network command requires the more specific routes to be in the BGP table.
- D. The aggregate-address command can summarize routes learned from other routing protocols besides BGP.
- E. The aggregate-address command has the option to suppress some of the more specific routes.

Answer: E

QUESTION 263:

Which two of these statements regarding BGP and admin distances are correct?
(Choose two.)

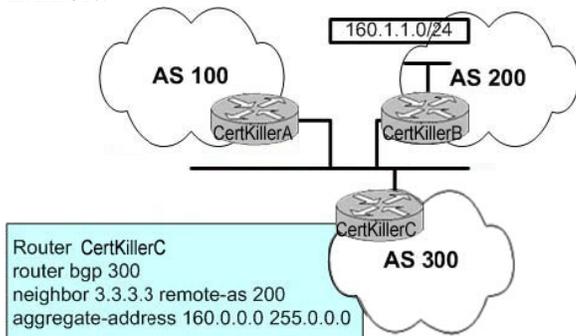
- A. BGP admin distance does not influence the path selection algorithm, but does influence if BGP-learned routes are placed in the IP routing table.
- B. External BGP routes have a default admin distance of 5.
- C. External BGP routes have a default admin distance of 20.
- D. Internal BGP routes have a default admin distance of 5.
- E. Internal BGP routes have a default admin distance of 20.
- F. BGP admin distances do not matter, since all BGP routes are automatically installed in the IP routing table.

Answer: A, C

QUESTION 264:

Refer to the exhibit. Which two of these statements are correct? (Choose two.)

Exhibit:



- A. Router Certkiller C will aggregate the address and advertise only the summary address to its neighbors.
- B. Router Certkiller C will advertise the 160.0.0.0/8 prefix and all of the more specific prefixes.
- C. Router Certkiller C cannot aggregate the 160.0.0.0/8 prefix if it does not have a more specific prefix in its BGP table.
- D. To aggregate the 160.0.0.0/8 prefix, router Certkiller C must originate that prefix from

within AS 300.

E. Router Certkiller C will only aggregate the 160.0.0.0/8 prefix for any updates it is sending to AS 200.

Answer: B, C

QUESTION 265:

Which one of these statements regarding intraconfederation EBGP sessions is correct?

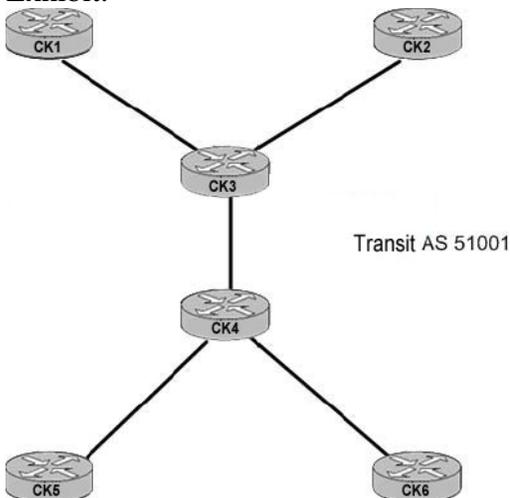
- A. Intraconfederation EBGP neighbors must be directly connected.
- B. An intraconfederation EBGP session behaves like an IBGP session when propagating routing updates.
- C. Member-AS numbers are removed when a router sends a BGP update over an intraconfederation EBGP session.
- D. Updates from an intraconfederation EBGP neighbor are subject to the BGP split horizon rule.
- E. Intraconfederation EBGP sessions must be established over loopback interfaces.

Answer: B

QUESTION 266:

Examine the topology shown above. Within Transit AS 51001, CK3 is setup as a route reflector for CK1 and CK2 , and CK4 is setup as a route reflector for CK5 and CK6 . In this case, how many IBGP sessions are required with the Transit AS 51001?

Exhibit:



- A. 4
- B. 5
- C. 10
- D. 15
- E. 30

F. 36

Answer: B

QUESTION 267:

During the autonomous system number migration process, which BGP feature allows a BGP router to act as a router within one autonomous system to some BGP neighbors but also appear to be in another autonomous system to other neighbors?

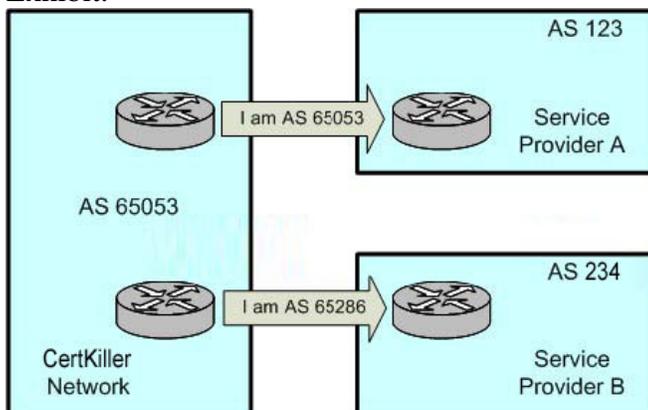
- A. remove-private-as
- B. local-as
- C. as-path prepending
- D. AS override
- E. Site-of-Origin (S00)

Answer: B

QUESTION 268:

Refer to the exhibit. Service provider A (AS 123) has an EBGP session to Certkiller where the AS number 65053 is used at Certkiller 's end. Service provider B (AS 234) also has an EBGP session to Certkiller where the AS number 65286 is used at Certkiller 's end. The customer decides to use AS 65053 internally. All router BGP configuration lines have 65053 as the AS number. The customer uses AS number 65286 only when establishing the EBGP session to AS 234. In this case, which is the correct configuration?

Exhibit:



- A. The neighbor {ip-address} local-as {private-as} command is needed on Certkiller 's router that is connected to Service Provider B.
- B. The neighbor {ip-address} local-as {private-as} command is needed on the Service Provider router in AS 234.
- C. The neighbor {ip-address} local-as {private-as} command is needed on Certkiller 's router that is connected to Service Provider A.
- D. The neighbor {ip-address} local-as {private-as} command is needed on the Service

Provider router in AS 123.

Answer: A

QUESTION 269:

Refer to the BGP configurations and the show outputs in the diagram. What are two reasons why the 197.1.0.0/16 and 192.168.1.0/30 prefixes are not in the BGP table of WGR1? (Choose two.)

Exhibit:

```
router bgp 1
no synchronization
bgp log-neighbor-changes
network 192.168.1.0
network 197.1.0.0 mask 255.255.0.0
neighbor 192.168.1.14 remote-as 22
auto-summary

WGR1#show ip route
Codes: C - connected, S - static, I - IGMP, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route

Gateway of last resort is wrtice

C 197.1.8.0/24 is directly connected, Loopback1
C 197.1.1.0/24 is directly connected, Loopback0
192.168.1.0/30 is subnetted, 3 subnets
C 192.168.1.12 is directly connected, Serial2/0.2
C 192.160.1.0 is directly connected, Serial2/0.1
C 192.168.1.16 is directly connected, Serial2/0.3
B 200.20.0.0/16 [20/0] via 192.168.1.14, 00:03:55
B 200.22.0.0/16 [20/0] via 192.168.1.14, 00:03:55

WGR1#show ip bgp
BGP table version is 26, local router ID is 197.1.8.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
*> 192.168.1.0    0.0.0.0             0         32768 i
*> 200.20.0.0/16 192.168.1.14        0          0 22 20 i
*> 200.22.0.0/16 192.168.1.14        0          0 22 i
```

- A. auto-summary is enabled.
- B. The 197.1.0.0/16 prefix is not in the routing table of WGR1.
- C. The 192.168.1.0/30 prefix is not in the routing table of WGR1.
- D. The network 192.168.1.0 command defaults to the classful mask.
- E. The aggregate-address 197.1.0.0 255.255.0.0 summary-only command should be used instead of the network 197.1.0.0 mask 255.255.0.0 command.

Answer: B, D

QUESTION 270:

If there is no exact match in the local routing table, what three conditions would result in BGP introducing the classful network 172.0.0.0 with the mask 255.0.0.0 into the BGP table? (Choose three.)

- A. no auto-summary is configured under the router bgp {as no.} command
- B. network 172.0.0.0 is configured under the router bgp {as no.} command
- C. any classful network must exist in the BGP table

- D. auto-summary is enabled under the router bgp {as no.} command
- E. the routing table contains the subnet 172.172.172.0 with a subnet mask of 255.255.255.0
- F. a classless network must exist in the BGP table

Answer: B, D, E

QUESTION 271:

Refer to the show outputs in the exhibit. R1 is using a route-map to perform AS-path prepending for only a set of prefixes. When the route-map is applied, R1 only sends the prefixes with the AS-path prepended. All other prefixes in the BGP table are no longer advertised out to the BGP neighbor. What is causing this problem?

Exhibit:

```
R1
route-map set-as permit 26
 match ip address 26
  set as-path prepend 42 26
route-map set-as deny 500
!
access-list 26 permit 0.26.0.0 255.0.255.255

R1#show ip bgp neighbors 192.168.1.17 advertised-routes
BGP table version is 45, local router ID is 199.199.199.199
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 128.26.0.0       10.0.0.0          0           32768 i

R1#show ip bgp
BGP table version is 45, local router ID is 199.199.199.199
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 10.0.0.0         192.168.20.233    0           32768 i
* 99.0.0.0          192.168.1.17      0           0 1 99 i
*                   192.168.20.22     0           0 22 1 99 i
*> 128.26.0.0       0.0.0.0           0           32768 i
*> 128.37.0.0       0.0.0.0           0           32768 i
```

- A. The access-list is missing the access-list 26 permit any any statement.
- B. The route-map is applied in the wrong direction.
- C. The route-map set-as deny 500 should be route-map set-as permit 500 instead.
- D. The BGP session on the R1 router has not been reset since the route-map was applied.
- E. The access-list 26 permit 0.26.0.0 255.0.255.255 statement is wrong.

Answer: C

QUESTION 272:

What are three benefits of an MPLS VPN? (Choose three)

- A. It provides a network-based VPN service.
- B. It provides equivalent security to Frame Relay.

- C. It eliminates the need for multiple routing protocols.
- D. It offers a more scalable solution than other VPN technologies.

Answer: A,B,D

QUESTION 273:

When running basic MPLS in conjunction with VPNs, how many labels does each packet contain?

- A. Each packet contains one label that identifies the VPN.
- B. Each packet contains at least two labels. One label identifies the path to the egress router and one that identifies the VPN.
- C. Each packet contains at least three labels. One label identifies the ingress router, one identifies the egress router and one identifies the path that will be taken.
- D. Each packet contains at least three labels. One label identifies the ingress router, one label identifies the path to the egress router, and one identifies the VPN.

Answer: B

QUESTION 274:

When running MPLS in the frame mode over an Ethernet, how does the receiving device identify that the frame contains MPLS information?

- A. MPLS frames are sent over a reserved session.
- B. MPLS is identified in the protocol port of the Layer 3 header.
- C. MPLS frames are sent in a special frame with a multicast address similar to CDP.
- D. The Ether Type of PID in the Layer 2 header identifies the frame as an MPLS frame.

Answer: D

Explanation:

1) According to the "Implementing Cisco MPLS" Student Guide (Text Part Number: 97-1153-01) Volume 1, version 1.0, page 2-31.

A label does not contain any information about the layer 3 protocol being carried in a packet. A new protocol identifier is used for MPLS-enabled Layer 3 protocol.

The following list shows the ethertype values used to identify Layer 3 protocols with most Layer 2 encapsulations.:

- Unlabeled IP unicast:PID=0x0800 identifies that the frame payload is an IP packet
- Labeled IP unicast: PID=0x8847 identifies that the frame payload is a unicast IP packet with a least one label preceding the IP header. The Bottom-of-stack bit indicates when the IP header actually starts.
- Labeled IP multicast: PID=0x8848 identifies that the frame payload is a multiicast IP packet with at least one label preceding the IP header. The Bottom-of-stack bit indicates when the IP header actually starts.

Summary: Protocol identifier in a layer 2 header specifies that the payload starts with a label (labels) and is followed by an IP header.

2) According to the RFC 3032 - MPLS label stack encoding

<<http://www.rfc-editor.org/rfc/rfc3032.txt>>

Transporting Labeled Packets over LAN Media:

Exactly one labeled packet is carried in each frame.

The label stack entries immediately precede the network layer header, and follow any data link layer headers, including, e.g., any 802.1Q headers that may exist.

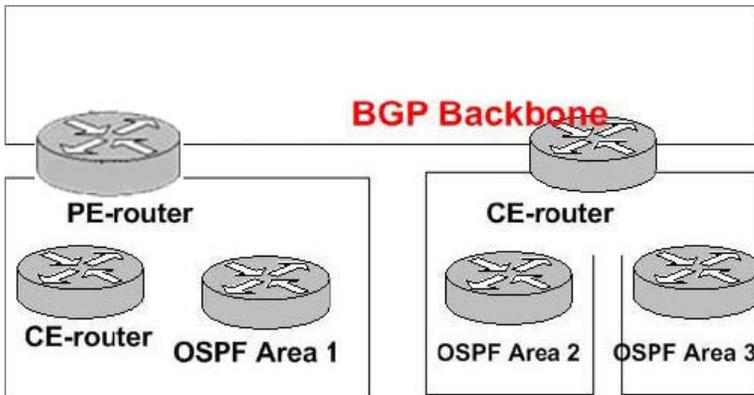
The ethertype value 8847 hex is used to indicate that a frame is carrying an MPLS unicast packet.

The ethertype value 8848 hex is used to indicate that a frame is carrying an MPLS multicast packet.

Not B: Option B cannot be correct as the question is asking frames on ethernet, this obviously means identification should be done at layer 2.

QUESTION 275:

Exhibit:



Given the information shown in the exhibit, which two statements are true? (Choose two)

- A. BGP is used to propagate routes between sites.
- B. OSPF is used to propagate routes between sites.
- C. Isolated copies of the customer's IGP run at every site.
- D. Redistribution between customer IGP and the backbone OSPF is performed at every PE-router.

Answer: A,D

QUESTION 276:

The VPN named my_vpn is operating on interface s0/0 of a PE-router. The CE-PE routing protocol is OSPF.

Why is the OSPF super-backbone needed in MPLS VPN environments?

- A. To ensure that the customer's OSPF traffic has priority over the backbone OSPF

routing updates.

- B. To ensure that the backbone internal OSPF routes are not inserted as external OSPF routes into the customer's VPN.
- C. To ensure that the customer's internal OSPF routes on one site are not inserted as external OSPF routes into other sites on the same VPN.
- D. To ensure that the customer's internal OSPF routes are not inserted as external OSPF routes into the provider backbone as internal OSPF routes.

Answer: C

QUESTION 277:

Which IOS command displays the LC-ATM capabilities of an IOS device?

- A. show mpls atm capability
- B. show mpls atm-lc capability
- C. show tag-switching lc-atm capability
- D. show tag-switching atm-tdp capability

Answer: D

QUESTION 278:

Which two commands are needed to define an MPLS VPN routing context in RIP for VPN my_vpn? (Choose two.)

- A. router (config) # router rip
- B. router (config-router) # Network 10.0.0.0
- C. router (config-router) # address-family ipv4
- D. router (config-router) # address-family ipv4 vrf my_vpn
- E. router (config-router) # redistribute bgp 65001 metric transparent

Answer: A,D

QUESTION 279:

In Cisco IOS release 12.1, how many VPN OSPF process can run simultaneously in an MPLS VPN PE-router?

- A. 1.
- B. 28
- C. 255
- D. The number of active processes is controlled by the memory available.

Answer: B

QUESTION 280:

Which command enables extended community propagation for VPNv4 MP-BGP sessions?

- A. router (config-router) # ip vpnv4 send-community both
- B. router (config-router-af) # ip vpnv4 send-community both
- C. router (config-router-af) # neighbor 172.16.1.2 send-community both
- D. router (config-router) # vpnv4 neighbor 172.16.1.2 send-community both

Answer: C

QUESTION 281:

In reference to MPLS VPNs, what is a routing protocol context?

- A. Routing protocol contexts are specified in the MPLS RFC as OSPF, BGP, and ISIS.
- B. It is how separate isolated copies of VPN routing instances are created by the IOS.
- C. It is the interface parameters and timers values used to determine which routes are exported.
- D. It is the display keyword used with the show ip route vpnv4 command to display the routing protocol parameters for a particular VRF

Answer: B

Explanations

1) According to the "Implementing Cisco MPLS" Student Guide (Text Part Number: 97-1154-01) Volume 2, version 1.0, page 8-6.

-Routing context=routing protocol run in one vrf

-Routing contexts were introduced in Cisco IOS software to support the need for separate isolated copies of VPN routing protocols. They can be implemented as separate routing processes (OSPF), similar to traditional Cisco IOS software implementation, or as separate isolated instances of the same routing protocol (EBGP, RIPv2).

2) According to the book "MPLS and VPN Architectures, CCIP Edition " by Jim Guichard , Ivan Pepelnjak.

In chapter 5: To support overlapping VPNs, the routing protocol must be limited to a single VPN routing and forwarding (VRF) table. Each PE router must be configured so that any routing information learned from an interface can be associated with a particular VRF. This is done through the standard routing protocol process and is known as the routing context. A separate routing context is used per VRF.

Some routing protocols (for example, RIP) support several instances (or routing contexts) of the same protocol, with each instance running in a different VRF. Other protocols (for example, OSPF) require a separate copy of the routing protocol process for each VRF.

QUESTION 282:

The VPN named my_vpn is operating on interface s0/0 a PE-router. The CE-PE routing protocol is RIP. The MPLS backbone IGP is OSPF. However, when you review the configuration for the PE-router, you find that BGP has been configured. What explains this configuration?

- A. VPN routes are always imported and exported using MP-BGP.
- B. This router is configured improperly. BGP is not needed in this configuration.
- C. This router is configured improperly E-BGP is the required routing protocol between PE and CE-routers.
- D. BGP should be configured on all routers supporting MPLS to ensure backward-compatibility reasons for earlier versions of IOS.

Answer: A

QUESTION 283:

Which statement about Cisco Express Forwarding's (CEF's) default operations is true?

- A. CEF is enabled by default on all Cisco routers.
- B. CEF's default operation are router dependent.
- C. CEF is disabled by default on all Cisco routers.
- D. CEF is enabled at a global level, but is disabled at an interface level.

Answer: D

QUESTION 284:

Given the following MPLS configuration to enable frame-mode MPLS over ATM:
interface atm 0/0.2 point-to-point

```
pvc auto  
ip unnumbered loopback 0  
tag-switching ip
```

Which command is using incorrect syntax?

- A. pvc auto
- B. tag-switching ip
- C. ip unnumbered loopback 0
- D. interface atm 0/0.2 point-to-point

Answer: A

QUESTION 285:

Which two are the proper syntax for enabling MPLS on an interface? (Choose two)

- A. router(config-if)# mpls ip
- B. router(config-if)# mpls enable
- C. router(config-if)# tag-switching
- D. router(config-if)# tag-switching ip
- E. router(config)# tag-switching int s0/0

Answer: A,D

QUESTION 286:

Which statement about MPLS VPN implementations and traditional peer-to-peer VPN implementations is true?

- A. MPLS and traditional peer-to-peer VPNs require the service provider to participate in the customer routing.
- B. MPLS and traditional peer-to-peer VPNs require that the customer routing is transparent to the service provider.
- C. MPLS VPNs require the use of link-state routing protocols, traditional peer-to-peer VPN implementations require distance vector protocols.
- D. MPLS VPNs are constructed using dynamic routing protocols; traditional peer-to-peer VPN implementations are constructed using static routes.

Answer: B

QUESTION 287:

How many routing table does a PE have?

- A. A PE has one routing table for each VRF
- B. A PE has one global routing table for all VRFs
- C. A PE has a global routing table and additional routing table for each VRF
- D. A PE has a CEF global routing table and additional routing table for each VRF

Answer: C

QUESTION 288:

In frame-mode MPLS, which statement is true?

- A. MPLS inserts a 32-bit label after the Layer 3 header.
- B. MPLS replaces the Layer 2 header with a 32-bit label.
- C. MPLS replaces the Layer3 header with a 32-bit label.
- D. MPLS inserts 32-bit label between the Layer 2 and Layer 3 headers.

Answer: D

QUESTION 289:

The ATM LSR forwards based upon the _____.

- A. Label
- B. VPI / VCI
- C. Route defined in the routing table.
- D. Layer 2 address of the destination.

Answer: B

QUESTION 290:

How does a PE-router operating in the default frame mode distribute labels?

- A. The PE-router distributes a label for every destination in its FIB.
- B. The PE-router only distributes labels for destinations inside the MPLS domain.
- C. The PE-router only distributes labels for destinations outside the MPLS domain.
- D. The PE-router does not distribute labels. Label distribution is only done by the core LSR.

Answer: A

QUESTION 291:

Which three statements about the traditional OSPF routing model are true?
(Choose three)

- A. Networks running OSPF can be divided into areas.
- B. OSPF was designed to support hierarchical networks.
- C. A single area is a physical site or logical division of that site.
- D. All areas must be physically connected to the backbone area (Area 0)
- E. OSPF implementations consisting of multiple areas must be interconnected by a backbone area.

Answer: A,B,E

QUESTION 292:

Why would a system administrator use the tag-switching atm maxhops global configuration command?

- A. In large ATM-MPLS networks, the LFIB can become too large. Because of this, it is necessary to limit the maximum diameter of the MPLS network.
- B. TDP relies exclusively on hop-count carried in the TDP request and replay packets to detect loops during the downstream-on-demand label allocation.

C. LDP does not support using Router ID to detect routing information in ATM networks during the downstream-on demand label allocation process.

D. Because end-to-end delay can cause problems with some voice applications, it is necessary to limit the maximum diameter of the MPLS network.

Answer: B

Explanation:

1) According to the "Implementing Cisco MPLS" Student Guide (Text Part Number: 97-1153-01) Volume 1, version 1.0, page 6-11.

tag-switching atm maxhops: Router ID can be used in LDP to detect routing information loops in ATM networks during downstream-on-demand label allocation process. TDP does not support this option and relies exclusively on hop-count carried in the TDP requests and reply packets to detect loop during the downstream-on-demand label allocation.

2) According to cisco: see the following website

http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_command_reference_chapter09186a00

tag-switching atm maxhops To limit the maximum hop counts to a value you have specified, use the tag-switching

atm maxhops global configuration command. Use the no form of this command to ignore the hop count.

tag-switching atm maxhops [number]

no tag-switching atm maxhops

Syntax Description

number(Optional) Maximum hop count.Defaults

The default is 254.

Command Modes

Global configuration

Command History

Release Modification 11.1 CTThis command was introduced.Usage Guidelines

When an ATM TSR receives a BIND REQUEST, it does not send a BIND back if the value in the request is equal to the maxhops value. Instead, the ATM TSR or TSR returns an error that specifies that the hop count has been reached.

When an ATM-TSR initiates a request for a tag binding, it includes a parameter specifying the maximum number of hops that the request should travel before reaching the edge of the ATM Tag Switching region. This is used to prevent forwarding loops in setting up tag paths across the ATM region.

Examples

The following example sets the hop count limit to 2:

tag-switching atm maxhops 2

Related Commands

Command Description show tag-switching atm-tdp bindingsDisplays the requested entries from the ATM TDP tag binding database.

QUESTION 293:

The VPN named my_vpn is operating on interface s0/0 of a PE-router. The CE-PE routing protocol is BGP. The MPLS backbone IGP is OSPF.

How should you verify routing information is being exchanged between the PE-routers that support this VPN?

- A. On the local PE-router, you use the show ip cef vrf command and verify that there is a label for the remote PE-router's next-hop.
- B. On the local PE-router, you use the show ip bgp summary command and verify that the BGP session to the remote PE-router is up.
- C. On the local PE-router, you use the show ip bgp vpn4 command and verify that routes from the remote PE-router have been received.
- D. From the CE-router, you use the ping vrf command to ping the remote CE-router. If the CE-routers cannot ping each other, then the PE-routers are not exchanging routing information.

Answer: C

QUESTION 294:

Which VRF parameters must be specified for a VRF to become operational?

- A. RD
- B. Interface
- C. Import RT.
- D. Export RT.

Answer: A

QUESTION 295:

BGP address families are used to configure which three route exchange mechanisms? (Choose three)

- A. Propagating Internet routes.
- B. VPNv4 routes that are propagated across an MPLS/VPN backbone.
- C. P-router to P-router routes that are propagated across an MPLS/VPN backbone.
- D. PE-CE routing protocol to exchange VPN routes between provider edge routers and customer edge routers.

Answer: A, B, D

Explanation:

According to the "Implementing Cisco MPLS" Student Guide (Text Part Number: 97-1154-01) Volume 2, version 1.0, page 8-31.

BGP address families: The BGP process in a MPLS VPN-enabled router performs 3 separate tasks:

- Global BGP routes(internet routing) are exchanged as in traditional BGP setup.
 - VPNv4 prefixes are exchanged through MP-BGP.
 - VPN routes are exchanged with CE routers through per-VRF EBGP sessions.
- Address families (routing contexts) are used to configure these three tasks in the same BGP process(because only one BGP process can be configured per router), and the routing contexts (called address families from the router configuration perspective) are used to configure all three independent route exchange mechanisms.
-

QUESTION 296:

You should run frame-mode MPLS over an ATM network _____. (Choose two)

- A. In situations where ATM SVCs are not available.
- B. On low end routers that cannot support the number of virtual circuits needed to run cell mode.
- C. In situations where the labeled packets must traverse ATM networks that do not support MPLS.
- D. In a migration scenario where the network is moving a classic IP over ATM to an MPLS aware ATM network.

Answer: B,C

QUESTION 297:

What is the first step to enable MPLS on an IOS-based device?

- A. Enable MPLS at a global level.
- A. B. Enable tag-switching at the global level.
- C. Enable MPLS switching on the affected interfaces.
- D. Enable Cisco Express Forwarding (CEF) at the global and affected interfaces.

Answer: D

QUESTION 298:

Given the following command:

Tag-switching advertise-tags for list1 to list2

Which statement is true?

- A. List1 specifies for which networks labels will be advertised.
- B. List1 specifies for which LSR neighbor labels will be advertised.
- C. List2 specifies the ending network address for which labels will be advertised.
- D. List1 specifies the starting network address for which labels will be advertised.

Answer: A

QUESTION 299:

How are route targets used to build virtual routing tables in the PE-routers?

- A. Route targets are prepended to each entry in the VRF to identify which VPN owns the route.
- B. Route targets are prepended to the customer's IP address to create unique address spacing for each VPN.
- C. Every customer route exported from a VRF is tagged with appropriate export route targets. VPN routes received by a PE-router are matched against import route targets configured in a CRF.
- D. Every customer route received from the CE is matched against import route targets. If there is match, the route is added to the VRF. Every route received from the neighbor PE-router is entered in the VRF. It is then matched against export route targets.

Answer: C

Explanation: Route distinguisher is prepended to IP address to create VPN address, not Route target.

QUESTION 300:

When using MPLS on an MPLS aware ATM network, which two statements are true? (Choose two)

- A. VPI and VCI are used as labels.
- B. PVCs must be manually created between end devices.
- C. The ATM devices become aware of Layer 3 by running a Layer 3 routing protocol.
- D. When running BGP, it is no longer necessary to establish logical connection between BGP neighbors.

Answer: A,B

QUESTION 301:

Which two statements about TDP and LDP are true? (Choose two)

- A. TDP and LDP populate the same LFIB.
- B. TDP operates over UDP while LDP operates over TCP.
- C. TDP and LDP both operate over TCP, but use different port numbers.
- D. While TDP and LDP are functionally equivalent, they use a different label format.

Answer: A, C

Explanation: According to the "Implementing Cisco MPLS" Student Guide (Text

Part Number: 97-1153-01) Volume 1, version 1.0, page 2-51.

TDP vs. LDP:

- MPLS and Tag switching are equal on the data plane (for me it means that they populate the same FIB and LFIB)
- The only difference is on the control plane, where tag switching uses Cisco proprietary TDP and MPLS uses standard LDP
- TDP and LDP are functionally equivalent but not compatible
- TDP use UDP and TCP port number 711
- LDP uses UDP and TCP port number 646

Not D: LDP and TDP use the same label format.

QUESTION 302:

When redistributing routes from MP-BGP to OSPF, which three are loop prevention mechanisms? (Choose three)

- A. Tag field.
- B. Down bit.
- C. Time To Live field
- D. Site of Origin field

Answer: B, C, D

Explanation:

According to the "Implementing Cisco MPLS" Student Guide (Text Part Number: 97-1154-01) Volume 2, version 1.0, pages 9-17 to 9-26

OSPF-BGP routing loops:

-OSPF down bit: An additional bit (down bit) has been introduced in the options field of the OSPF LSA header. PE routers set the down bit when redistributing routes from MP-BGP into OSPF. PE routers never redistribute routes with the down bit set into MP-BGP.

-OSPF Tag field:

The tag field in external OSPF routes is used to detect cross-domain routing loops. PE routers set the tag field to the BGP As-number when redistributing non- OSPF routes from MP-BGP into OSPF. The tag field is propagated between OSPF domains when the external OSPF routes are redistributing between OSPF domains. PE routers never redistribute OSPF routes with the tag field equal to their BGP AS-number into MP-BGP. And you know that Site of Origin is used in BGP to prevent loops to occur.

QUESTION 303:

Which command is used to display summary information on ATM tag bindings?

- A. show tag-switching vpi summary
- B. show tag-switching vpi bindings
- C. show tag-switching atm-tdp summary

D. show tag-switching atm-tdp bindings

Answer: D

QUESTION 304:

The VPN named my_vpn is operating on interface s0/0 of a PE-router. The RD is 65000:101. The CE-PE routing protocol is RIP. The MPLS backbone IGP is OSPF. Which command should be used to display all labels associated with this VPN?

- A. show mpls vrf my_vpn fib.
- B. Show ip bgp rd 65000:101 labels
- C. Show vrf my_vpn labels summary
- D. Show tag-switching forwarding vrf my_vpn

Answer: D

QUESTION 305:

Which two are VRF route-limiting options supported by IOS? (Choose two.)

- A. The maximum hop command limits the total number of hops that can be contained in a VRF.
- B. The maximum AS-in command limits the total number of Ass that can be contained in a VRF.
- C. The neighbor maximum-prefix command limits the number of routes that an individual BGP peer can send.
- D. The maximum routes command limits the total number of routes in a VRF, regardless of whether they are received from CE-routers or from other PE-router via MP-IBGP.

Answer: C, D

QUESTION 306:

You need selective VRF export so only a subset of the routes _____.

- A. From a neighbor PE-router is exported to the CE-router.
- B. From a neighbor PE-router is exported to the VRF on the local PE.
- C. Exported from a neighbor CE-router is entered into the VRF on the local PE.
- D. Exported from a VRF is one PE-router is imported into a VRF in another PE-router.

Answer: D

QUESTION 307:

Which IOS command should you use to display the contents of Label information

Base (LIB or TIB)?

- A. show tag tdp tib
- B. show tag tdp label
- C. show tag tdp detail
- D. show tag tdp bindings

Answer: D

QUESTION 308:

As related to VPNs, what is a P-network?

- A. The P-network is the part of the network under the SP's control.
- B. The P-network is the part of the network under the PTO's control.
- C. The P-network is the part of the network that is controlled using P-mib.
- D. The P-network is the part of the network that has been partitioned into an SP domain.

Answer: A

QUESTION 309:

Assuming TTL propagation is enabled, how does MPLS ensure the integrity of the TTL loop detection mechanism in IP?

- A. The TTL field in the MPLS header is set to 255. This field is decremented at each hop in the MPLS cloud. If the TTL has not expired, its value is copied back into the IP header.
- B. The TTL field in the IP header is copied to the TTL field in the MPLS header. This field is passed transparently through the MPLS cloud. The TTL value is copied back into the IP header.
- C. The TTL field in the IP header is copied to the TTL field in the MPLS header. This field is decremented by one in the MPLS cloud. If the TTL has not expired its value is copied back into the IP header.
- D. The TTL field in the IP header is copied to the TTL field in the MPLS header. This field is decremented at each hop in the MPLS cloud. If the TTL has not expired, its value is copied back into the IP header.

Answer: D

QUESTION 310:

Which statement about TDP/LDP neighbors is true?

- A. Neighbor discovery is done automatically under both TDP and LDP.
- B. Neighbor configuration is done manually when using TDP. It is automatic when using

LDP.

C. On point-to-point links, neighbor assignment is not needed. However, on broadcast links, you must assign neighbors as you do in BGP.

D. Because LDP operates under TCP, neighbors must be assigned. Because TDP uses UDP it operates in a broadcast mode and does not use neighbors.

Answer: A

QUESTION 311:

Which two enable LC-ATM control on a Catalyst interface? (Choose two)

- A. router (config-if) # ld pip
- B. router (config-if) # td pip
- C. router (config-if) # mpls ip
- D. router (config-if) # tag-switching ip

Answer: C, D

QUESTION 312:

Which three statements about MPLS VRFs are true? (Choose three)

- A. Only one interface can be assigned to a single VRF.
- B. Only one VRF can be assigned to a single interface.
- C. A VRF is a routing and forwarding instance for a VPN.
- D. VPN sites share a VRF if they are in an overlapping VPN.
- E. A single VPN site or many VPN sites can share the same VRF as long as these sites share exactly the same connectivity requirements.

Answer: B, C, E

Explanation:

B: You can have multiple interfaces associated with a VRF. But only one VRF associated with a particular interface (meaning an interface cannot be associated with multiple VRFs).

E (not D): According to the "Implementing Cisco MPLS" Student Guide (Text Part Number: 97-1154-01) Volume 2, version 1.0, page 7-64

Impact of complex VPN topologies on Virtual Routing Tables:

- A virtual routing table in a PE router can be used only for sites with identical connectivity requirements.

QUESTION 313:

When should you disable redistribution of IPv4 routing updates between MP-BGP neighbors?

- A. When the network only supports MPLS VPNs.
- B. When the network is using OSPF or ISIS instead of I-BGP as the network IGP.
- C. When you want to reduce the number of labels that the P-router needs to store.
- D. When you want to reduce the number of routes that the P-routers must store in their global BGP routing table.

Answer: B

QUESTION 314:

When is the conditional label distribution used?

- A. To block access between two MPLS VPNs.
- B. To advertise a label only when a network is available.
- C. To prevent advertising of a label when a link is overloaded.
- D. To save memory space by not advertising a label for the MPLS WAN links.

Answer: D

QUESTION 315:

What is the impact of MPLS VPN on CE-routers?

- A. The CE-routers are not MPLS VPN aware.
- B. The CE-routers must support the same IGP as the P-network.
- C. The CE-routers must use EBGP as the CE to PE routing protocol.
- D. The CE-routers must be upgraded to version 12.0 or later to support MPLS.

Answer: A

QUESTION 316:

Which two planes make up the MPLS architecture? (Choose two)

- A. Packet plane
- B. Control plane
- C. Routing plane
- D. Forwarding plane

Answer: B, D

QUESTION 317:

MPLS supports which three applications? (Choose three)

- A. VPN
- B. CDN
- C. QoS
- D. Broadcast
- E. Traffic engineering

Answer: A, C, E

QUESTION 318:

Which IOS command is used to display the entries from the ATM TDP tag-binding database?

- A. show tag-switching vpi database
- B. show tag-switching vpi bindings
- C. show tag-switching atm-tdp database
- D. show tag-switching atm-td.p bindings

Answer: D

QUESTION 319:

Which two commands specify an Rd for VRF my_vpn? (Choose two.)

- A. router (config-vrf) # rd 65000:010
- B. router (config-vrf) # rd 192.168.2.1:010
- C. router (config-int) # rd 192.168.2.1:010
- D. router (config) # ip vrf my_vpn rd 65000:010
- E. router (config) # ip vrf my_vpn rd 192.168.2.1:010

Answer: A, B

QUESTION 320:

Which command configures the propagation of a static VRF route between PE-routers?

- A. router(config-if)# ip route vrf my_vpn 10.0.0.0 255.0.0.0 10.250.0.2
- B. router(config-router-af)# ip route 10.0.0.0 255.0.0.0 10.250.0.2 s0/0.
- C. router(config-router)# ip route vrf 10.0.0.0 255.0.0.0 10.250.0.2 s0/0.
- D. router(config)# ip route vrf my_vpn 10.0.0.0 255.0.0.0 10.250.0.2 s0/0.

Answer: D

Explanation:

According to the "Implementing Cisco MPLS" Student Guide (Text Part Number: 97-1154-01) Volume 2, version 1.0, page 8-58

Configuring Per-VRF static routes:

```
router (config)#
```

```
ip route vrf name static route parameters
```

-This command configures per-VRF static routes.

-The route is entered in the VRF table.

-You must always specify the outgoing interface, even if you specify the next hop.

Sample router configuration:

```
ip route vrf Customer_ABC 10.0.0.0 255.0.0.0 10.250.0.2 serial 0/0
```

```
!
```

```
router bgp 12703
```

```
address-family ipv4 vrf Customer_ABC
```

```
redistribute static
```

According to the above example, answer D is the right one.

QUESTION 321:

Given the following access list:

```
ip access-list 100 deny tcp any any eq 646
```

```
ip access-list 100 permit ip any any
```

Which statement is true?

- A. The access list blocks Telnet session and should be applied as an input filter on all of the VT interfaces for security reasons.
- B. The access list blocks LDP traffic and should be applied as an input filter on the PE-router's interface connected to the CE to block an attempt by an outside source from exchanging labels with the PE-router.
- C. The access list blocks TCP and should be applied as an input filter on the PE-router's interfaced to the CE to block an attempt by an outside source from using programs link traceroute to identify the core routers.
- D. The access list blocks UDP and should be applied as an input filter on the PE-router's interfaced to the CE to block an attempt by an outside source from using programs link traceroute to identify the core routers.

Answer: B

QUESTION 322:

What is the P-router perception of end-to-end MPLS VPN routing?

- A. The P-router is not MPLS VPN aware.
- B. The P-router forwards packets based upon the RT contained in the label.

- C. The P-router perceives the end-to-end MPLS VPN routing as an extension of its own IGP routing.
- D. The P-router provides separate routing instances for its own IGP and the end-to-end MPLS VPN routing.

Answer: A

Explanation: P-routers are not MPLS-VPN aware.

QUESTION 323:

On ingress, a label is imposed to a packet. Which process is responsible for this function?

- A. LDP process.
- B. Control plane process
- C. Penultimate hop process.
- D. Forwarding plane process.

Answer: B

QUESTION 324:

What is the purpose of the Down bit in an OSPF LSA header?

- A. The Down bit is used to prevent redistribution loops between MP-BGP and OSPF.
- B. The Down bit is used to indicate the route that is being advertised has become unavailable.
- C. The Down bit is used to indicate an interface has been shut down or become unavailable.
- D. The Down bit is used to indicate the route that is being advertised is coming from an area that is lower in the hierarchical model.

Answer: A

QUESTION 325:

How many route targets can you configure on a single VRF?

- A. Two (one import and on export).
- B. The maximum is only limited by the router's memory.
- C. One export and as many imports as the router's memory allows.
- D. One import and as many exports as the router's memory allows.

Answer: B

QUESTION 326:

Which command series can be used to troubleshoot TDP session establishment?

- A. The debug tag-switching tdp session.
- B. The debug tag-switching tcp session.
- C. The debug tag-switching q931 session.
- D. The debug tag-switching neighbor session.

Answer: A

QUESTION 327:

What is the impact of complex VPN topologies such as Central Services VPN on the VRF tables in the PE-routers?

- A. Complex VPN topologies might require more than one VRF per VPN.
- B. Complex VPN topologies might require more than one VRF per interface.
- C. Complex VPN topologies might require the use of multiple routing protocols to separate VPN address spaces.
- D. Complex VPN topologies might require the use of the public address space to ensure there is no overlap in the address spaces.

Answer: A

Not D: D is incorrect. In the same VPN, no matter complex or not, there should not be overlapping address. Does public address solve this problem? Probably, but you need to plan in advance (buy these addresses then use them), the same can be achieved with private addresses if you DO plan in advance.

QUESTION 328:

Which statement is true about the hardware requirements of MPLS?

- A. Because you do not need to run a routing protocol on P-routers, they require less memory than routers supporting classic IP routing.
- B. Because of the additional processing and memory requirements needed to build the LFIB, MPLS is only available on high end routers.
- C. MPLS is available on low end routers, but their use is limited because of the additional processing and memory requirements needed to build the LFIB.
- D. Because P-routers do not need to carry routes outside the MPLD domain, they require less memory than routers that support the same application using classic IP routing.

Answer: C

QUESTION 329:

Which two enable LC-ATM control on a router interface? (Choose two)

- A. router (config-if) # ldp ip
- B. router (config-if) # tdp ip
- C. router (config-if) # mpls ip
- D. router (config-if) #tag-switching ip

Answer: C, D

QUESTION 330:

How does the export route map affect the VRF export process?

- A. A route map can be specified for each VRF to filter routers exported from a CE-router to a PE-router.
- B. A route map can be specified for each VRF to filter routes exported from one PE-router to another PE-router.
- C. A route map can be specified for each VRF to attach additional route targets to routes exported from a PE-router to a CE-router.
- D. A route map can be specified for each VRF to attach additional route targets to routes exported from one PE-router to another PE-router.

Answer: D

QUESTION 331:

Which command should you use to display the contents of LFIB?

- A. show tag routes
- B. show tag tdp lfib
- C. show tag bindings
- D. show tag forwarding-table

Answer: D

QUESTION 332:

Exhibit:

MISSING

What is the interaction between the routers in Area 2 and the super-backbone?

- A. The OSPF super-backbone is completely transparent to OSPF Area 2.
- B. The super-backbone appears as a BGP domain to the routers in OSPF Area 2.
- C. The super-backbone appears as another OSPF area to the routers in OSPF Area 2.
- D. The super-backbone appears as another OSPF domain to the routers in OSPF Area 2.

Answer:

QUESTION 333:

How do you test end-to-end data flow between PE-routers?

- A. Use the ping vrf command from the local PE-router to ping the remote PE router's loopback address.
- B. Use the telnet command from the local PE-router to access the remoter PE-router's loopback address.
- C. Use the traceroute command from the local PE-router to the remoter PE-router's loopback address.
- D. Because PE-PE traffic is done via label switching over an LSP, end-to-end data flow cannot be tested. You must test from CE to CE-router.

Answer: C

Explanation:

Not A: You cannot issue "ping vrf" to a loopback, when other end's loopback is not part of VPN.

QUESTION 334:

How are routing protocol context implemented in OSPF?

- A. Each routing context is implemented by redistributing into MBGP.
- B. Each routing context is implemented by assigning it to an interface.
- C. Each routing context is implemented as a separate routing process.
- D. Each routing context is implemented as a separate isolated instance of the same routing protocol.

Answer: C

Explanation:

According to the "Implementing Cisco MPLS" Student Guide (Text Part Number: 97-1154-01) Volume 2, version 1.0, page 8-6.

-Routing context=routing protocol run in one vrf

-Routing contexts were introduced in Cisco IOS software to support the need for separate isolated copies of VPN routing protocols. They can be implemented as separate routing processes (OSPF), similar to traditional Cisco IOS software implementation, or as separate isolated instances of the same routing protocol (EBGP, RIPv2).

This is the same issue as the question 201.

QUESTION 335:

What is the proper syntax for enabling Cisco Express Forwarding (CEF)?

- A. router(config)# ip cef
- B. router(config)# cef ip
- C. router(config)# cef enable
- D. router(config-if)# cef ip enable

Answer: A

QUESTION 336:

Which two statements about MPLS VPN implementations and traditional peer-to-peer VPN implementations are true? (Choose two)

- A. In MPLS VPNs, the PE-routers are aware of the customer routes.
- B. In MPLS VPNs, the PE-routers are not aware of the customer routes.
- C. In traditional peer-to-peer VPN implementations, the PE-routers are aware of the customer routes.
- D. In traditional peer-to-peer VPN implementations, the PE-routers are not aware of the customer routes.

Answer: A, C

A: Since in MPLS VPN, PE and CE still need to exchange routing information then propagate through MP-iBGP. in peer-to-peer VPNS, PE and P routers ARE in customer's network, of course they exchange routing information. Peer-to-peer routers are visible to customers because they are INSIDE the customer network, so B cannot be correct.

QUESTION 337:

Which two are IP-based overlay VPN technologies? (Choose two.)

- A. MPLS.
- B. Frame Relay
- C. IP Security (IPSec)
- D. Generic Route Encapsulation (GRE)

Answer: C, D

Explanation:

QUESTION 338:

Which IOS command should you use to list routers in an LSR's adjacency table?

- A. show tag tdp nei
- B. show tag tdp adj

- C. show tag tdp sum
- D. show tag tdp detail

Answer: A

Explanation:

The TDP neighbors and the status of individual TDP sessions also can be monitored with the show tag tdp neighbor command. The command displays the TDP identifiers of the local and remote routers, the IP addresses and the TCP port numbers between the TDP connection is established, the connection uptime and the interfaces through which the TDP neighbor was discovered, as well as all the interface IP addresses used by the TDP neighbor.

Reference: MPLS and VPN Architectures (Cisco Press) page 33

QUESTION 339:

What are two mandatory configuration steps to enable MPLS in an all Cisco networks?
(Choose two)

- A. router (config)# mpls
- B. router (config)# ip cef
- C. router (config-if)# mpls ip
- D. router (config-if)# mpls label protocol ldp

Answer: C

Explanation:

ip cef

To enable Cisco Express Forwarding (CEF) on the route processor card, use the ip cef command in global configuration mode.

ipcef [distributed]

mpls ip (interface configuration)

To enable MPLS forwarding of IPv4 packets along normally routed paths for a particular interface, use the mpls ip interface configuration command.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_command_reference_chapter09186a008008

QUESTION 340:

How does the import route map affect the VRF import process?

- A. The import route map overrides the route target import filter.
- B. A route must pass either the route target import filter or the import route map to be imported.
- C. The import route map overrides the route target import filter and controls the import of

routes.

D. A route has to pass the route target import filter first and then the import route map to be imported.

Answer: B

Explanation:

Use an import route map when an application requires finer control over the routes imported into a VRF than provided by the import and export extended communities configured for the importing and exporting VRF.

The import map command associates a route map with the specified VRF. You can use a route map to filter routes that are eligible for import into a VRF, based on the route target extended community attributes of the route. The route map might deny access to selected routes from a community that is on the import list.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_command_reference_chapter09186a008017

QUESTION 341:

What is the proper command to associate interface s0/0 with a VRF named my_vpn?

- A. Router (config-vrf)# int s0/0
- B. Router (config-if)# vrf my_vpn
- C. Router (config)# if vrf my_vpn int s0/0
- D. Router (config-if)# ip vrf forwarding my_vpn

Answer: D

Explanation:

After you define all relevant VRFs on the PE router, you must tell the PE router which interfaces belong to which VRF and, therefore, should populate the VRF with routes from connected sites. More than one interface can belong to the same VRF. You can do this by using the ip vrf forwarding interface-mode command, which associates the interface with the name VRF. Both main and sub-interfaces can be defined within a VRF.

Reference: MPLS and VPN Architectures (Cisco Press) page 204

QUESTION 342:

You need VRF route to limit the effect of _____. (Choose two)

- A. Route loops on the MPLS VPN backbone.
- B. Malicious behavior on the MPLS VPN backbone.
- C. Excessive hop count on the customer's network.
- D. Configuration errors on the MPLS VPN backbone.

Answer: A C

Explanation:

QUESTION 343:

Which command configures the redistribution of static VRF routes between PE-routers?

- A. router (config)# redistribute static
- B. router (config-if)# redistribute static
- C. router (config-router)# redistribute static
- D. router (config-router-af)# redistribute static

Answer: C

Explanation:

Router(config-router)#redistribute static
Redistributes VRF static routes into the VRF BGP table.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products_feature_guide09186a00800e977b.html

QUESTION 344:

The VPN named my_vpn is operating on interface s0/0 of a PE-router. The CE-PE routing protocol is RIP. The MPLS backbone IGP is OSPF. However, when you review the configuration for the PE-router, you find that BGP has been configured. What explains this configuration?

- A. VPN routes are always imported and exported using MP-BGP.
- B. This router is configured improperly. BGP is not needed in this configuration.
- C. This router is configured improperly. E-BGP is the required routing protocol between PE and CE-routers.
- D. BGP should be configured on all routers supporting MPLS to ensure backward-compatibility reasons for earlier version of IOS.

Answer: C

Explanation:

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

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products_feature_guide09186a00800e977b.html#111

QUESTION 345:

Which command specifies an RT for VRF my_vpn?

- A. Router (config-vrf)# route-target both 12703:15
- B. Router (config-router-af)# route-target import 12703:15
- C. Router (config)# ip vrf my_vpn route-target import 12703:15
- D. Router (config-if)# ip vrf my_vpn route-target import 12703:15

Answer: A

Explanation:

route-target

To create a route-target extended community for a VRF, use the route-target VRF submode command. To disable the configuration of a route-target community option, use the noform of this command.

route-target{import | export | both} route-target-ext-community
noroute-target {import | export | both}route-target-ext-community

Syntax Description

import Imports routing information from the target VPN extended community.

export Exports routing information to the target VPN extended community.

both Imports both import and export routing information to the target VPN extended community.

route-target-ext-community Adds the route-target extended community attributes to the VRF's list of import, export, or both (import and export) route-target extended communities.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products_feature_guide09186a00800e977b.html#111

QUESTION 346:

How could you check for potential MTU size issues on the path taken by a PE-to-PE LSP?

- A. Because MPLS packets are label switched, MTU problems can only be detected by the user applications.
- B. Use the ping vrf command with packet size set to the largest MTU along the path and DF bit set from the local PE-router to ping the remote PE-router.
- C. Use the ping vrf command with packet size set to the smallest MTU along the path

and DF bit set from the local PE-router to ping the remote PE-router.
D. Because MPLS packets are label switched, packets are automatically fragmented and reassembled by the PE-routers.
Therefore, there are no potential MTU issues.

Answer: A D

QUESTION 347:

Which three statements about MPLS VRFs are true? (Choose three)

- A. Only one interface can be assigned to a single VRF.
- B. Only one VRF can be assigned to a single interface.
- C. A VRF is a routing and forwarding instance for a VPN.
- D. VPN sites share a VRF if they are in an overlapping VPN.
- E. A single VPN site or many VPN sites can share the same VRF as long as these sites share exactly the same connectivity requirements.

Answer: C D E

Explanation:

1. A forwarding table is derived from the routing table and is based on CEF technology.
2. A set of interfaces that use the derived forwarding table.
3. Rules that control the import and export of routes from and into the VPN routing table. These rules were introduced to support overlapping VPNs and are explained later in this chapter.
4. A set of routing protocols/peers, which inject information into the VPN routing table. This includes static routing.
5. Router variables associated with the routing protocol that is used to populate the VPN routing table.

The combination of the VPM IP routing table and associated VPN IP forwarding table is called VPN routing and forwarding instance (VRF).

Reference: MPLS and VPN Architectures (Ciscopress) page 165 + 166

QUESTION 348:

The VPN called my_vpn is operating on interface s0/0 of a PE-router. The CE-PE routing protocol is RIP. The MPLS backbone IGP is OSPF.

How could you verify that the routes received from the remote CE-router are being redistributed into RIP?

- A. Verify the routing table in the local CE-router.
- B. Use the show ip cef vrf command on the local PE-router and verify the adjacency table in the CEF cache.
- C. Use the debug ip routing command in the local PE-router and verify which routes are advertised to the local CE-router.

D. Use the show ip routing advertised command on the local PE-router and verify that the routes are being advertised to the CE-router.

Answer: B

Explanation:

To display the Cisco Express Forwarding (CEF) forwarding table associated with a Virtual Private Network (VPN) routing/forwarding instance (VRF), use the show ip cef vrfcommand in privileged EXEC mode.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_command_reference_chapter09186a008017

QUESTION 349:

The FIB table is different from other fast-switching caches. It does NOT contain the information about the outgoing interface and the corresponding Layer 2 header. Where is this information stored?

- A. ARP table
- B. Neighbor table
- C. Adjacency table
- D. The above statement is false.
- E. The Layer 2 header is part of the FIB.

Answer: C

Explanation:

Adjacency Tables

Nodes in the network are said to be adjacent if they can reach each other with a single hop across a link layer. In addition to the FIB, CEF uses adjacency tables to prepend Layer 2 addressing information. The adjacency table maintains Layer 2 next-hop addresses for all FIB entries.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_configuration_guide_chapter09186a00800c

QUESTION 350:

What does the NEGOTIATED field in the output of the show tag-switching atm-tdp capability command indicate?

- A. It indicates that the VC merge capability has been successfully negotiated.
- B. It indicates a set of options that both LDP peer devices have agreed to share on this interface.
- C. It indicates that both TDP peer devices have completed the negotiation phase and have

agreed on the TDP/LDP distribution.

D. It indicates that the negotiated allocation scheme is UNDIR, if and only if, both peer devices have UNIDIR capability (otherwise it is BIDIR).

Answer: B

Explanation:

Negotiated - Set of options that both LDP peer devices have agreed to share on this interface. For example, the VPI or VCI allocation on either peer device remains within the negotiated ranges.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_command_reference_chapter09186a008008

\

QUESTION 351:

Why would you explicitly define the VPI/VCI when deploying cell-mode MPLS?

- A. To reduce the number of labels required.
- B. Because cell-mode MPLS only runs over defined VPI/VCIs.
- C. To ease a migration by connecting an MPLS-enabled ATM network to the old ATM network.
- D. Because MPLS is needed and parts of the private network cross a non-MPLS aware network.

Answer: A

Explanation: A is the only possible answer. A technique (VC merge) actually exists to help reduce the number of labels (and thus VC) created. by mapping multiple incoming label values to a specific outgoing label, the number of VCs can be reduced. of course i am not sure the meaning of the word "explicitly" since VC merge seem to do all the mapping automatically.

Incorrect Answers:

B: B is false, as the reference states ...with some hardware platforms supporting only approximately 1000 VCs. The very small number of VCs supported over an ATM interface makes these circuits a scarce resource that must be tightly controlled. if 1000 is only a small number, and if optionB is true, then according to the question, all 1000 VCs must be explicitly defined. which is not possible, so we may "assume" that VCs can be dynamically assigned.

C: C doesn't seem to be correct. Explicitly defining VCs require careful designs.

D: Option D is wrong, since a LSP requires all devices to be able to distribute labels, if non-MPLS aware switches are

used, they cannot fulfill the downstream-on-demand feature. in this situation, only Frame (but not Cell) mode can be used.

QUESTION 352:

What does the REMOTE field in the output of the show tag-switching atm-tdp summary command indicate?

- A. The number of ATM labels assigned by this LSR on this interface.
- B. The number of ATM labels assigned to outgoing labels on this interface.
- C. The number of ATM labels contained in incoming labels on this interface.
- D. The number of ATM labels assigned by the neighbor LSR on this interface.

Answer: D

Explanation:

remote- The number of ATM labels assigned by the neighbor LSR on this interface.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_command_reference_chapter09186a008008

QUESTION 353:

What are two methods to achieve differentiated QoS in MPLS? (Choose two)

- A. By creating a separate Label Switch Path for each class.
- B. By using the experimental bits to identify the different classes.
- C. By using the DiffServ Codepoints to identify the different classes.
- D. By using the TOS bits in the Layer 3 header to identify the different classes.

Answer: B C

Explanation:

Two general approaches are used to mark MPLS traffic for QoS handling within an MPLS network in the first method, the DiffServ "coloring" information is carried in the experimental (EXP) field of the MPLS shim header. This field allows for eight different QoS markings. Label Switched Paths (LSPs) using this approach are called E-LSPs, signifying that QoS information is inferred from the EXP field. Alternatively, IETF specifications allow for a second method of carrying the DiffServ information. Here, the label associated with each MPLS packet carries the portion of the DiffServ marking that specifies how a packet should be queued. The dropping precedence portion of the DiffServ marking is carried either in the EXP field, if an MPLS shim header is being used, or on fields available for this purpose on underlying technologies (for example, CLP bit for ATM and DE bit for Frame Relay). Switching paths within the MPLS network using this approach are called L-LSPs, signifying that QoS information is inferred, in part, from the MPLS

label.

LSPs supporting DiffServ may be established with bandwidth reservation. That is, bandwidth requirements for a label switched path could be signaled at LSP establishment time. Bandwidth reservation could be used to perform admission control on the DiffServ resources that have been provisioned. Though admission control can be performed on an LSP basis, the QoS design within the MPLS network is DiffServ-based, taking advantage of the scalability benefits implicit in that QoS architecture.

Reference:

http://www.cisco.com/en/US/tech/CK436/CK428/technologies_white_paper09186a00800a4455.shtml

QUESTION 354:

Which statement about configuring frame MPLS over an ATM PVC is true?

- A. The VPI/VCI must be manually assigned using a value from the label pool.
- B. The VCI/VPI are assigned dynamically by LDP using a value from the label pool.
- C. The VPI/VCI must be manually assigned using a value supported by the local ATM switch.
- D. The interface automatically learns the VPI/VCI from the local ATM switch and associates it with the defined MPLS.

Answer: B

Explanation:

Label switching Controlled ATM interface (LC-ATM interface) is an interface on a router or an ATM switch in which the VPI/VCI value is assigned through MPLS control products (TDP or LDP).

Reference: MPLS and VPN Architectures (Cisco Press) page 50

QUESTION 355:

When using IP over ATM, which three statements are true? (Choose three)

- A. Layer 3 PVCs must be established to provide routing paths.
- B. A Layer 2 hub and spoke topology is usually used because it is easier to manage.
- C. Layer 2 devices have no knowledge of Layer 3 routing information. Virtual circuits must be established.
- D. Layer 2 topology may be different than the Layer 3 topology, resulting in suboptimal paths and link utilization.

Answer: A C D

Explanation:

When trying to fit MPLS architecture into the limitations of ATM technology, a number of obstacles must be overcome:

1. There is no mechanism for direct exchange of IP packets between two adjacent MPLS nodes over an ATM interface. All data exchange over an ATM interface must take place over an ATN virtual circuit (VC).

2. ATM switches cannot perform MPLS label lookup or layer 3 lookup. The only capability of an ATM switch is to map the incoming VC in a cell to an outgoing VC and an outgoing interface.

Reference: MPLS and VPN Architectures (Ciscopress) page 49

QUESTION 356:

Which MPLS plane is responsible for packet forwarding?

- A. packet plane
- B. control plane
- C. routing plane
- D. forwarding plane

Answer: D

Explanation:

MPLS compliments IP technology. It is designed to leverage the intelligence associated with IP Routing, and the Switching paradigm associated with Asynchronous Transfer Mode (ATM). MPLS consists of a Control Plane and a Forwarding Plane. The Control Plane builds what is called a "Forwarding Table," while the Forwarding Plane forwards packets to the appropriate interface (based on the Forwarding Table).

Reference:

http://www.cisco.com/en/US/tech/CK436/CK428/technologies_white_paper09186a00800b010f.shtml

QUESTION 357:

To whom does a PE-router advertise labels, when operating in the default frame mode?

- A. The PE-router advertises a label to all of its neighbors.
- B. The PE-router only advertises labels to its neighbors in the MPLS domain.
- C. The PE-router only advertises labels to its neighbors outside the MPLS domain.
- D. The PE-router does not advertise labels.

Label advertisement is only done by the core LSR.

Answer: B

Explanation:

The Ingress Edge-LSR receives an IP packet, classifies the packet into a forward equivalence class (FEC), and labels the packet with the outgoing label stack corresponding to the FEC. For unicast destination-based IP routing, the FEC corresponds to a destination subnet and the packet classification is a traditional layer 3 lookup in the

forwarding table.

Reference: MPLS and VPN Architectures (Cisco Press) page 25

QUESTION 358:

What is penultimate hop popping?

- A. A process run on the ingress router that assigns the label to the packet.
- B. A process run on the P-router that removes a label before forwarding the packet to the egress router.
- C. A process run on the ingress router that sets the TTL value to 255 to hide the core routers from a traceroute.
- D. A process run on the egress router that pops (sets) the TTL value in the Layer 3 header to the value that is contained in the top label TTL value.

Answer: B

Explanation:

Traffic Engineering Components

A router capable of supporting MPLS is known as Label Switching Router (LSR). The LSR, found just before the last LSR in the MPLS clouds, is known as the penultimate hop. The end-to-end MPLS path is known as Label Switched Path (LSP). LSP is originated at the head-end router and terminates at the tail-end router.

Reference:

http://www.cisco.com/en/US/tech/CK436/CK428/technologies_white_paper09186a00800a4472.shtml

QUESTION 359:

Which two statements about cell-mode MPLS loop prevention mechanisms are true?
(Choose two)

- A. MPLS relies on the loop prevention mechanism MP-BGPB.
- Loop detection algorithms in LDP ensure that loops are not present.
- C. MPLS relies on the loop prevention mechanisms embedded in the IGP.
- D. MPLS labels contain a Time-To-Live (TTL) field that prevents packets from looping indefinitely.

Answer: B, D

Not A: A is not correct too since running cell mode MPLS doesn't require neither MPLS VPN nor MP-iBGP.

Not C: C is not correct as ATM switches cannot run IGP.

QUESTION 360:

How does a core LSR operating in the default frame mode advertise labels?

- A. The core LSR advertises a label for every destination in its FIB.
- B. The core LSR only advertises labels for destinations inside the MPLS domain.
- C. The core LSR only advertises labels for destinations outside the MPLS domain.
- D. The core LSR does not advertise labels. Label advertisements is only done by the PE-router.

Answer: A

Explanation:

Core LSRs receive this labeled packet and use label forwarding tables to exchange the inbound label in the incoming packet with the outbound label corresponding to the same FEC.

Reference: MPLS and VPN Architectures (Cisco Press) page 25

QUESTION 361:

Which statement about an overlapping VPN is true?

- A. A site participates in more than one VPN.
- B. A site has a link to both the intranet and the Internet.
- C. The VPN contains both private and public address spaces.
- D. The VPN contains two or more overlapping address spaces.

Answer: A

Explanation:

To support connectivity requirements, the MPLS/VPN architecture supports the concepts of sites, where a VPN is made up of one or multiple sites. A VPN is essentially a collection of sites sharing common routing information, which means that a site may belong to more than one VPN if it holds routes from separate VPNs.

Reference: MPLS and VPN Architectures (Cisco Press) page 169

QUESTION 362:

How many routing tables reside on a P-router?

- A. The P-router has a single global routing table.
- B. A P-router has one routing table for each VRF.
- C. A P-router has a global routing table for each VRF.
- D. A P-router has a global routing table and one additional routing table for each VRF.

Answer: A

Explanation:

A, obviously. since it is a P router but not a PE router. It is used to transfer the route to other PE routers as part of the path. It does not "inject" the routes to any of their routing

tables. Since it is just part of a path, it needs ONLY a global routing table to inter-connect PE devices.

Not D: Option D is true for PE-router.

QUESTION 363:

How are customer routes exchanged across the P-network?

- A. LDP is used to exchange customer routes across the P-network.
- B. IBGP is used to exchange customer routes across the P-network.
- C. OSPF is used to exchange customer routes across the P-network.
- D. MP-BGP is used to exchange customer routes across the P-network.

Answer: D

Explanation: It must be D, OSPF can only be used as in IGP within P network and does not have the capabilities to carry any customer route.

QUESTION 364:

As related to VPNs, what is a CE-router?

- A. A CE-router is a router in the C-network with a link to the P-network.
- B. A CE-router is a router in the P-network that provides the link to the customer's site.
- C. A CE-router is a router in the P-network that serves as the customer's entry point to the SP network.
- D. A CE-router is a router in the customer's site that provides the address translation necessary to ensure that there is no overlap in the addressing scheme.

Answer: C

Explanation:

The customer connects to the service provider network through a Customer Premises Equipment device. The CPE is usually a Packet Assembly and Disassembly (PAD) device that provides plain terminal connectivity, a bridge, or router. The CPE device is also sometimes called a Customer Edge (CE) device.

Reference: MPLS and VPN Architectures (Cisco Press) page 131

QUESTION 365:

Which two attributes in BGP are used to help with implementation of the OSPF super-backbone? (Choose two)

- A. MED
- B. weight
- C. AS path

- D. site of origin
- E. extended community

Answer: D E

Explanation:

* The VPN information is transported using BGP extended communities from PEs to other PEs and is re-injected into the OSPF areas as Summary Network (type 3) Link-State Advertisements (LSAs).

The MPLS VPN Super Backbone also enables customers to use multiple area 0 backbones on their sites. Each site can have a separate area 0 as long as it is connected to the MPLS VPN Super Backbone. The result is the same as a partitioned area 0 backbone.

Reference:

http://www.cisco.com/en/US/tech/CK4_36/CK7_98/technologies_configuration_example09186a0080093f82.shtml

QUESTION 366:

The VPN named my_vpn is operating on interface s0/0 of a PE-router. The CE-PE routing protocol is OSPF.

The MPLS backbone IGP is OSPF.

Which statement is true about the interaction between the customer's OSPF routes and the backbone's OSPF routes?

- A. On the PE-router, two instances of OSPF are run: one for the VPN and one for the backbone IGP.
- B. On the PE-router, three instances of OSPF are run: one for the VPN, one for the backbone IGP, and one to carry the PE-PE routes.
- C. This configuration is acceptable for MPLS implementations. However, when MPLS VPNs are implemented, OSPF cannot be used as the CE-PE routing protocol and the backbone IGP.
- D. On the PE-router, a single instance of OSPF is run. However, the customer's routes are kept separate for the backbone route because the RD is prepared to the customer's routes.

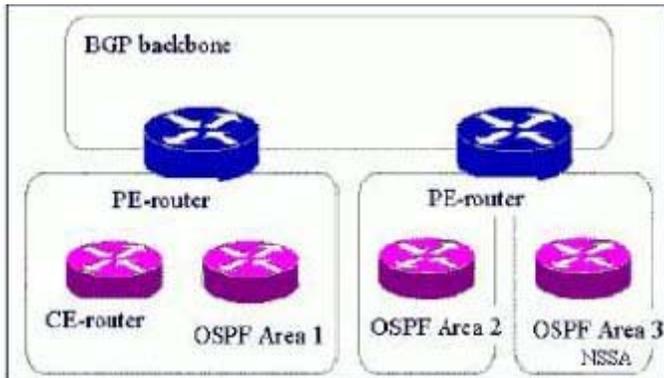
Answer: C

Explanation:

It is important to note that the MPLS/VPN backbone is not a real OSPF area 0 backbone. No adjacencies are formed between PE routers - only between PE and CE routers. MP-iBGP is used between PE routers, and all OSPF routes are translated into VPN-IPv4 routes. This means that the redistribution of routes into BGP does not cause these routes to become external OSPF routes when advertised to other member sites of the same VPN. Reference: MPLS and VPN Architectures (Cisco Press) page 235

QUESTION 367:

Exhibit:



Which statement is true about route summarization for routes announced to the backbone by Area 1?

- A. By default, routes will not be summarized.
- B. By default, routes will be summarized on the network boundaries.
- C. By default, routes will be summarized to Area 3 (NSSA) but not to Area 2.
- D. By default, routes will be summarized based upon the summarization schedule configured in the CE-router.

Answer: B

Explanation:

Summarizing is the consolidation of multiple routes into one single advertisement. This is normally done at the boundaries of Area Border Routers (ABRs). Although summarization could be configured between any two areas, it is better to summarize in the direction of the backbone. This way the backbone receives all the aggregate addresses and in turn will injects them, already summarized, into other areas. There are two types of summarization:

1. Inter-area route summarization
1. External route summarization

Reference: <http://www.cisco.com/warp/public/104/3.html#12.0>

QUESTION 368:

What is the influence of the Down bit on the OSPF route section process?

- A. OSPF routes with the Down bit set are never entered in the M-BGP routing table.
- B. OSPF routes with the Down bit set are never entered in the backbone's IGP routing table.
- C. OSPF routes with the Down bit set are passed up the hierarchical model. They are blocked from being passed down the hierarchical model.
- D. OSPF routes with the Down bit set are marked as unavailable. After the third consecutive update with the Down bit set, they are removed from the

routing table.

Answer: B

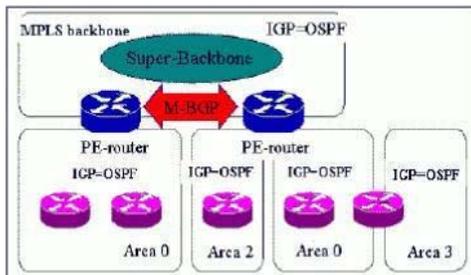
Explanation:

Because inter-site routes may be advertised into a site from various locations, it is necessary to provide a mechanism that allows a PE router to understand whether the route has actually originated from within the attached site, or whether it was injected by another PE router. This mechanism is provided through use of the down-bit, which is an extension to the OSPF protocol and is part of the Options field of the generic LSA header. Any summary LSAs that are generated by the PE routers will have the down-bit set within the LSA.

Reference: MPLS and VPN Architectures (Cisco Press) page 240

QUESTION 369:

Exhibit:



What is the interaction between a super-backbone and Area 3?

- A. The super-backbone appears as a NSSA to the non-backbone OSPF routers of Area 3.
- B. The super-backbone appears as BGP domain to the non-backbone OSPF routers of Area 3.
- C. The super-backbone appears as another OSPF area to the non-backbone OSPF routers of Area 3.
- D. The super-backbone appears as Area 0 (backbone area) to the non-backbone OSPF routers of Area 3.

Answer: D

Explanation:

Traditionally, an elaborate OSPF network consists of a backbone area (area 0) and a number of areas connected to this backbone via an area border router (ABR). By using an MPLS backbone for VPN with OSPF on the customer's site, you can introduce a third level in the hierarchy of the OSPF model. This third level is called the MPLS VPN super backbone. In simple cases, the MPLS VPN super backbone is combined with the traditional area 0 backbone. This means that there is no area 0 backbone on the customer network, since the MPLS VPN super backbone plays the same role as the area 0 backbone. This is

shown in the diagram below:

Reference: http://www.cisco.com/warp/public/121/mpls_ospf2.html

QUESTION 370:

How are OSPF route attributes propagated across the MPLS VPN backbone? (Choose two)

- A. OSPF priority is propagated in the Experimental Bits.
- B. OSPF cost or external metric is propagated in the BGP MED attribute.
- C. Metrics for external OSPF router is propagated in the Hop Count field.
- D. OSPF area, route type, and metric type are propagated in an extended BGP community.

Answer:

Explanation: Pending: ?; D. Send your suggestion to [feedback@ Certkiller .com](mailto:feedback@Certkiller.com)
The redistribution of VRF OSPF routes into MP-iBGP is achieved through use of the address family within the BGP configurations.

Reference: MPLS and VPN Architectures (Cisco Press) page 235

QUESTION 371:

Which statement about operating frame MPLS over an ATM PVC is true?

- A. LDP sessions are established over a reserved PVC.
- B. The attached ATM switch is aware of the MPLS traffic.
- C. A TDP or LDP neighborship is established between the two PVC endpoint routers.
- D. Switching is performed based on the VPI/VCI value in the ATM header established by LDP.

Answer: D

Explanation:

In certain circumstances, such as during a transition to a full IP=ATM MPLS environment, or if transit ATM switches do not support MPLS, it might be desirable to run MPLS in Frame-mode across ATM PVCs. This is perfectly valid configuration although it suffers from the same scaling issues (due to the high number of VCs) as running IP over ATM in an overlay mode.

QUESTION 372:

In frame-relay MPLS, where is the label imposed?

- A. Between Layer 1 and Layer 2 headers.
- B. Between Layer 2 and Layer 3 headers.

- C. In the VPI/VCU fields.
- D. In the VPI field alone.

Answer: C

QUESTION 373:

What is a benefit of CEF switching?

- A. CEF supports IP source prefix-based switching using the FIB.
- B. CEF uses less memory than fast switching uses.
- C. CEF is less CPU intensive than fast switching is.
- D. CEF provides Netflow statistics with minimum CPU overhead.
- E. CEF allows multiple data planes to share a common control plane.

Answer: C

QUESTION 374:

What does the global configuration command interface atm 0/0.1 mpls do?

- A. enables frame-mode MPLS operations on the ATM subinterface
- B. creates a cell-mode LC-ATM subinterface
- C. enables both LDP and TDP on the ATM subinterface
- D. enables LDP only on the ATM subinterface
- E. creates an ATM subinterface that can support both frame-mode and cell-mode operations

Answer: B

QUESTION 375:

Which of the following best describes the function of the OSPF tag field?

- A. prevents cross-domain routing loops for external LSAs.
- B. prevents cross-domain routing loops for internal LSAs.
- C. prevents intra-domain routing loops for external LSAs.
- D. prevents intra-domain routing loops for internal LSAs.

Answer: A

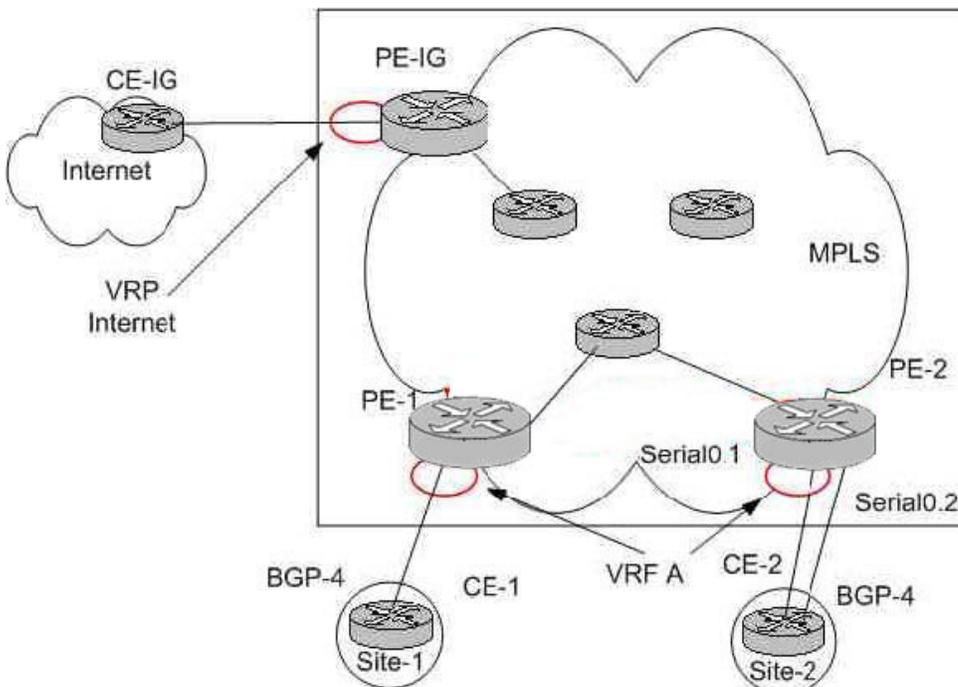
QUESTION 376:

Which statement is true about the use of EBGP as the CE-PE routing protocol in MPLS VPN implementations?

- A. The CE router must be the BGP next hop.
- B. The egress P router must be the BGP next hop.
- C. ebgp-multihop must be configured on the MP-BGP sessions between the PE routers.
- D. next-hop-self must be configured on the MP-IGBP sessions between the PE routers.
- E. The BGP next hops announced using the core IGP can be summarized to reduce the size of the core routing table.

Answer: D

QUESTION 377:



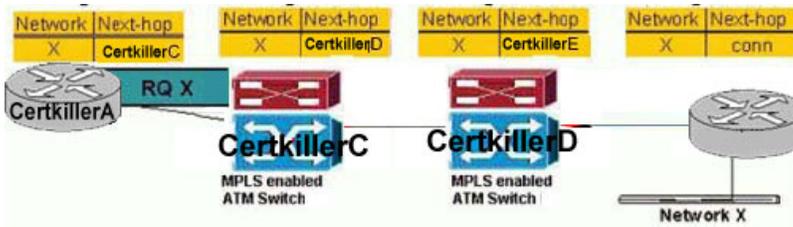
In the diagram, the Internet is accessed through a dedicated Internet VPN. Which routing table will PE-2 use to forward packets from Site-2 to Site-1?

- A. Global routing table
- B. VRF A routing table
- C. The VPNv4 routing table
- D. VRP Internet routing table

Answer: C

QUESTION 378:

Exhibit



Study the exhibit which displays an MPLS cell-mode network that is using default label distribution.

When the label request for Network X is sent from router Certkiller A to the MPLS-enabled ATM Switch Certkiller C, which of the following statements are true?

- A. The MPLS-enabled ATM switch Certkiller C will reply with the local label if it already has the next-hop label. Or it will not reply with the local label if it does not have the next-hop label and it will forward the label request to the MPLS-enabled ATM Switch Certkiller D.
- B. The MPLS-enabled ATM switch Certkiller C will reply with the local label then it will send its own label request to MPLS-enabled switch Certkiller D.
- C. Since MPLS-enabled ATM Switch Certkiller C cannot perform IP routing lookup, it will forward the label request to router Certkiller E (an ATM edge LSR) and Router Certkiller E will then reply with its local label to Router CertkillerA (also an ATM edge LSR).
- D. Since the MPLS-enabled ATM Switch Certkiller C cannot perform IP routing lookup, it will reject the label request from Certkiller A.

Answer: A

QUESTION 379:

In a managed CE router VPN implementation, which three of the following are RT configuration tasks or tasks that would need to be completed? (Choose three)

- A. Configure NMS import / export route target in NMS VRF.
- B. Configure NMS import / export route target in customer VRF.
- C. Import routes with NMS RT into customer VRF.
- D. Export routes with NMS RT into customer VRF.
- E. Export loopback addresses from customer VRF with RT NMS_Client.
- F. Import loopback addresses from customer VRF with RT NMS_Client.

Answer: A, C, E

QUESTION 380:

When the show mpls atm-idp command is used, which keyword will display specified entries from the ATM database?

- A. summary

- B. bindings
- C. capability
- D. database

Answer: B

QUESTION 381:

What does the show mpls atm-ldp capability command display?

- A. Capabilities negotiated between MPLS-ATM and LDP neighbors for label-controlled ATM (LC-ATM) interfaces.
- B. Capabilities specified from the ATM label binding database.
- C. Capabilities of the Multiprotocol Label Switching (MPLS) label forwarding information base (LFIB).
- D. Capabilities of tunnels announced to the Interior Gateway Protocol (IGP), including interface, destination, and bandwidth.

Answer:

QUESTION 382:

DRAG DROP

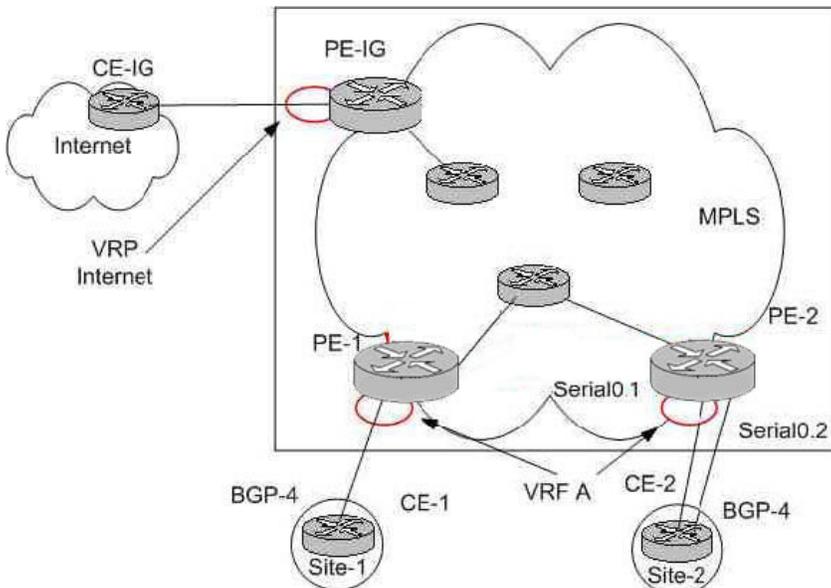
Match each item to its correct description.

VPNv4 address	96 bits long.
Route Target (RT)	Uses to make potentially overlapping IPv4 addresses globally unique.
Route Distinguisher (RD)	Used to prevent loops.
Site of Origin (SOO)	Attached to the VPNv4 routes to indicate VPN membership.

Answer:

96 bits long.	VPNv4 address
Uses to make potentially overlapping IPv4 addresses globally unique.	Route Distinguisher (RD)
Used to prevent loops.	Route Target (RT)
Attached to the VPNv4 routes to indicate VPN membership.	Site of Origin (SOO)

QUESTION 383:



In the diagram, the Internet is accessed through a dedicated Internet VPN. With which router or routers must PE-2 establish an address-family IPv4 BGP neighbor relationship?

- A. CE-2
- B. PE-1
- C. PE-IG
- D. PE-1 and PE-IG
- E. CE-1, PE-1 and PE-IG
- F. CE-1, CE-2, PE-1 and PE-IG

Answer: A

QUESTION 384:

What is the difference in implementation between a managed CE services MPLS VPN and a central services MPLS VPN?

- A. RD assignment
- B. Selective routes export
- C. Selective routes import
- D. MP-BGP route redistribution filtering
- E. CE-PE routing process
- F. None

Answer: B

QUESTION 385:

DRAG DROP

Match the following VPN type to its description.

Simple VPN	Some sites can participate in more than one vpn.	Place here
Overlapping VPN	All sites can communicate with a Common Server but not with each other.	Place here
Central services VPN	Sometimes it is called a "rainbow" VPN implementation.	Place here
Managed network VPN	Every site can communicate with every other site.	Place here

Answer:

Some sites can participate in more than one vpn.	Overlapping VPN
All sites can communicate with a Common Server but not with each other.	Central services VPN
Sometimes it is called a "rainbow" VPN implementation.	Simple VPN
Every site can communicate with every other site.	Managed network VPN

QUESTION 386:

DRAG DROP

Match the benefits on the left to the appropriate Internet access method on the right by dragging and dropping each benefit to the proper access method. Benefits that are not appropriate to either method should remain in the left column.

This method is the easiest to implement	Internet access through a VPN
This method is well known in the SP environment.	Place here
Only the SP PE routers are aware of Internet routing in the method.	Internet access through a separate subinterface
With this method, the P-Net does not need to carry any of the Internet routes.	Place here
	Place here

Answer:

QUESTION 387:

Which three pieces of information are configured within the VRF configuration mode - Router(config-vrf#)? (Choose three)

- A. RD
- B. Import RT
- C. Export RT
- D. Routing protocol address-family
- E. Multiprotocol BGP (MPBGP) neighbor

Answer: A, B, C

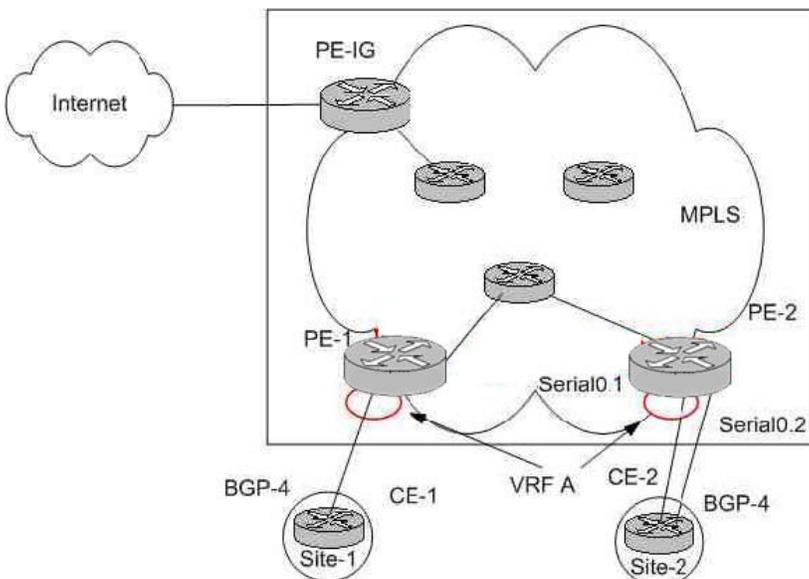
QUESTION 388:

What best describes wholesale Internet access service?

- A. an ISP that uses the IP transport infrastructure of another service provider to reach the customers.
- B. an ISP that provides broadband access (like DSL or cable) directly to the customers.
- C. an ISP that provides both MPLS VPN and Internet access simultaneously to the customers
- D. an ISP that provides a managed firewall service to the customers
- E. and ISP that provides managed CE services to customers

Answer: A

QUESTION 389:



In the diagram, Internet access is through a dedicated subinterface implementation. What indicates that all of the routes to PE-2 must establish a global BGP neighbor relationship?

- A. CE-1
- B. PE-1
- C. PE-IG
- D. CE-2 and PE-IG
- E. CE-1 and PE-2
- F. PE-1 and PE-2

Answer:

QUESTION 390:

If aggregation (summarization) were to be used on a network with ATM LSRs. What would result?

- A. LSPs would be broken in two.
- B. There would be extra LFIB entries.
- C. The size of the LFIB table would increase.
- D. There would be extra LIB entries

Answer: A

QUESTION 391:

Which condition must be met for the receiving PE router to install VPNv4 routes into a VRF?

- A. If at least one RD attached to the VPNv4 route matches at least one import RD configured in the VRF.
- B. If at least one RT attached to the VPNv4 route matches at least one import RT configured in the VRF.
- C. If at least one RD attached to the VPNv4 route matches at least one export RD configured in the VRF.
- D. If at least one RT attached to the VPNv4 route matches at least one export RT configured in the VRF.

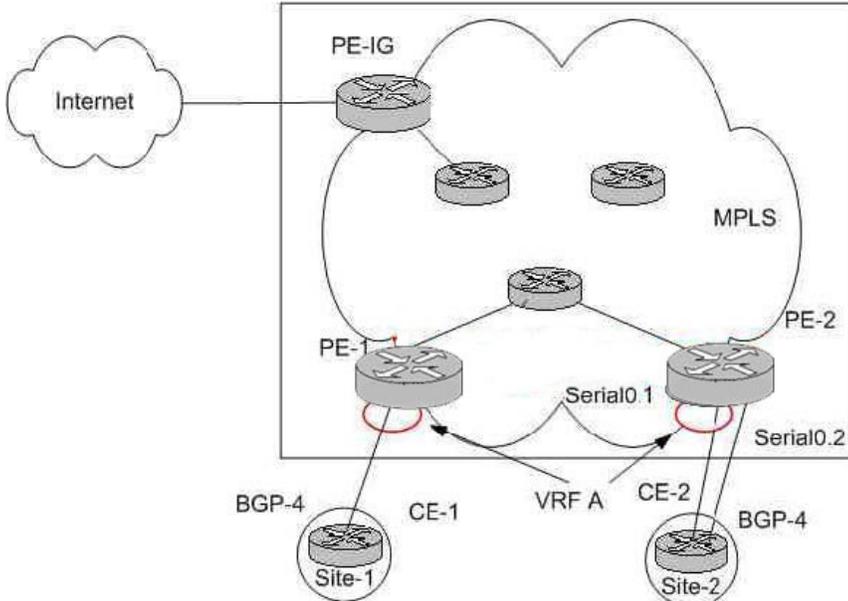
Answer: B

QUESTION 392:

Which one of the following is true regarding MPLS independent control label allocations?

- A. The LSR can always assign a label for a destination prefix, even if it has no downstream label.
- B. The LSR can assign a label for a destination prefix only if it has already receives a label from the next-hop LSR, otherwise, it must request a label from the next-hop LSR.
- C. The LSR will assign a label to a destination prefix only when asked for a label by an upstream LSR.
- D. The label for a destination prefix is allocated and advertised to all LDP peers, regardless of whether the LDP peers are upstream or downstream LSRs for the destination prefix.
- E. The LSR stores the receives label in its LIB, even when the label is not received from the next-hop LSR.
- F. The LSR stores only the labels received from the next-hop LSR, all other labels are ignored.

Answer: A

QUESTION 393:

Refer to the graphic. CE-2 is using a dedicated sub-interface implementation for Internet access. In this case, PE-2 will need to establish VPNv4 BGP neighbor relationship with which other router(s)?

- A. PE-1 only
- B. PE-IG only
- C. PE-1 and PE-IG
- D. PE-1 and CE-2

Answer: A

QUESTION 394:

Which type of overlay VPN is simplest to implement?

- A. Fully-meshed
- B. Hub-and-spoke
- C. Partial-mesh
- D. Redundant hub-and-spoke

Answer:

QUESTION 395:

What is the main implementation option when using the global IP routing table to provide Internet access?

- A. Use a separate logical interface for Internet access that is not placed in a VRF.
- B. Use the same logical interface for Internet and VPN access and disable packet filtering.
- C. Import the Internet routes from the global IP routing table into the VRF.
- D. Export the Internet routes from the global IP routing table into the VRF.
- E. Use packet filtering between the VRF and the VPNv4 MPBGP RIB (routing information base).

Answer: A

QUESTION 396:

Which well-known tunnel technology is fast, simple to implement, supports multiple routed protocols, but it provides no security?

- A. GRE
- B. IPSec
- C. Peer-to-peer
- D. Layer 2 forwarding
- E. Layer 2 tunneling

Answer: A

QUESTION 397:

How is route target (RT) information attached to a VPNv4 route?

- A. Using the MPLS Label field.
- B. Using the MPLS EXP bits.
- C. Using route tags.
- D. Using Extended BGP communities.
- E. Using the Route Distinguisher (RD) field.

Answer: D

QUESTION 398:

How can you configure a router that carries both Internet and VPNv4 routes so that it will only propagate the VPNv4 routes and not the Internet routes to the 10.1.1.1 BGP neighbor?

- A.

```
router bgp 65010
no bgp default ipv4-unicast
neighbor 10.1.1.1 remote-as 65010
address-family vpnv4
neighbor 10.1.1.1 activate
```

! output omitted
B. router bgp 65010
no bgp default ipv4-unicast
neighbor 10.1.1.1 remote-as 65010
neighbor 10.1.1.1 activate
address-family vpnv4
neighbor 10.1.1.1 activate
! output omitted
C. router bgp 65010
neighbor 10.1.1.1 remote-as 65010
address-family vpnv4
neighbor 10.1.1.1 activate
! output omitted
D. router bgp 65010
address-family vpnv4
neighbor 10.1.1.1 activate
! output omitted

Answer: A

Explanation:

Wrong answers:

B: The default address family has been disabled so you have to specify the address family you which to activate before.

C: It will establish address-family ipv4 with the neighbor so it will send the full bgp table

D: Neighbor statement wasn't specified (remote-as)

QUESTION 399:

In situations where cost prohibits having separate physical links for VPN and Internet traffic, what are the two recommended methods for creating two logical links over a single physical link? (Choose two)

- A. The use of Frame Relay subinterfaces.
- B. The use of 802.1q trunking on LAN interfaces.
- C. The use of GRE tunnel interfaces.
- D. The use of PPP multilink interfaces.
- E. The assignment of two VRFs on the same physical interface.

Answer: A, B

QUESTION 400:

Which VPN implementation is achieved with technology such as Frame Relay and ATM?

- A. Layer 1 overlay VPN
- B. Layer 2 overlay VPN
- C. Layer 3 overlay VPN
- D. Peer-to-peer VPN

Answer: B

QUESTION 401:

Which two statements are true about the label stack used for MPLS VPNs? (Choose two)

- A. The first label in the label stack is the VPN label assigned by the egress PE router.
- B. The first label in the label stack is the LDP label used to reach the egress PE router.
- C. The first label in the label stack is the VPN label used to reach the egress PE router.
- D. The second label in the label stack is the VPN label assigned by the egress PE router.
- E. The second label in the label stack is the LDP label used to reach the egress PE router.
- F. The second label in the label stack is the VPN label assigned by the ingress PE router and tells the ingress PE router how to forward the incoming VPN packet.

Answer: B, D

QUESTION 402:

DRAG DROP

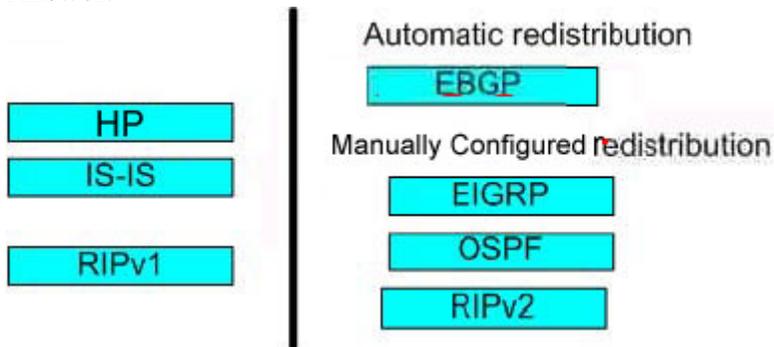
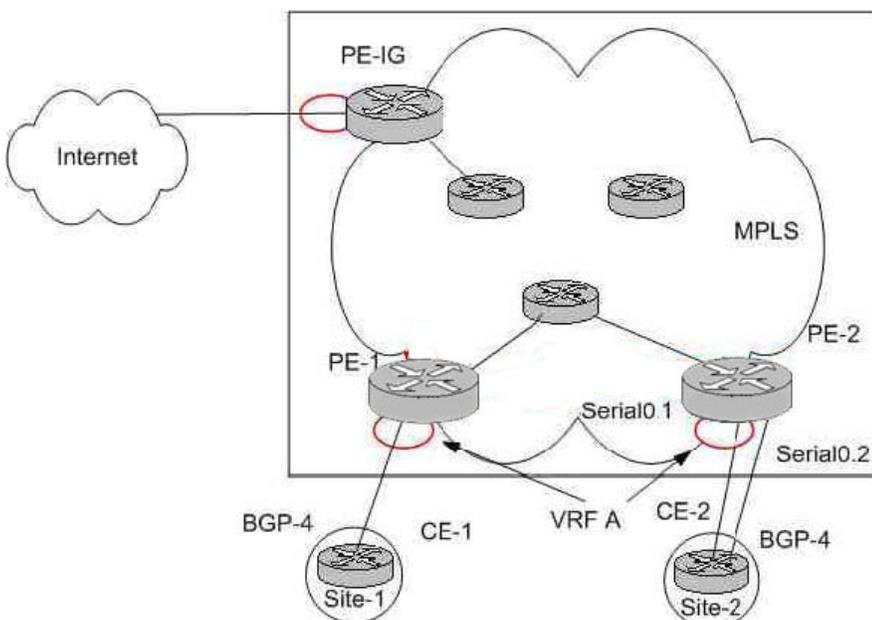
Drag the VRF-aware on the left to the method used on the right that redistributes routes from MP-BGP into the protocol VRF instance.

EBGP	Place here
EIGRP	
IGRP	
IS-IS	
OSPF	
RIPv1	
RIPv2	

Automatic redistribution

Manually Configured redistribution

Answer:

**QUESTION 403:**

In the diagram, Internet access is through a dedicated subinterface implementation. Which of its routing tables will PE-2 use to forward packets from Site-2 to Site-1?

- A. The global routing table.
- B. The VPNv4 routing table.
- C. The VRF A routing table.
- D. The VRF A FIB.
- E. The PE-IG routing table.

QUESTION 404:

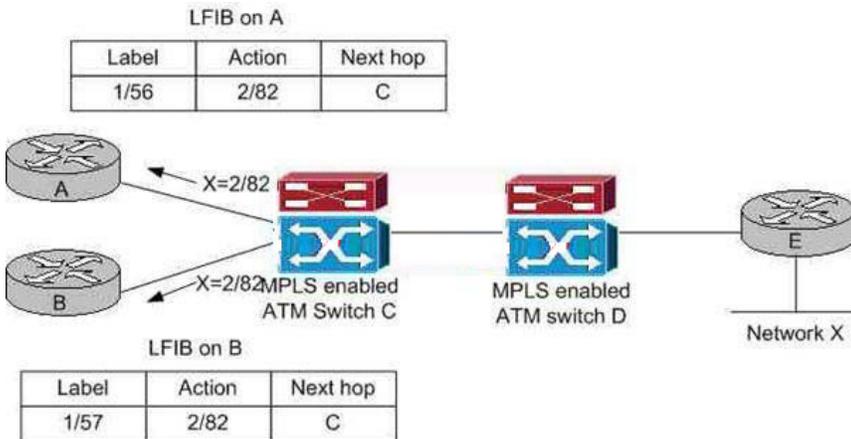
Which term is used to describe a VPN in which the service provider does not participate in the customer routing?

- A. MPLS simple VPN

- B. overlay VPN
- C. shared PE peer-to-peer VPN
- D. dedicated PE peer-to-peer VPN
- E. MPLS complex VPN

Answer: B

QUESTION 405:



Refer to the graphic.

What is a potential problem if the MPLS-enabled ATM switch C reuses an already allocated local label (2/82) for the second upstream LDP peer (router B)?

- A. Cell interleaving issues at egress of MPLS-enabled ATM switch C.
- B. Security concerns at the ingress of the MPLS-enabled ATM switch C. Another router connected to MPLS enabled ATM switch C can send a packet and spoof the label (2/82).
- C. Label lookup issues at the ingress of MPLS-enabled ATM switch C if per-interface label allocation is not used.
- D. Label lookup issues at the ingress of MPLS-enabled ATM switch C if per-platform label allocation is not used.

Answer: A

QUESTION 406:

In MPLS VPN implementations, how are the second (VPN) labels in the label stack propagated from the egress PE router to the ingress PE router?

- A. The core IGP.
- B. MP-IBGP VPNv4 routing updates.
- C. MP-IBGP IPv4 routing updates.
- D. EBGP IPv4 routing updates
- E. LDP

F. TDP

Answer: B

QUESTION 407:

In the set extcommunity rt 123:750 additive command, what does the parameter additive do?

- A. Replaces the existing set of route targets with a new set.
- B. Adds the new route target to the existing route target list but does not replace any existing route targets.
- C. Adds the new route target to the existing route target list and replaces the most current route target.
- D. Adds the new route target to the existing route target list and replaces the least current route target.

Answer: B

QUESTION 408:

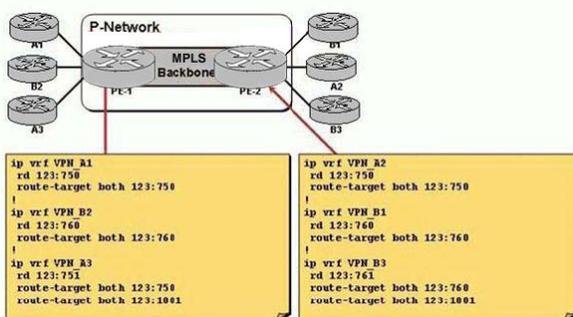
What is true about how MPLS implementations support different applications?

- A. The applications only differ in the control plane.
- B. The forwarding plane is customized for each application.
- C. MPLS VPNs require RSVP.
- D. In general a FEC is assigned to a label.

Answer: A

QUESTION 409:

Exhibit:



Refer to the exhibit.

In which type of MPLS VPN are customers A and B participating?

- A. Overlapping MPLS VPN.

- B. Simple MPLS VPN.
- C. Central services MPLS VPN.
- D. Overlay MPLS VPN.
- E. Managed CE service MPLS VPN.

Answer: A

QUESTION 410:

Which MPLS VPN implementation allows selected sites in one simple VPN to communicate with selected sites of a second VPN?

- A. Central services VPN
- B. Managed CE router services VPN
- C. Overlapping VPN
- D. Managed PE router services VPN

Answer:

QUESTION 411:

Exhibit

```
hostname Certkiller1
!
ip cef
!
ip vrf Customer_Certkiller
  rd 115:300
  route-target both 115:300
!
Access list 10 permit 192.463.3.0.0.0.255
!
route-map CertkillerMAP permit 10
  match ip address 10
  set extcommunity rt 115:301
!
!output omitted
```

Study the configuration shown in the exhibit. Router Certkiller 1 is supposed to selectively attach an additional RT of 115:301 to all export routes matched by access-list 10. The configuration is currently not working as intended.

Which two items are wrong in the configuration? Select two.

- A. The route-map is missing the route-map Certkiller MAP permit 20 statement.
- B. The vrf configuration is missing the export map RTMAP command.
- C. The vrf configuration is missing the route-target export 115:301 command.
- D. The rd 115:300 command is not correct.

- E. The set extcommunity rt 115:301 command is not correct.
- F. The route-target both 115:300 command is not correct.

Answer: B, E

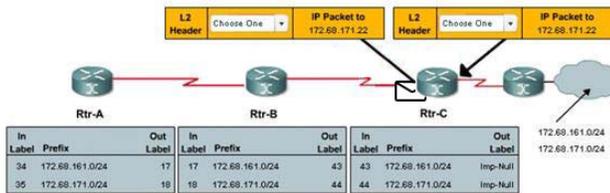
QUESTION 412:

DRAG DROP

For the packet destined to 172.68.171.22 from Rtr-B:

1. What label will be used at the ingress to Rtr-C?
2. What label will be used at the egress from Rtr-C?

Use the drop down box and choose the correct label for each packet.



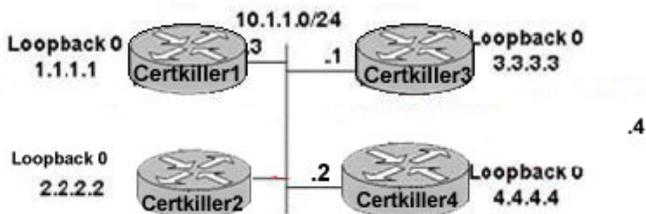
Answer:

Explanation:

The first from left should be 44
 as the Rtr-B's outgoing label is 44 for that particular prefix.
 The 2nd should be Imp null
 as that's happening the incoming packets with label 44.

QUESTION 413:

Exhibit



Study the exhibit. All four routers on the Certkiller Ethernet LAN sent out LDP hello messages.

What is the set of TCP sessions that will be initiated?

- A. Certkiller 1 to Certkiller 2, Certkiller 3 and Certkiller 4
- B. Certkiller 2 to Certkiller 1, Certkiller 3 and Certkiller 4
- C. Certkiller 3 to Certkiller 1, Certkiller 2 and Certkiller 4
- D. Certkiller 4 to Certkiller 1, Certkiller 2 and Certkiller 3

Answer: D

QUESTION 414:

In order for MPLS to be implemented on ATM switches, what requirements must the ATM switch meet? Select two.

- A. become Layer 3 aware by running a routing protocol
- B. use MPLS LDP or TDP to distribute and receive MPLS label information
- C. use BGP to exchange MPLS VPN labels in the data plane
- D. use RSVP to exchange MPLS traffic-engineering labels in the data plane
- E. establish a full mesh of Layer 2 ATM virtual circuits between all the ATM switches in the MPLS domain
- F. use cell-mode MPLS and insert MPLS label in the ATM AAL5 header

Answer: A, B

QUESTION 415:

When a VPN solution is being implemented in which each customer site has its own independent Internet access, which three statements are true? Select three.

- A. Each customer site has to be individually secured against unauthorized Internet access.
- B. Each customer site has optimal traffic flow to and from the Internet sites.
- C. Each customer site requires at least some public IP addresses.
- D. Each customer edge (CE) router will forward the VPN and Internet traffic toward the Internet.
- E. A central firewall service will be required.
- F. The MPLS VPN service provider backbone will be over utilized, because the same traffic can cross the backbone of the service provider twice, first as VPN traffic and then as Internet traffic (or vice versa).

Answer: A, B, C

QUESTION 416:

Which component of MPLS architecture is a simple label-based forwarding engine that is independent of the type of routing protocol or label exchange protocol?

- A. control plane
- B. data plane
- C. routing plane
- D. forwarding plane

Answer: B

QUESTION 417:

Which statements are correct about configuring MPLS over an ATM virtual path on a cell-mode router? Select three.

- A. An LC-ATM interface must be created
- B. The ATM VPI value must be set to the virtual path number.
- C. The virtual path number has to match between the peers.
- D. The control VC must use a VPI value of 1 to match the default VPI value.

Answer: A, B, C

QUESTION 418:

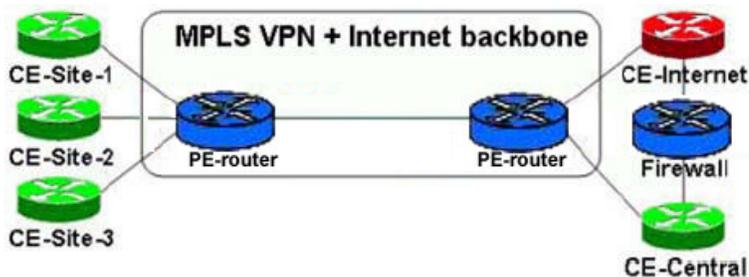
For what purpose can the ip vrf sitemap route-map-name command be used?

- A. applies a route map for setting the SOO
- B. applies a route map for setting the Down Bit
- C. applies a route map for setting the Routing Bit
- D. applies a route map for setting the RD
- E. applies a route map for setting the RT
- F. applies a route map for setting the VRF name

Answer: A

QUESTION 419:

Exhibit



The exhibit shows a topology where the customer, Certkiller .com, is using Internet access through a central firewall at their central site. If Certkiller .com's MPLS VPN and the Internet access traffic share the same service provider backbone, which of the following statements are true? Select two.

- A. Customer Internet traffic can cross the service provider backbone twice, first as VPN traffic and then as Internet traffic.
- B. Each customer site CE router must perform NAT if the customer is using private IP addresses.
- C. Suboptimal routing of the customer's VPN traffic will occur.
- D. The PE routers will need to exchange both VPNv4 and IPv4 BGP routing updates.

E. BGP must be used as the routing protocol between the CE and PE routers.

Answer: A, D

QUESTION 420:

A PE router is running EIGRP with its MPLS VPN customer. What is wrong with the following EIGRP configuration on the PE router?

```
router eigrp 100
address-family ipv4 vrf Certkiller vpn
network 10.0.0.0
redistribute bgp 65100 metric 10000 100 255 1 1500
! output omitted
```

- A. The autonomous-system as-number command is missing under the address-family
- B. The network 10.0.0.0 command should be configured under the router eigrp configuration mode and not under the address-family.
- C. The redistribute bgp 65200 metric 10000 100 255 1 1500 command should be configured under the router eigrp configuration mode and not under the address-family.
- D. The address-family ipv4 vrf Certkiller vpn command should be changed to address-family vpnv4 vrf Certkiller vpn.
- E. The EIGRP AS number (100) does not match the BGP AS number (65100).

Answer: A

QUESTION 421:

Which one of the following MPLS label allocation methods is the least secure?

- A. per-interface label space
- B. per-platform label space
- C. per-VRF label space
- D. per-LIB label space
- E. per-LFIB label space

Answer: B

QUESTION 422:

What can happen if a P router (LSR) within an LSP performs route summarization and suppresses the more specific networks?

- A. Route summarization will cause every LSR along the LSP to perform a routing lookup.

- B. Route summarization will help reduce the size of the LIB and LFIB and will not affect MPLS label switching operations.
- C. Route summarization will cause the P router (LSR) that performs the route summarization to perform PHP (penultimate-hop popping).
- D. Route summarization will break an LSP into two segments, and the P router (LSR) that performs the route summarization will need to perform a routing lookup.

Answer: D

QUESTION 423:

What is the correct configuration to limit the route export to only the loopback address of 10.1.1.1?

- A. ip vft VPN_A
!output omitted
export map NMS
!
route-map NMS
match ip address 1
set extcommunity rt 123:100 additive
!
access-list 1 permit 10.1.1.1 0.0.0
- B. ip vft VPN_A
!output omitted
export map NMS
!
route-map NMS
match ip address 1
set extcommunity rt 123:100 no-export
!
access-list 1 permit 10.1.1.1 0.0.0
access-list 1 permit any
- C. ip vft VPN_A
!output omitted
export map NMS
!
route-map NMS
match ip address 1
set community rt 123:100 additive
!
access-list 1 permit 10.1.1.1 0.0.0
- D. ip vft VPN_A
!output omitted
export map NMS
!

```
route-map NMS
match ip address 1
set community rt 123:100 no-export
!
access-list 1 permit 10.1.1.1 0.0.0
access-list 1 permit any
```

Answer: A

QUESTION 424:

Exhibit

Certkiller7 # show mpls ldp bindings

```
tib entry: 150.1.0.0/16, rev 18
  local binding: tag: 20
  remote binding: tsr: 192.168.1.17:0, tag: 21
  remote binding: tsr: 192.168.1.97:0, tag: 20
tib entry: 150.1.11.32/28, rev 19
  local binding: tag: 21
  remote binding: tsr: 192.168.1.17:0, tag: imp-null
  remote binding: tsr: 192.168.1.97:0, tag: 21
tib entry: 192.168.1.17/32, rev 17
  local binding: tag: 19
  remote binding: tsr: 192.168.1.17:0, tag: imp-null
  remote binding: tsr: 192.168.1.97:0, tag: 19
```

Study the exhibit. Based on the show command output, which four statements are correct? Select four.

- A. Certkiller 7 will announce a label of 20 to its LDP neighbors for prefix 150.1.0.0/16.
- B. To reach prefix 150.1.11.32/28 via 192.168.1.17, Certkiller 7 will pop the label.
- C. To reach prefix 192.168.1.17/32 via 192.168.1.97, Certkiller 7 will use a label of 19.
- D. To reach prefix 150.1.0.0/16 via 192.168.1.97, Certkiller 7 will use a label of 20.
- E. Certkiller 7 will announce an implicit null label to its LDP neighbors for prefix 192.168.1.17/32.
- F. To reach prefix 192.168.1.17/32 via 192.168.1.17, Certkiller 7 will use a label of 19.

Answer: A, B, C, D

QUESTION 425:

Using OSPF as the CE-PE routing protocol, when and which router will set the down bit?

- A. The PE router will set the down bit when redistributing routes from MP-BGP into OSPF.

- B. The PE router will set the down bit when redistributing routes from OSPF into MP-BGP.
- C. The CE router will set the down bit when propagating the OSPF route to the PE router.
- D. The PE router will set the down bit when propagating the OSPF route to the CE router.
- E. The P router will set the down bit when propagating the route to another P router.
- F. The PE router will set the down bit when propagating the route to another PE router over MP-BGP.

Answer: A

QUESTION 426:

An OSPF LSA type 5 route is redistributed into MP-BGP. That same route is then redistributed back from MP-BGP into OSPF on another PE router. In this case, that OSPF route will appear as what LSA type on the destination CE router?

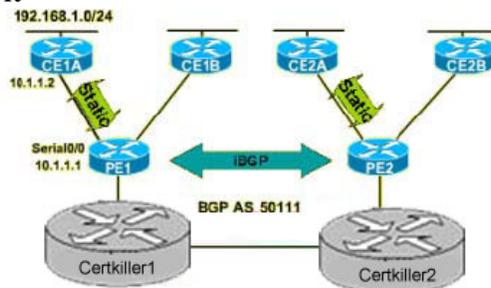
- A. LSA type 1
- B. LSA type 2
- C. LSA type 3
- D. LSA type 4
- E. LSA type 5
- F. LSA type 7

Answer: E

QUESTION 427:

DRAG DROP

Exhibit



Select first command here
Select second command here
Select third command here
Select fourth command here

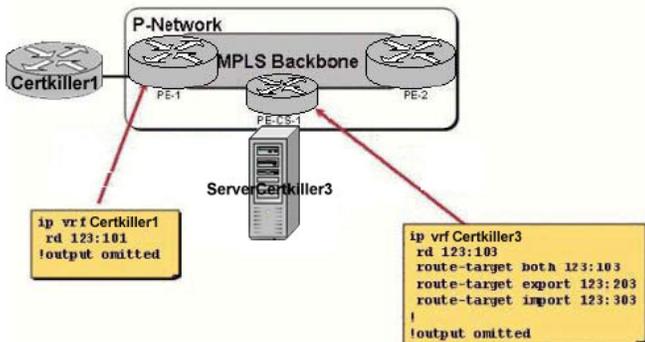
CE1A and CE2A belong to the Certkiller_A VPN. Select the correct commands for

the PE1 router to establish static routing from PE1 to the CE1A 192.168.1.0/24 network by selecting them below. Some of the commands must be entered in the correct order. Also enable propagation of the static route into MP-iBGP. Assume the MP-iBGP session between PE1 and PE2 have been configured already.

Answer:

QUESTION 428:

Exhibit



Study the exhibit. Certkiller 1 is supposed to participate in a central service MPLS VPN. What import and export RTs are required in the Certkiller 1 VRF?

- A. route-target both 123:101
route-target export 123:303
route-target import 123:203
- B. route-target both 123:101
route-target export 123:203
route-target import 123:303
- C. route-target both 123:101
route-target both 123:103
- D. route-target both 123:203
route-target both 123:303
- E. route-target both 123:103
route-target export 123:303
route-target import 123:203
- F. route-target both 123:103
route-target export 123:203
route-target import 123:303

Answer: A

QUESTION 429:

Which two of the following statements regarding LDP are true? Select two

- A. LDP can also be used between nonadjacent routers using multicast LDP hello messages.
- B. LDP does not require periodic hello messages once the LDP session has been established between the LDP peers.
- C. LDP hello messages use TCP packets with a destination port of number of 646.
- D. Multiple sessions can be established between a pair of LSRs if they use multiple label spaces.
- E. Per-platform label space can be identified by label space ID of 0 in the LDP identifier field.

Answer: D, E

QUESTION 430:

Which statement is true about overlapping VPNs?

- A. Sites that participate in more than one VPN import routes with RTs from any VPN in which they participate and export routes with RTs for all VPNs in which they participate.
- B. Sites that participate in more than one VPN import routes with RTs for all VPNs in which they participate and export routes with RTs for none of the VPNs in which they participate.
- C. Sites that participate in more than one VPN only import routes with RTs for all VPNs in which they participate.
- D. None of the above statements are correct regarding overlapping VPNs.

Answer: A

QUESTION 431:

Which IOS show command can be used to examine all the routes that have been redistributed into MP-BGP?

- A. show ip bgp vpv4 all
- B. show ip bgp
- C. show ip route vrf name
- D. show ip protocol vrf name
- E. show ip route bgp
- F. show ip bgp sum

Answer: A

QUESTION 432:

What is a major drawback of using traditional IP routing over an ATM network when connecting multiple sites?

- A. Each ATM switch in the path has to perform Layer 3 routing lookup.
- B. ATM virtual circuits have to be established between the different sites.
- C. There is high ATM management overhead between the ATM switch and the router at each site.
- D. Each ATM switch has to be manually configured to participate in Layer 3 routing.
- E. There is high PNNI overhead.-

Answer: B

QUESTION 433:

What is used to provide virtual intra-area OSPF connectivity across MPLS VPN backbone so the customer's traffic will prefer the MPLS VPN backbone over a backdoor OSPF link between the customer's sites?

- A. OSPF virtual-link
- B. OSPF sham-link
- C. OSPF nssa-link
- D. OSPF stub-link
- E. OSPF demand-circuit
- F. OSPF LSA1-link

Answer: B

QUESTION 434:

What are three drawbacks of a peer-to-peer VPN using a shared provider edge (PE) router? (Choose three)

- A. A full mesh of virtual circuits is required between the customer sites.
- B. All the customers have to share a common IP address space.
- C. Optimal routing between customer sites cannot be guaranteed.
- D. The shared PE router has to know all routes for all customers.
- E. Packet filters are required on the PE routers.

Answer: B, D, E

QUESTION 435:

What are the two major design models for offering Internet access and MPLS VPN services using the same MPLS backbone? (Choose two)

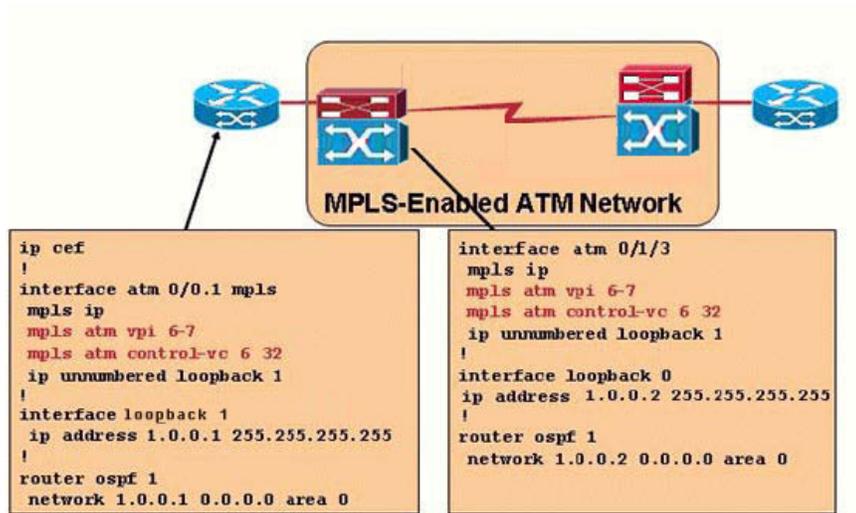
- A. Implementing Internet access as another VPN.
- B. Implementing Internet access through the global IP routing table on the PE routers.
- C. Implementing Internet access through the global VRF routing table on the PE routers.
- D. Implementing Internet access using Wholesale Internet Access Service.

E. Implementing Internet access using packet filtering between the VRF and the global IP routing table on the CE routers.

Answer: A, B

QUESTION 436:

Exhibit



Refer to the exhibit. Which two of the following statements about the MPLS configurations are true? (Choose two)

- A. The VPI range being configured is the default VPI range.
- B. The router is missing the mpls label protocol ldp configuration command on its ATM 0/0.1 subinterface to make it an LC-ARM enabled subinterface.
- C. There is a problem with the configuration because the control VC should be set to 0 32 instead.
- D. The ATM switch is using VC merge since VC merge is enabled by default.
- E. For MPLS label allocations, both VPI 6 and 7 can be used.

Answer: D, E

QUESTION 437:

Which show command is used to verify the proper redistribution of the MPLS VPN EIGRP customer routes into a per-VRF instance if MP-BGP?

- A. show ip bgp vpnv4 vrf vrf-name
- B. show ip bgp
- C. show ip route bgp
- D. show ip route eigrp
- E. show ip vrf vrf-name
- F. show ip vrf detail vrf-name

Answer: A

QUESTION 438:

DRAG DROP

Drag the number of bit(s) on the left to the field of the MPLS header that uses them on the right.

1	Time To Live Field Field length in bits
2	
3	Stack Field Field length in bits
4	
8	Experimental Field Field length in bits
12	
16	
20	Label Field Field length in bits
24	

Answer:

2	Time To Live Field 8
4	Stack Field 1
12	Experimental Field 3
16	
24	Label Field 20

QUESTION 439:

What are the three default frame-mode MPLS label allocation and distribution modes? (Choose three)

- A. per-platform label space
- B. per-interface label space
- C. unsolicited label distribution
- D. downstream-on-demand label distribution
- E. liberal label retention
- F. conservative label retention

Answer: A, C, E

QUESTION 440:

What are the two tunneling technologies used to implement Layer 3 overlay VPNs over an IP backbone? (Choose two)

- A. IPSec
- B. GRE
- C. L2TP
- D. PPTP
- E. AToM
- F. MPLS TE

Answer: A, B

QUESTION 441:

Which IOS command can be used to examine the MPLS VPN label stack for a prefix?

- A. show ip route vrf [vrf-name] [ip-prefix]
- B. show ip cef [ip-prefix]
- C. show ip cef vrf [vrf-name] [ip-prefix]
- D. show ip vrf detail
- E. show ip bgp vpnv4 all
- F. show mpls interfaces

Answer: C

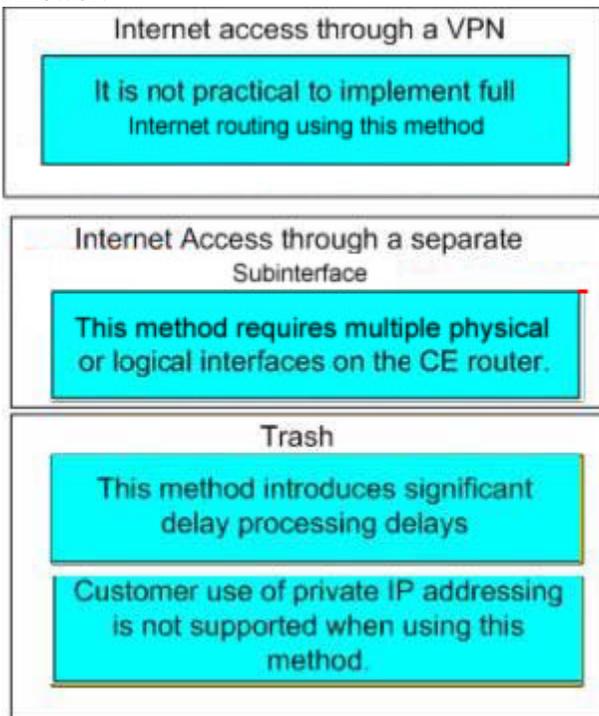
QUESTION 442:

DRAG DROP

Match the drawbacks on the left to the appropriate MPLS Internet access methods on the right by dragging and dropping the drawbacks to the proper access method. Drawbacks that are not appropriate to either method should be placed in the trash.

This method introduces significant delay processing delays	Internet access through a VPN
It is not practical to implement full Internet routing using this method.	Internet Access through a separate subinterface
This method requires multiple physical or logical interfaces on the CE router.	Trash
Customer use of private IP addressing is not supported when using this method.	

Answer:



QUESTION 443:

Which one of the following is a problem solved by VC merge?

- A. Slow MPLS converges issue when using cell-mode MPLS.
- B. Limited number of VCs supported by the router and MPLS-enabled ATM switch hardware.
- C. Security concern when the same local label value is reused on different interfaces.
- D. Loop detection since the ATM header does not contain a TTL field.

Answer: B

QUESTION 444:

A pair of MPLS routers, operating in the default frame-mode, are connected over two parallel links. Which of the following describes what will happen when they try to establish LDP sessions?

- A. Only one LDP session will be established.
- B. Multiple LDP sessions, one for each link, will be established and the TCP traffic will be load balanced between the links.
- C. Multiple LDP sessions, one for each link, will be established, but only the one with the highest IP address will be active.
- D. Multiple LDP sessions, one for each link, will be established, but only the one with the best metric will be active.

Answer: A

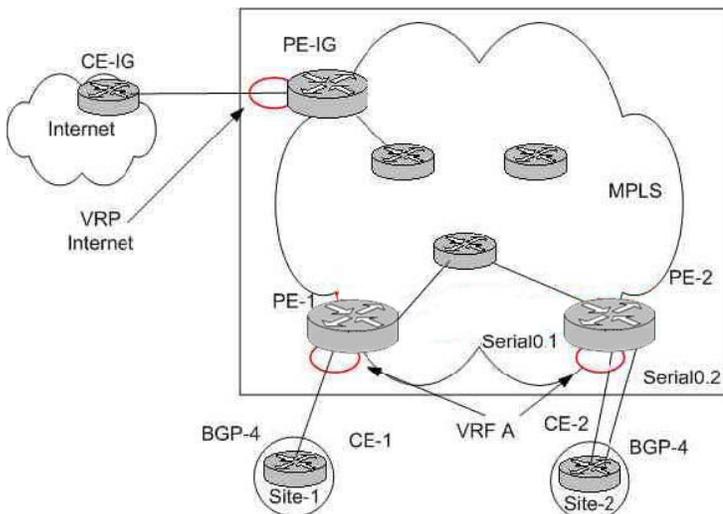
QUESTION 445:

What does the Cisco ODAP for MPLS VPN feature enable?

- A. VRF aware NAT services
- B. VRF aware on-demand IP address pool
- C. VRF aware Hot Standby Router Protocol
- D. VRF aware OSPF domain area protocol
- E. VRF aware OSPF database access procedures
- F. VRF aware DHCP relay services

Answer: B

QUESTION 446:



In the diagram, the Internet is accessed through a dedicated Internet VPN implementation. With which two routers must PE-2 establish a VPNv4 BGP neighbor relationship? (Choose two)

- A. CE-1
- B. CE-2
- C. PE-1
- D. CE-IG
- E. PE-IG

Answer: C, E

QUESTION 447:

To troubleshoot an MPLS VPN, the show ip route [vrf-name] command should be entered on which router(s)?

- A. P
- B. CE
- C. PE
- D. CE or PE
- E. P or PE
- F. P or PE or CE

Answer: C

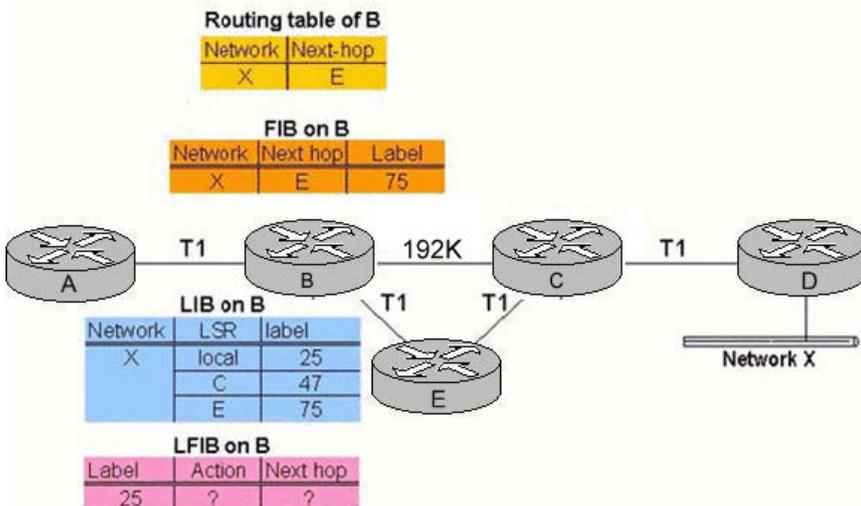
QUESTION 448:

Which two statements are true about a service provider managed central firewall service? (Choose two)

- A. Traffic between the customer sites and the Internet are typically restricted only by the security functions of the service provider managed central firewall.
- B. Traffic between sites of an individual customer must flow inside the customer's VPN.
- C. Traffic between the different customers (Inter-customer VPN traffic) should be allowed.
- D. Customers that still want to retain their own private IP addresses inside their network must use NAT on the service provider managed central firewall.
- E. Service providers typically use public IP addresses as the address space between the service provider managed central firewall and the customers.

Answer: A, B

QUESTION 449:

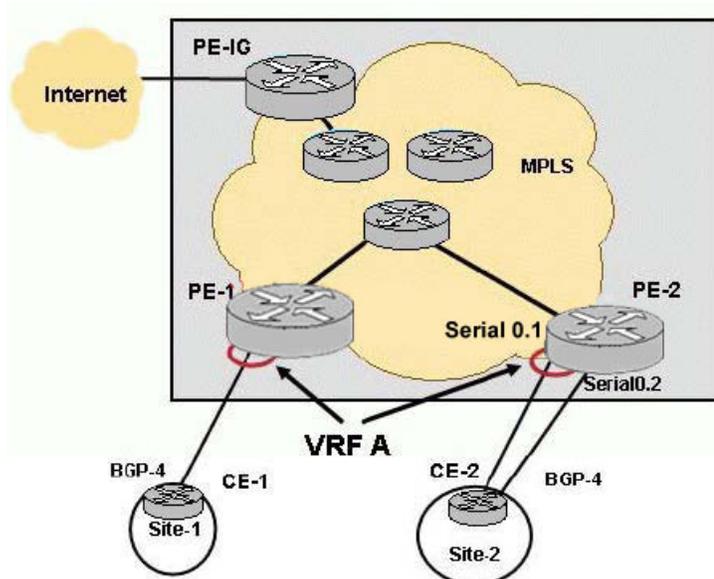


Refer to the graphic. At router B, complete the LFIB for a label of 25.

- A. Action: label of 47 and Next hop: router C
- B. Action: label of 75 and Next hop: router E
- C. Action: label of 25 and Next hop: router C
- D. Action: label of 25 and Next hop: router E
- E. Action: 47 and 75 and Next hop: router C and router E

Answer: B

QUESTION 450:



On a dedicated subinterface implementation, PE-2 must establish an address-family vrf IPv4 BGP neighbor relationship with which router?

- A. CE-1
- B. CE-2
- C. PE-1
- D. PE-IG
- E. CE-1 and CE-2
- F. PE-1 and PE-IG

Answer: B

QUESTION 451:

What is wrong with the following BGP router configuration?

```
router bgp 65001
no bgp default ipv4-unicast
!
address-family vpnv4
```

```
neighbor 192.168.1.1 remote-as 65001
neighbor 192.168.1.1 update-source loopback 0
neighbor 192.168.1.1 activate
neighbor 192.168.1.1 next-hop-self
neighbor 192.168.1.1 send-community both
```

- A. The neighbor 192.168.1.1 remote-as 65001 and the neighbor 192.168.1.1 update-source loopback 0 commands should be moved under the bgp router configuration mode.
- B. The neighbor 192.168.1.1 activate command should be moved under the bgp router configuration mode.
- C. The no bgp default ipv4-unicast command needs to be removed.
- D. The neighbor 192.168.1.1 remote-as 65001, the neighbor 192.168.1.1 update-source loopback 0, and the neighbor 192.168.1.1 activate commands should be moved under the bgp router configuration mode.
- E. Nothing is wrong. The router will be able to propagate VPNv4 routes to the 192.168.1.1 MP-BGP neighbor.

Answer: A

QUESTION 452:

When nonadjacent LDP neighbors are used for implementing an MPLS traffic engineering solution, how are the nonadjacent LDP neighbors discovered?

- A. using multicast CR-LSP (constraint-based LSP)
- B. using unicast CR-LSP (constraint-based LSP)
- C. using multicast LDP hello messages
- D. using unicast LDP hello messages
- E. using multihop MP-IBGP
- F. using multihop MP-EBGP

Answer: D

QUESTION 453:

DRAG DROP

Drag and drop the correct commands in the proper order to redistribute the RED VPN routes from the BGP VRF into the EIGRP VRF.

router eigrp 1	Command 1 (place here)
router bgp 65001	
no auto-summary	Command 2 (place here)
redistribute eigrp 1	
address-family ipv4 vrf RED	Command 3 (place here)
address-family vpnv4 vrf RED	
redistribute eigrp 1 metric transparent	
redistribute bgp 65001 metric transparent	
redistribute bgp 65001 metric 1000 100 255 1 1500	

Answer:

Command 1 router eigrp 1
Command 2 address-family ipv4 vrf RED
Command 3 redistribute bgp 65001 metric 1000 100 255 1 1500

QUESTION 454:

Refer to the exhibit. Based on the show outputs, which condition could be preventing the P1 router from establishing TDP adjacency with its neighbor over the s0/0.211 and s0/0.212 subinterfaces?

```
P1#sh mpls ldp neighbor

P1#sh mpls interfaces
Interface          IP           Tunnel   Operational
Serial0/0.211      Yes (tdp)    No       Yes
Serial0/0.212      Yes (tdp)    No       Yes
```

- A. The s0/0.211 and s0/0.212 subinterfaces line protocol are in the down state.
- B. The P1 router cannot establish a TCP session with its neighbors.

- C. The P1 router is missing the mpls label protocol LDP command.
- D. The show mpls tdp neighbor command needs to be used to view the TDP neighbor status.

Answer: B

QUESTION 455:

DRAG DROP

Drag the item on the left to the correct table it pertains to on the right.

part of the MPLS control plane	<p style="text-align: center;">LIB</p> <div style="background-color: #FFFF00; padding: 2px; text-align: center;">(place here)</div> <div style="background-color: #FFFF00; padding: 2px; text-align: center;">(place here)</div>
shadow copy of the IP routing table and contains outgoing label information	<p style="text-align: center;">FIB</p> <div style="background-color: #FFFF00; padding: 2px; text-align: center;">(place here)</div> <div style="background-color: #FFFF00; padding: 2px; text-align: center;">(place here)</div>
used to forward incoming labeled packets which are sent out as labeled packets	<p style="text-align: center;">LFIB</p> <div style="background-color: #FFFF00; padding: 2px; text-align: center;">(place here)</div> <div style="background-color: #FFFF00; padding: 2px; text-align: center;">(place here)</div>
MPLS data plane table that maps an incoming label to te outgoing label	
part of the DEF architecture	
stores the local and outgoing label per destination prefix	

Answer:

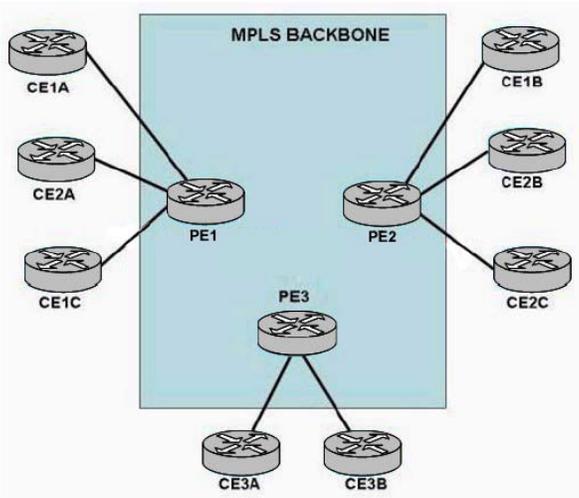
<p style="text-align: center;">LIB</p> <div style="background-color: #00FFFF; padding: 2px; text-align: center;">part of the MPLS control plane</div> <div style="background-color: #00FFFF; padding: 2px; text-align: center;">stores the local and outgoing label per destination prefix</div>
<p style="text-align: center;">FIB</p> <div style="background-color: #00FFFF; padding: 2px; text-align: center;">shadow copy of the IP routing table and contains outgoing label information</div> <div style="background-color: #00FFFF; padding: 2px; text-align: center;">part of the CEF architecture</div>
<p style="text-align: center;">LFIB</p> <div style="background-color: #00FFFF; padding: 2px; text-align: center;">used to forward incoming labeled packets which are sent out as labeled packets</div> <div style="background-color: #00FFFF; padding: 2px; text-align: center;">MPLS data plane table that maps an incoming label to te outgoing label</div>

QUESTION 456:

Refer to the exhibit. How many different VRFs are required to support the given connectivity requirements? Sites CE1A and CE1B require connectivity to each other. Sites CE2A and CE2B require connectivity to each other. Site CE1C requires connectivity to sites CE1A, CE1B, CE3A, and CE3B. Site CE2C requires connectivity to sites CE2A, CE2B, CE3A, and CE3B. Sites CE3A and CE3B require

connectivity to each other and to CE1C and CE2C.

Exhibit:



- A. 3 VRFs
- B. 4 VRFs
- C. 5 VRFs
- D. 6 VRFs
- E. 8 VRFs

Answer: C

QUESTION 457:

With MPLS VPNs, if the CE-PE routing protocol is BGP, which statement below is true?

- A. The different customer sites must use different autonomous system numbers.
- B. Manual route redistribution between EBGP and MP-BGP is required.
- C. The SOO extended BGP community can be used to prevent routing loops.
- D. A BGP virtual sham-link is required between the customer site's CE routers to ensure optimal routing between the customer sites.
- E. The customer's autonomous system number must match the MPLS VPN service provider's autonomous system number.

Answer: C

QUESTION 458:

BGP is used as the CE-PE routing protocol in an MPLS VPN. The customer routes are successfully propagated from the CE router to the ingress PE router, but they are not showing up on the egress PE router. What is the most likely cause of the problem?

- A. There is a BGP-to-MPBGP redistribution error on the ingress PE router.

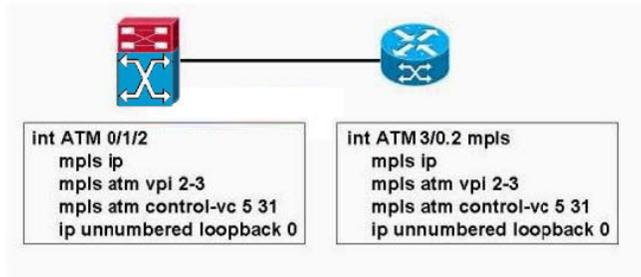
- B. There is a MPBGP-to-BGP redistribution error on the egress PE router.
- C. One of the P routers is performing route summarization, breaking the path between the ingress and the egress PE routers into two LSPs.
- D. RTs attached to the CE routes exported by the ingress router are not matched by at least one of the import RTs on the egress PE router.
- E. A route reflector is used to propagate the routes from the ingress PE router to the egress PE router, thus breaking the LSP.

Answer: D

QUESTION 459:

Refer to the exhibit. A diagram of a router connected to an MPLS-enabled ATM switch via an LC-ATM MPLS interface, and a partial configuration for the MPLS-enabled ATM switch and router are shown. Which statement describes what is incorrect about the configuration shown?

Exhibit:



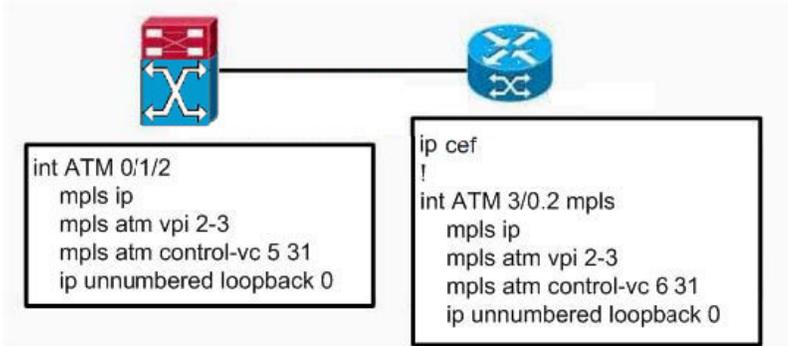
- A. CEF has not been enabled on the router.
- B. The VPI range of 2-3 is invalid.
- C. The control VPI/VCI has not been set to 0/32 on the router.
- D. VC-merge has not been enabled on the ATM switch interface.
- E. The router has not been configured to specifically use LDP.

Answer: A

QUESTION 460:

Refer to the exhibit. A diagram of a router connected to an MPLS-enabled ATM switch via an LC-ATM MPLS interface, and the partial configuration for the MPLS-enabled ATM switch and router are shown. Which statement describes what is incorrect about the configuration?

Exhibit:



- A. The control VPI/VCI numbers do not match.
- B. The ATM VPI range of 2-3 is invalid.
- C. CEF has not been enabled on the router ATM 3/0.2 subinterface.
- D. VC-merge has not been enabled on the switch ATM 0/1/2 interface.
- E. The router has not been configured to specifically use LDP.

Answer: A

QUESTION 461:

What are three characteristics of overlay VPNs? (Choose three.)

- A. Service provider infrastructure appears as point-to-point links to the customer routers.
- B. Routing protocols run directly between the customer routers.
- C. Implementing optimum routing between customer sites requires a partial mesh of virtual circuits.
- D. Service provider PE routers use route filtering to isolate between different customers.
- E. Service provider does not participate in customer routing.

Answer: A, B, E

QUESTION 462:

DRAG DROP

Drag and drop the correct commands in the proper order to redistribute the RED VPN routes from the EIGRP vrf into the BGP vrf.

router eigrp 1	Command 1 (place here)
router bgp 65001	
no auto-summary	Command 2 (place here)
redistribute eigrp 101	
autonomous-system 101	Command 3 (place here)
address-family ipv4 vrf RED	
address-family vpnv4 vrf RED	
redistribute eigrp 101 metric 1000 100 255 1 100	
redistribute bgp 65001 metric 1000 100 255 1 100	

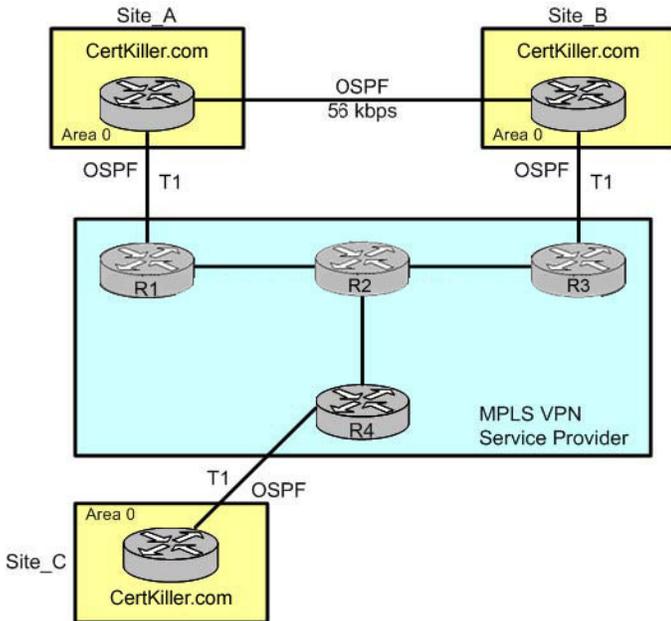
Answer:

router eigrp 1	Command 1 router bgp 65001
no auto-summary	Command 2 address-family ipv4 vrf RED
autonomous-system 101	Command 3 redistribute eigrp 101
address-family vpnv4 vrf RED	
redistribute eigrp 101 metric 1000 100 255 1 100	
redistribute bgp 65001 metric 1000 100 255 1 100	

QUESTION 463:

Refer to the exhibit. Customer XYZ is using an MPLS VPN to connect its locations of Site_A, Site_B, and Site_C. Each site uses OSPF as the CE-PE routing protocol. After a new direct 56 kbps link was added between Site_A and Site_B, all applications degraded in response time only between Site_A and Site_B. What is the most probable cause and its solution?

Exhibit:



- A. A routing loop needs to be corrected by an OSPF down bit.
- B. A routing loop needs to be corrected by the OSPF tag field.
- C. Suboptimal routing needs to be corrected by a virtual link.
- D. Suboptimal routing needs to be corrected by a sham link.

Answer: D

QUESTION 464:

When configuring Internet access using a separate MPLS VPN, which three statements are correct? (Choose three.)

- A. The Internet backbone is separate from the MPLS VPN backbone.
- B. Two dedicated physical or logical links between the PE and the CE routers are required.
- C. An Internet gateway is connected as a CE router to the MPLS VPN backbone.
- D. An Internet gateway shall insert full Internet routing into the Internet VPN to achieve optimal routing.
- E. The customer's Internet access is enabled by combining the Internet VPN with the Customer VPN using overlapping VPN topology.

Answer: A, C, E

QUESTION 465:

When configuring Internet access from a MPLS VPN using a separate subinterface, which two configurations are required on the PE router? (Choose two.)

- A. Use the ip vrf forwarding command on the subinterface that is used for MPLS VPN

access.

- B. Use the ip vrf forwarding command on the subinterface that is used for Internet access.
- C. The CE router is configured as a BGP neighbor in both the global BGP process and inside the VPN in the VRF.
- D. The PE-to-CE VPN and Internet routing are configured using two separate address-family ipv4 vrf commands, one for VPN routing and one for Internet routing.

Answer: A, C

QUESTION 466:

What is the purpose of the global configuration command, ip dhcp relay information option vpn?

- A. enables the DHCP relay agent to insert the VPN suboptions to the BOOTP request
- B. enables the DHCP relay agent to convert the broadcast DHCP request to a unicast DHCP request to a shared DHCP server
- C. enables the DHCP relay agent to perform VRF-aware NAT before forwarding the DHCP request to a shared DHCP server
- D. enables ODAP (On-Demand Address Pool) on the DHCP relay agent

Answer: A

QUESTION 467:

Using MPLS unicast IP forwarding, what will happen if an LSR receives an unlabeled incoming packet?

- A. It will process switch the packet by doing a routing table lookup.
- B. It will forward the packet using the LFIB.
- C. It will forward the packet using the LIB.
- D. It will forward the packet using the FIB.
- E. It will drop the packet immediately.

Answer: D

QUESTION 468:

When routes between MP-BGP and RIP are being redistributed, what does the option metric transparent do?

- A. copies the BGP MED into the RIP hop count
- B. copies the BGP local preference into the RIP hop count
- C. copies the RIP hop count into the BGP local preference
- D. copies the RIP hop count into the BGP MED
- E. uses the default seed metric when redistributing from MP-BGP into RIP

F. uses the default seed metric when redistributing from RIP into MP-BGP

Answer: A

QUESTION 469:

For which purpose is the command `mpls ldp maxhops` used?

- A. In large ATM-MPLS networks, the LFIB can become too large and it may be necessary to limit the maximum diameter of the MPLS LSPs.
- B. Because downstream-on-demand label allocation uses hop count to control loop detection, it maybe necessary to limit the maximum diameter of the MPLS network.
- C. Because end-to-end delay can cause problems with some voice applications, it may be necessary to limit the maximum diameter of the MPLS network.
- D. When interconnecting large frame mode MPLS and cell mode networks it may be necessary to limit the maximum network diameter to prevent forwarding loops.

Answer: D

Explanation:

Usage Guidelines

When an ATM label switch router (LSR) initiates a request for a label binding, it sets the hop count value in the Label Request message to

1. Subsequent ATM LSRs along the path to the edge of the ATM label switching region increment the hop count before forwarding the Label Request message to the next hop.

When an ATM LSR receives a Label Request message, it does not send a Label Mapping message in response, nor does it propagate the request to the destination next hop if the hop count value in the request equals or exceeds the maxhops value. Instead, the ATM LSR returns an error message that specifies that the maximum allowable hop count has been reached. This threshold is used to prevent forwarding loops in the setting up of label switch paths across an ATM region.

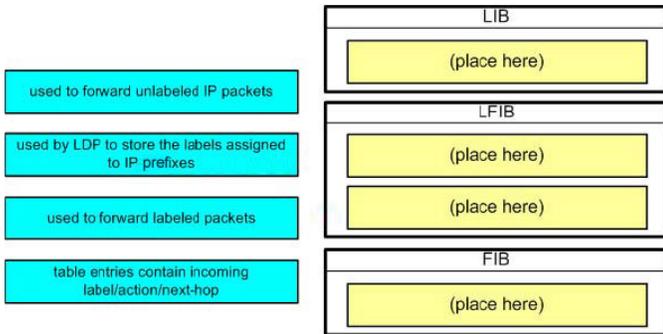
source:

http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/products_command_reference_chapter09186a008046

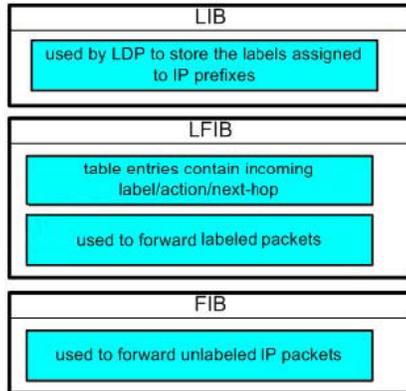
QUESTION 470:

DRAG DROP

Drag the description on the left to the name of the table on the right to which it best corresponds.



Answer:



QUESTION 471:

What is wrong with the following MP-BGP to EIGRP route redistribution configuration?

```
router eigrp 100
address-family ipv4 vrf XYZ
autonomous-system 200
redistribute bgp 51001
```

- A. The redistribute command should be under the router eigrp 100 process and not under the address-family.
- B. The redistribute command is missing the metric option.
- C. The redistribute command is missing the subnet option.
- D. The redistribute command is missing the metric transparent option.
- E. The redistribute command is missing the vrf option option.

Answer: B

QUESTION 472:

Which statement correctly describes managed CE routers service?

- A. The export map {route-map-name} command is used within the customer's VRF so only the loopback interface on the CE routers are exported to the central server VPN.
- B. The managed CE routers service allows the service provider to have access to the CE

router and to all of the networks behind the CE router.

C. The managed CE routers service uses the same VRF and RD design as overlapping MPLS VPN.

D. The managed CE routers service uses the same VRF and RD design as central services MPLS VPN combined with simple MPLS VPN, except for the CE router Route Target marking process during route import.

Answer: A

QUESTION 473:

DRAG DROP

Drag the term on the left to its description on the right. Not all the terms on the left are used.

Frame-mode MPLS	Enables routers to store all received labels even if they are not being used.
Liberal retention mode	Determines the network convergence time.
IGP	Uses liberal retention mode.
Cell-mode MPLS	
Conservative retention mode	
LDP	

Answer:

	Liberal retention mode
	IGP
	Frame-mode MPLS
Cell-mode MPLS	
Conservative retention mode	
LDP	

QUESTION 474:

Which statement represents the overall impact of convergence in a frame-mode MPLS network when there is a failure on a link that an LSP is using?

A. Overall convergence is not affected because, by default, frame-mode MPLS uses liberal retention mode.

B. The LSPs must be rebuilt by the LDP, which minimally impacts convergence regardless of the number of LSPs.

C. The LSPs must be rebuilt by the LDP, and the impact of convergence depends upon the number of LSPs using the failed link.

D. The LSPs must be rebuilt along with the LFIB and LIB on each LSR on the path, which can dramatically impact convergence.

Answer: A

QUESTION 475:

Which three statements correctly describe how labels are used in cell-mode MPLS? (Choose three.)

- A. The ATM header VPI/VCI field is used for forwarding decisions.
- B. The 32-bit MPLS label is preserved in the frame but is not used in the ATM network.
- C. The original MPLS label is present only in the first cell of a packet.
- D. The 32-bit MPLS label is inserted between the ATM header and the ATM AAL5 header.
- E. Labels in cell-mode MPLS are not a scarce resource like in frame-mode MPLS.
- F. The TTL field in the MPLS label is copied into the ATM header TTL field by the ingress edge LSR.

Answer: A, B, C

QUESTION 476:

When using MPLS unicast IP forwarding, what will happen if an LSR receives an incoming labeled packet but the LSR can't find that incoming label in its LFIB?

- A. The packet will be forwarded using the FIB.
- B. The packet will be forwarded using the LIB.
- C. The packet will be process switched by performing a route lookup in the routing table.
- D. The packet will be forwarded using the LFIB with an imp-null outgoing label.
- E. The packet will be dropped even if the IP destination exists in the FIB.

Answer: E

QUESTION 477:

What is the purpose of the area {area-id} sham-link {source-address} {destination-address} cost {number} command?

- A. to create an OSPF virtual link between two CE routers
- B. to allow the customer sites to use the same OSPF area ID
- C. to enable optimized routing over the MPLS backbone when there is a direct backdoor link between two CE routers
- D. to allow the MPLS backbone to act as an OSPF Superbackbone
- E. to prevent routing loops when using OSPF as the CE-PE routing protocol
- F. to allow the customer sites to use area 0 as the OSPF area ID

Answer: C

QUESTION 478:

Which two of the following describe the actions that the OSPF routing process on a PE router will take based on the down bit being set? (Choose two.)

- A. The routing bit will be set so the OSPF routes with the down bit set will be preferred over the MP-BGP learned routes.
- B. The routing bit will be cleared so the OSPF routes with the down bit set never enter the IP routing table.
- C. OSPF routes with the down bit set will not be redistributed from OSPF into MP-BGP.
- D. OSPF routes with the down bit set will be redistributed from OSPF into the MPLS VPN OSPF super-backbone as intra-area routes.
- E. OSPF routes with the down bit set will be redistributed from OSPF into the MPLS VPN OSPF super-backbone as type 5 LSAs.

Answer: B, C

QUESTION 479:

The following partial output is a PE router configuration related to the OSPF process. During troubleshooting, you notice that the OSPF routes from the remote sites are not showing up in the local CE router. What is a problem that can cause this issue?
router ospf 1 vrf ABC network 0.0.0.0 255.255.255.255 area 0 redistribute bgp 55111

- A. The output is missing the metric option in the redistribute command.
- B. The output is missing the metric-type option in the redistribute command.
- C. The output is missing the subnet option in the redistribute command.
- D. The output is missing the vrf option in the redistribute command.
- E. The OSPF network statement is not correct.
- F. The OSPF routing process is missing the address-family ipv4 vrf ABC configuration.

Answer: C

QUESTION 480:

An OSPF route is redistributed into MP-BGP with the following extended community: Extended Community: RT:1:10 OSPF RT:0:1:0 Which parameter(s) will the OSPF RT:0:1:0 indicate?

- A. OSPF cost value
- B. OSPF metric type
- C. OSPF down bit

- D. OSPF routing bit
- E. OSPF area ID and LSA type
- F. OSPF domain ID and OSPF version

Answer: E

QUESTION 481:

DRAG DROP

Drag the proper commands in the proper order to configure EBGP as the CE to PE routing protocol for VPN RED on a PE router.

router bgp 65001	
no bgp default ipv4 unicast	Command 1 (place here)
address-family ipv4 vrf RED	
address-family vpnv4 vrf RED	Command 2 (place here)
neighbor 10.10.10.1 active	
neighbor 10.10.10.1 remote as 65001	Command 3 (place here)
neighbor 10.10.10.1 remote as 65002	
neighbor 10.10.10.1 update-source loopback 0	

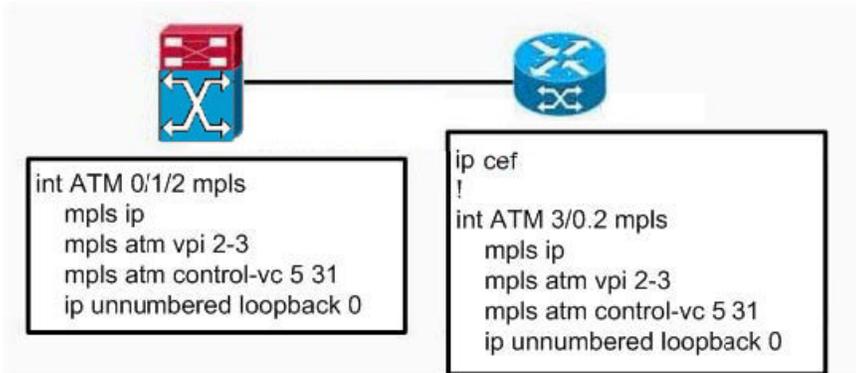
Answer:

no bgp default ipv4 unicast	Command 1 router bgp 65001
address-family vpnv4 vrf RED	Command 2 address-family ipv4 vrf RED
neighbor 10.10.10.1 active	
neighbor 10.10.10.1 remote as 65001	Command 3 neighbor 10.10.10.1 remote as 65002
neighbor 10.10.10.1 update-source loopback 0	

QUESTION 482:

Refer to the exhibit. A router connected to an MPLS-enabled ATM switch via an LC-ATM MPLS interface, and a partial configuration for the MPLS-enabled ATM switch and the router are shown. Which statement describes what is incorrect about the configurations shown?

Exhibit:



- A. The router subinterface is not specified as an MPLS cell-mode subinterface.
- B. CEF has not been enabled on the ATM interface of the MPLS-enabled ATM switch.
- C. The control VPI/VCI has not been set to 0/32.
- D. VC-merge has not been enabled on the ATM interface of the MPLS-enabled ATM switch.
- E. The router has not been configured to specifically use LDP.
- F. The ATM interface of the MPLS-enabled ATM switch has not been configured for cell-mode MPLS operations.

Answer: A

QUESTION 483:

An OSPF LSA type 1 route is redistributed into MP-BGP. That same route is then redistributed back from MP-BGP into OSPF on another PE router. In this case, which LSA type on the destination CE router will the OSPF route appear as?

- A. LSA type 1
- B. LSA type 2
- C. LSA type 3
- D. LSA type 4
- E. LSA type 5
- F. LSA type 7

Answer: C

QUESTION 484:

When using the redistribute ospf {process-id} command, which types of OSPF routes will be redistributed into BGP?

- A. all internal (interarea and intra-area) OSPF routes
- B. all external OSPF routes
- C. all external OSPF routes except type 7 LSAs
- D. both internal and external OSPF routes
- E. only classful OSPF routes without the subnet option

F. only external OSPF routes without the match option

Answer: A

QUESTION 485:

DRAG DROP

Drag and drop the proper commands in the proper order to create a MP-BGP session between two PE routers that will be used to support the RED VPN.

router bgp 65001	Command 1 (place here)
address-family vpnv4	Command 2 (place here)
autonomous-system 65001	Command 3 (place here)
address-family ipv4 vrf RED	Command 4 (place here)
neighbor 172.16.1.2 activate	Command 5 (place here)
neighbor 172.16.1.2 no-default-ipv4	
neighbor 172.16.1.2 remote-as 65001	
neighbor 172.16.1.2 remote-as 65002	
neighbor 172.16.1.2 update-source loopback 0	

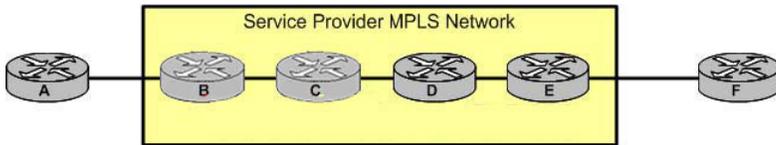
Answer:

	Command 1 router bgp 65001
autonomous-system 65001	Command 2 neighbor 172.16.1.2 remote-as 65001
address-family ipv4 vrf RED	Command 3 neighbor 172.16.1.2 update-source loopback 0
neighbor 172.16.1.2 no-default-ipv4	Command 4 address-family vpnv4
neighbor 172.16.1.2 remote-as 65002	Command 5 neighbor 172.16.1.2 activate

QUESTION 486:

Refer to the exhibit. It shows an MPLS network with four routers as well as two customer routers. Router E is advertising the network 192.168.12.0/24. Rather than propagating the 192.168.12.0/24 network, Router D aggregates the 192.168.12.0/24 route into 192.16.0.0/16. Which route or routes will Router B receive?

Exhibit:



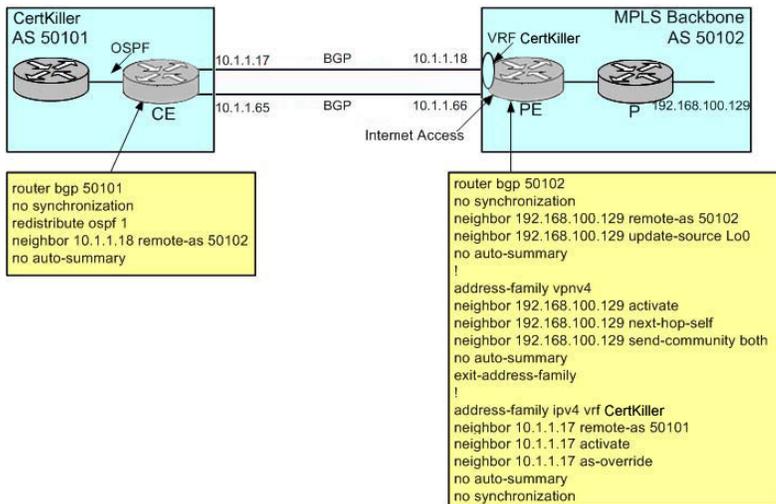
- A. only 192.168.0.0/16
- B. only 192.160.12.0/24
- C. both 192.168.0.0/16 and 192.160.12.0/24
- D. neither of the routes-the LSP is broken

Answer: A

QUESTION 487:

Refer to the exhibit. The MPLS VPN Customer A is using a separate interface for Internet access. However, with the current configurations shown, the CE router is not receiving any Internet routes from the PE router. Which two additional configuration commands can resolve the Internet connectivity issue? (Choose two.)

Exhibit:



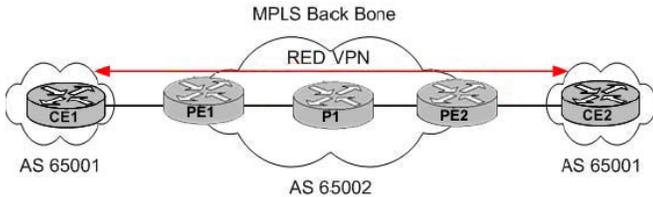
- A. At the CE router, under router bgp 50101, add the neighbor 10.1.1.66 remote-as 50102 command.
- B. At the CE router, under router bgp 50101, add the network 0.0.0.0 command.
- C. At the CE router, under router bgp 50101, add the ip route 0.0.0.0 0.0.0.0 10.1.1.66 command.
- D. At the PE router, under address-family ipv4 vrf Certkiller , add the neighbor 10.1.1.65 remote-as 50101 command.
- E. At the PE router, under address-family ipv4 vrf Certkiller , add the neighbor 10.1.1.17 default-originate command.
- F. At the PE router, under router bgp 50102, add the neighbor 10.1.1.65 remote-as 50101 command.

Answer: A, F

QUESTION 488:

Refer to the exhibit. The customer is experiencing routing problems because both sites are using an AS number of 65001. As the ISP, which neighbor command would you use to address this issue and on which two routers would you configure that command? (Choose three.)

Exhibit:



- A. CE1
- B. CE2
- C. PE1
- D. PE2
- E. allowas-in
- F. as-override

Answer: C, D, F

QUESTION 489:

A CE router running EBGP with the MPLS VPN service provider's PE router is not able to receive any BGP updates from the PE router, and the debug ip bgp updates on the CE router shows the following debug message. BGP(0): 150.1.11.18 rcv UPDATE about 150.1.12.16/28 -- DENIED due to: AS-PATH contains our own AS
What can be configured on the PE router to solve this problem?

- A. Configure the PE router as a route reflector with the CE router as one of its clients.
- B. Configure AS-Path Prepending on the PE router.
- C. Configure AS-Override on the PE router.
- D. Configure AllowAS-In on the PE router.
- E. Configure AllowAS-Out on the PE router.

Answer: C

QUESTION 490:

Which two of the following statements correctly describe the difference between Edge LSRs and LSRs? (Choose two.)

- A. LSRs have all their interfaces enabled for MPLS.
- B. LSRs can only perform label switching and not IP routing.
- C. Edge LSRs remove labels and forward IP packets out of the MPLS domain.
- D. Edge LSRs only forward labeled packets and perform label swapping.

E. Edge LSRs have all their interfaces enabled for MPLS.

Answer: A, C

QUESTION 491:

With MPLS VPN-aware NAT, what additional information is tracked inside the NAT translation table?

- A. RD information
- B. RT information
- C. VRF information
- D. Multi-protocol BGP prefixes
- E. MPLS Labels

Answer: C

QUESTION 492:

Which three fields must be included in MPBGP updates between PE routers in an MPLS network? (Choose three.)

- A. VPNv4 address
- B. label used for VPN packet forwarding
- C. AS path
- D. Site of Origin
- E. TE path designator

Answer: A, B, C

QUESTION 493:

Which three statements correctly describe a route target (RT)? (Choose three.)

- A. An RT is an extended BGP community.
- B. An RT is a 64-bit prefix prepended to a IPv4 address enabling different customers to share the same IP address pool.
- C. An RT is attached to IPv4 routing updates for indication of MPLS VPN membership.
- D. Multiple RTs can be attached to a single VPNv4 BGP routing update.
- E. Each VRF on a PE router can have multiple import RTs configured.
- F. Export RTs are manually configured on the CE routers.

Answer: A, D, E

QUESTION 494:

Within the VRF configuration mode, the import map and the export map commands are used for which purpose?

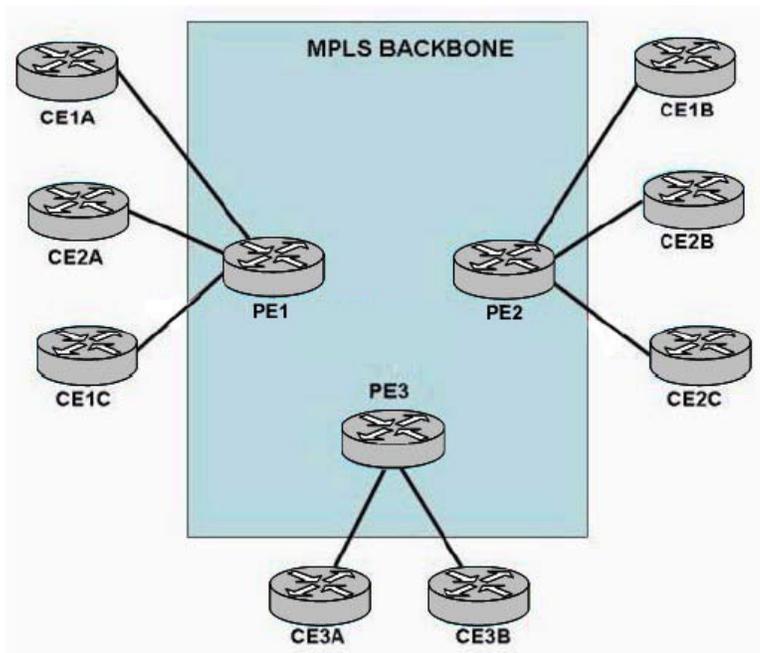
- A. to reference the import and export route-maps that are used to limit the maximum number of routes in a VRF
- B. to reference the import and export route-maps that are used to selectively specify additional criteria for route-targets import and export
- C. to reference the import and export route-maps that are used to filter the routes that are redistributed into MPBGP
- D. to reference the import and export route-maps that are used to set loop avoidance parameters such as SOO
- E. mandatory VRF commands to specify the import and export route-targets

Answer: B

QUESTION 495:

Refer to the exhibit and the following connectivity requirements. Which type of MPLS VPN is being implemented? Sites CE1A and CE1B require connectivity to each other. Sites CE2A and CE2B require connectivity to each other. Site CE1C requires connectivity to sites CE1A, CE1B, CE3A, and CE3B. Site CE2C requires connectivity to sites CE2A, CE2B, CE3A, and CE3B. Sites CE3A and CE3B require connectivity to each other and to CE1C and CE2C.

Exhibit:



- A. central services MPLS VPN
- B. overlapping MPLS VPN
- C. simple VPN
- D. combinations of overlapping and central services MPLS VPN
- E. combinations of simple and central services MPLS VPN

Answer: D

QUESTION 496:

Which four of the following questions need to be answered during the MPLS VPN troubleshooting process? (Choose four.)

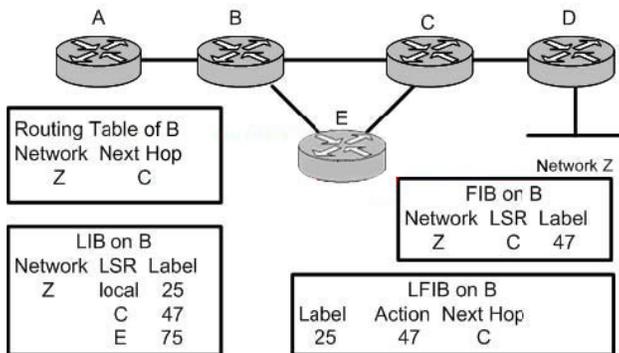
- A. Are the VPNv4 routes propagated to the other PE routers?
- B. Are the CE routes received by the PE and inserted into the PE VRF?
- C. Are the CE routes received by the PE router redistributed into the PE global BGP table?
- D. Are the VPNv4 routes redistributed from MPBGP into the PE-CE routing protocol?
- E. Are the VPNv4 routes inserted into the VRF on the other PE routers?
- F. Are the PE routers enabled for propagating IPv4 routing updates between them?

Answer: A, B, D, E

QUESTION 497:

Refer to the exhibit. Routers A and D are the edge LSRs. If the link between Routers B and C has failed, what best represents the state of the LFIB table on Router B after convergence?

Exhibit:



- A. Label 25, Action 75, Next Hop E
- B. Label 47, Action pop, Next Hop E
- C. Label 75, Action 25, Next Hop E
- D. Label 47, Action pop, Next Hop C

Answer: A

QUESTION 498:

The MPLS VPN provider is using a separate Internet VPN to provide Internet access to its MPLS VPN customers. Currently, all of its customers are not able to access the Internet. The partial provider's CE-Internet router configuration is shown below. Based on the configuration, what could be the cause of the problem?

```
ip route 0.0.0.0 0.0.0.0 10.1.1.1
!  
router bgp 51001  
network 0.0.0.0  
neighbor 10.2.2.2 remote-as 51002  
!  
PE neighbor router  
neighbor 10.2.2.2 prefix-list test1 out  
neighbor 10.3.3.3 remote-as 51001  
!  
Another Internet router  
neighbor 10.3.3.3 prefix-list test2 out  
!  
ip prefix-list test2 permit 0.0.0.0/0  
ip prefix-list test1 permit 0.0.0.0/0  
ge 1
```

- A. The test2 prefix-list should be applied to the 10.2.2.2 neighbor and the test1 prefix-list should be applied to the 10.3.3.3 neighbor.
- B. Both prefix-lists should be applied in the "in" direction instead of the "out" direction.
- C. The network command is missing the mask 0.0.0.0 option.
- D. The network command is missing the mask 255.255.255.255 option.
- E. The neighbor 10.2.2.2 activate and the neighbor 10.3.3.3 activate commands are missing.

Answer: A

QUESTION 499:

What benefit does AToM provide to the service provider's customers?

- A. By supporting Layer 2 VPNs, customers maintain control of their site-to-site routings over the WAN.
- B. By supporting Layer 3 VPNs, a full mesh of virtual circuits will not be required between the different customer sites to enable optimal routing.
- C. By supporting secured Layer 3 VPNs, customers do not have to deal with the complexity of configuring IPsec.
- D. By supporting MPLS traffic engineering over ATM, customers can better utilize their WAN link.
- E. By supporting Diff-Serv QoS, AToM allows customers to deploy voice/video applications across the WAN.

Answer: A

QUESTION 500:

Which of the following could be called a VPN identifier in the MPLS/VPN architecture?

- A. route target
- B. route distinguisher
- C. VRF
- D. VPN IPv4 address
- E. BGP site-of-origin (SOO) extended community attribute

Answer: A

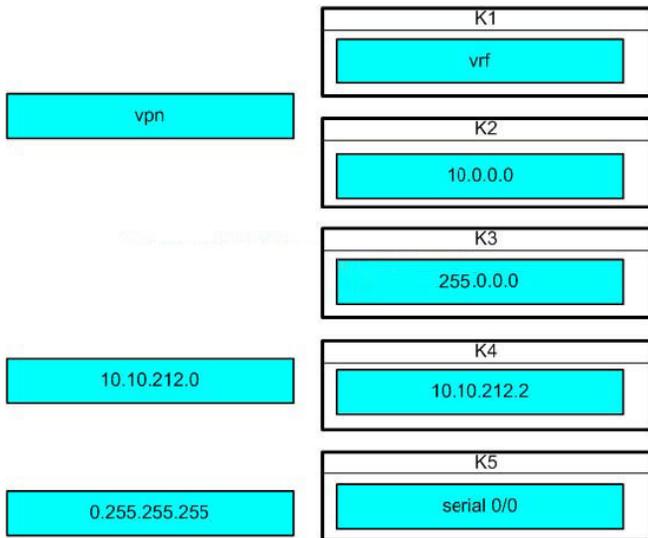
QUESTION 501:

DRAG DROP

Drag and drop the proper keywords to complete the command that will create a static route based upon the following. Destination network = 10.0.0.0/8, Next hop = 10.10.212.2, Output interface = serial 0/0, VPN = RED ip route [k1] RED [k2] [k3] [K4] [K5]

vrf	K1 (place here)
vpn	K2 (place here)
10.0.0.0	K3 (place here)
255.0.0.0	K4 (place here)
serial 0/0	K5 (place here)
10.10.212.0	
10.10.212.2	
0.255.255.255	

Answer:



QUESTION 502:

What are two concerns when implementing CEF switching? (Choose two.)

- A. increased CPU utilization from maintaining the FIB table
- B. increased memory requirement
- C. the requirement to disable other IOS features such as NBAR and MQC
- D. increased memory requirement on the VIP when implementing DCEF
- E. configuration complexity
- F. troubleshooting complexity because of the many tables that CEF maintains

Answer: B, D

QUESTION 503:

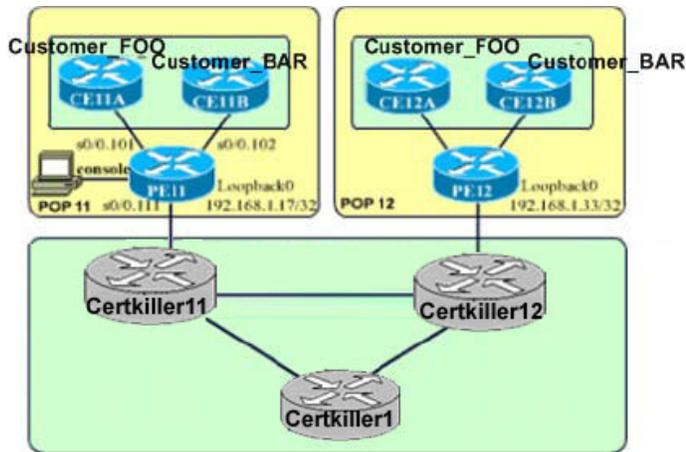
Within an MPLS domain, which of the following has the most impact on network convergence time?

- A. the LDP convergence time
- B. the IGP convergence time
- C. the IGP to MPLS redistribution process
- D. the time required to rebuild the LIB table
- E. the time required to cache the routing table entries into the FIB table

Answer: B

Case Study Certkiller .com, Scenario

Exhibit, Network Topology



Certkiller .com is providing MPLS VPN service to Customer. _FOO and Customer_BAR. Customer_FOO sites CE11A and CE12A belong in a simple VPN. Customer_BAR sites CE11B and CE12B belong in ** missing**

Case Study Certkiller .com, 6 Questions

QUESTION 504:

At the PE11 router, how many BGP routes are in the Customer_FOO VRF routing table?

- A. 1
- B. 2
- C. 3
- D. 4
- E. 5

Answer: C

QUESTION 505:

The PE11 router has how many BGP neighbors?

- A. 1
- B. 2
- C. 3
- D. 4
- E. 5

Answer: A

QUESTION 506:

What is the BGP next-hop and MED for the 10.2.12.15/28 prefix in the

Customer_BAR VRF? Select two.

- A. Next-hop = 192.168.1.33
- B. Next-hop = 150.1.11.33
- C. Next-hop = 192.168.1.17
- D. MED = 0
- E. MED = 1
- F. MED = 100

Answer: A, E

QUESTION 507:

What routing protocol is used between the PE11 and CE11A routes?

- A. EIGRP
- B. EBGp
- C. RIPv2
- D. Static Routing

Answer: C

QUESTION 508:

In the PE11 router global routing table, the EIGRP routes are used for what purpose?

- A. VPNv4 routing
- B. IGP routing within Certkiller .com network
- C. PE-CE routing
- D. To carry the customer routes across Certkiller .com network.
- E. Internet routing using a separate Internet VPN.

Answer: B

QUESTION 509:

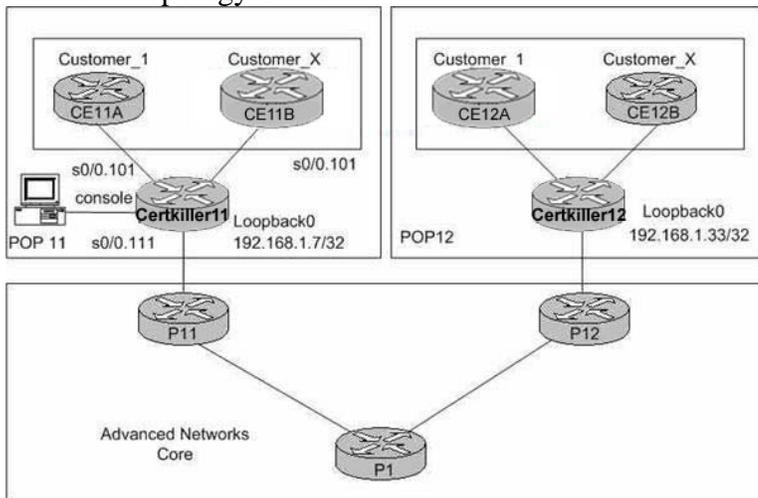
How many BGP routes are in PE11's global IP routing table?

- A. 0
- B. 1
- C. 2
- D. 6
- E. 12

Answer: A

Case Study Certkiller .com Japan, Scenario

Network Topology Exhibit.



You work as network administrator at the Nagasaki office of the Japanese branch of Certkiller .com. Certkiller .com is providing MPLS VPN Service to a customer. Customers sites CE11A and CE12A belong in a simple VPN. CE11B and CE12B are not used by the simulator.

Case Study Certkiller .com Japan, 5 Questions

QUESTION 510:

For Certkiller 11 to reach IP address 192.168.1.33 (Certkiller 12 loopback), what outgoing label will be imposed (used)?

- A. 16
- B. 17
- C. 18
- D. 19
- E. 20
- F. Implicit Null

Answer: C

QUESTION 511:

What local label will Certkiller 11 announce to its LDP peer for reaching the 192.168.1.17/32 prefix (Certkiller 11 own loopback)?

- A. 16
- B. 17
- C. 18

- D. 19
- E. 20
- F. Implicit Null

Answer: F

QUESTION 512:

What routing protocol is used between CE11A and Certkiller 11?

- A. EBGp
- B. IBGP
- C. Multi-Protocol BGP
- D. EIGRP AS 65010
- E. EIGRP AS 65020

Answer: E

QUESTION 513:

On Certkiller 11, how many prefixes are learned from Certkiller 12 (192.168.1.33) over the MPBGP session?

- A. 0
- B. 1
- C. 4
- D. 7

Answer: C

QUESTION 514:

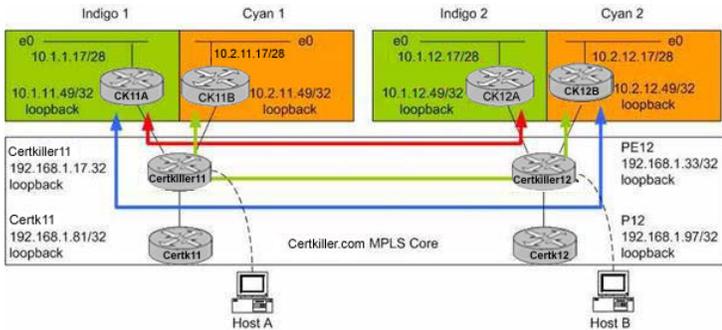
What label or labels will be imposed by Certkiller 11 to reach network 10.1.12.64/28 in the Customer_1 VRF?

- A. 18 23
- B. 20 18
- C. 18 only
- D. 20 only
- E. 23 only
- F. implicit Null

Answer: A

Case Study Certkiller .com Tallinn, Scenario

Network topology exhibit



You are a CCNP certified technician working at the Tallinn branch of Certkiller .com.

The colored lines represent VPNs membership. Certkiller .com is a service provider that is providing advanced VPN solutions using MPLS VPNs.

Currently Certkiller .com is activating three VPNs. The RED VPN is a sample VPN that connects the site Indigo 1 with Indigo 2.

Case Study Certkiller .com Tallin, 4 Questions

QUESTION 515:

Indigo Company has indicated that they cannot establish connectivity between Indigo 1 and Indigo2.

Testing has validated that there is MPLS connectivity between the Certkiller routers. Based upon the output of the show commands, what configuration changes must be made to correct the Indigo connectivity issue?

- A. On Certkiller 11, configure a BGP VPN4 neighbor session between Certkiller 11 and Certkiller 12.
- B. On Certkiller 12, configure a BGP VPN4 neighbor session between Certkiller 12 and Certkiller 11.
- C. On Certkiller 11, configure a rd of 1:10 in the Indigo_Cyan VRF.
- D. On Certkiller 12, configure a rd of 1:20 in the Cyan_Indigo VRF.
- E. On Certkiller 11, configure a rd of 1:20 in the Cyan_1 VRF.
- F. On Certkiller 12, configure a rd of 1:20 in the Indigo_2 VRF.

Answer: A

QUESTION 516:

What import and export route target would you add to the Cyan_Indigo vrf on Certkiller 12 to enable the overlapping VPN?

- A. 1:10
- B. 1:20
- C. 4:11
- D. 4:20

E. 4:21
F. 1:1001

Answer: F

QUESTION 517:

Based upon the information that you can obtain from Certkiller 11 and Certkiller 12, what is the routing protocol used between Certkiller 11 and the Indigo 1 site?

- A. BGP
- B. OSPF
- C. EIGRP
- D. RIPv2
- E. ISIS

Answer: A

QUESTION 518:

Based upon the information that you can obtain from Certkiller 12, what type of BGP routing updates will Certkiller 12 send to Certkiller 11?

- A. Certkiller 11 and Certkiller 12 will only exchange IPv4 routing updates.
- B. Certkiller 11 and Certkiller 12 will only exchange VPNv4 routing updates.
- C. Certkiller 11 and Certkiller 12 will only exchange IPv4 and VPNv4 routing updates.

Answer: B

QUESTION 519:

Based on this configuration, which two peering router neighbor statements are correct? (Choose two.)

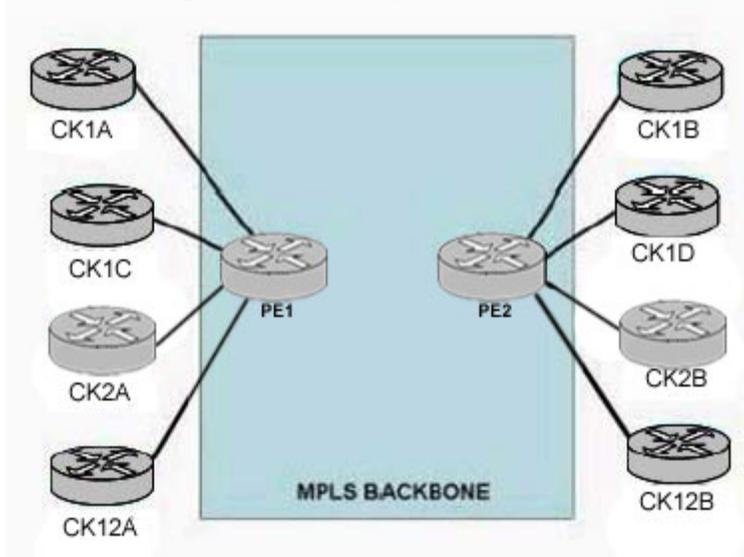
```
routerbgp 50001
neighbor192.168.1.1 remote-as 50001
neighbor10.1.1.1 remote-as 50002
neighbor10.1.1.1 local-as 50003
!output omitted
```

- A. EBGP - neighbor 10.1.1.2 remote-as 50003
- B. EBGP - neighbor 10.1.1.2 remote-as 50001
- C. EBGP - neighbor 10.1.1.2 remote-as 50001 andneighbor 10.1.1.2 local-as 50003
- D. IBGP - neighbor 192.168.1.2 remote-as 50001
- E. IBGP - neighbor 192.168.1.2 remote-as 50003
- F. IBGP - neighbor 192.168.1.2 remote-as 50003 and neighbor 192.168.1.2 local-as 50001

Answer: C, D

QUESTION 520:

Refer to the exhibit and the following connectivity requirements. How many different VRFs are required? Sites CK1 A, CK1 B, CK1 C, and CK1 D require connectivity among them. Sites CK2 A and CK2 B require connectivity between them. Site CK1 2A requires connectivity to sites CK1 A, CK1 B, CK1 C, CK1 D, and CK1 2B. Site CK1 2B requires connectivity to sites CK2 A, CK2 B, and CK1 2A.



- A. 2 VRFs
- B. 3 VRFs
- C. 4 VRFs
- D. 6 VRFs
- E. 8 VRFs
- F. 10 VRFs

Answer: C

QUESTION 521:

Given the AS-path of (51002 51003) 51001 i from the show ip bgp output, what is the origin?

- A. AS 51001
- B. AS 51002
- C. AS 51003
- D. (51002 51003)
- E. IGP
- F. IBGP

Answer: E

QUESTION 522:

What best describes the following configuration example of allowas-in?

```
routerbgp 100
address-familyipv4 vrf CustomerA
neighbor195.12.4.5 remote-as 123
neighbor195.12.4.5 activate
neighbor195.12.4.5 allowas-in 2
```

- A. permits incoming BGP updates defined by access-list 2
- B. permits incoming BGP updates defined by class-map 2
- C. permits incoming BGP updates defined by route-map 2
- D. permits incoming BGP updates with no more than two occurrences of AS 100 in the AS path
- E. permits incoming BGP updates with no more than two occurrences of AS 123 in the AS path

Answer: D