



Exam : 642-661

Title : Configuring BGP on Cisco Routers (BGP)

Ver : 09-26-07

---

**QUESTION 1:**

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 2:**

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

Explanation:


An autonomous system (AS) is a collection of networks under a single technical administration. Some other definitions refer to a collection of routers or IP prefixes, but in the end they all mean the same entity. The important principle is the technical administration, which means sharing the same routing protocol and routing policy. Legal and administrative ownership of the routers does not matter in terms of autonomous systems. Autonomous systems are identified by AS numbers, 16-bit unsigned quantities ranging from 1 to 65535. Public AS numbers are assigned by Internet registries. For customers that need AS numbers to run BGP in their private networks, a range of private AS numbers (64512 - 65535) has been reserved.

Reference: Configuring BGP on Cisco Routers volume 1 p.2-4

---

**QUESTION 3:**

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 4:

Exhibit:



How should the EBGP session between R1 and R2 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 R1 and R2 to reach the other router's loopback.
- D. Establish a single EBGP session using the loopback0 interface IP address on R1 and R2.
- E. Establish two EBGP sessions using both the S0 and S1 interface IP address on R1 and R2.

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 one EBGP path. Would need two if static routes were not defined and were going between different routers
2. E Not an efficient way of doing things requires double the tables and bandwidth for updates

#### QUESTION 5:

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 6:**

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?

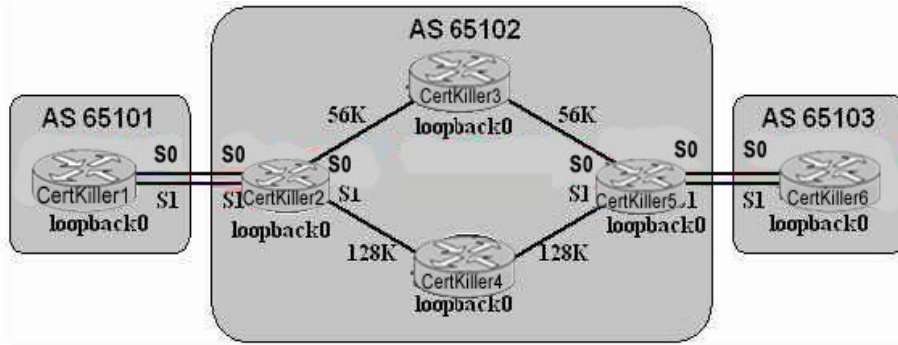
- A. 2.2.2.2
- B. 3.3.3.3
- C. 10.1.1.1
- D. 10.1.1.2

Answer: C

---

**QUESTION 7:**

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

## QUESTION 8:

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 9:**

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 10:**

BGP table

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

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

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 11:**

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 12:**

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  
!  
ip route 192.168.0.0 255.255.0.0 192.168.1.2
- 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  
!  
ip route 192.0.0.0 255.0.0.0 192.168.1.2
- D. 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  
!  
ip route 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 13:**

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

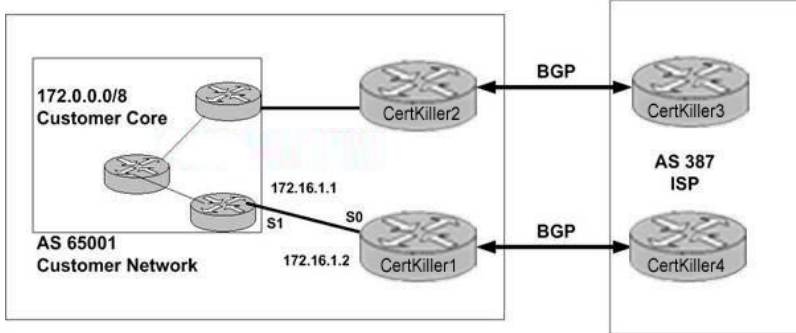
Wrong Answer

1. A is missing Next-hop-self command
2. B This is not done at the EBG level
3. D update-source-self invalid parameter update-source address is the proper neighbor parameter

---

**QUESTION 14:**

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. \_\_\_\_\_

!

router bgp 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 15:**

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 16:

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 17:**

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  
staticroute
```

```
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 bg 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 18:**

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 19:**

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 20:**

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 66.0.0.1
BGP state = Established, table version = 1, up for 00:14:52
Last lead 00:00:52, hold 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 00:16:02 due to User 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 21:**

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 notifications, 0 in queue
Sent 0 messages, 0 notifications, 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 22:**

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 23:**

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

Explanation:

BGP doesn't use reverse mask like access lists. The correct command would be network 10.1.0.0 mask 255.255.0.0

---

**QUESTION 24:**

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

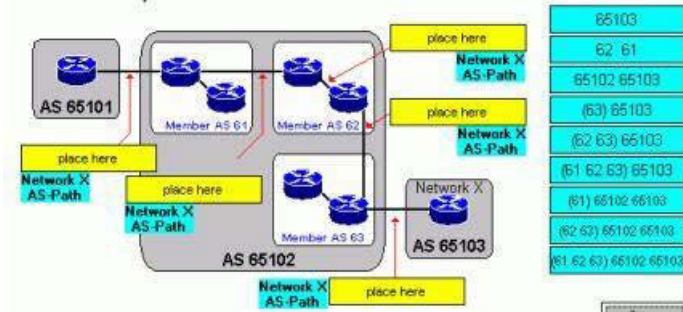
- 1. B Lowest RID
  - 2. D default is one
- 

**QUESTION 25:**

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 26:

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 27:

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.)

```
router bgp 65111
neighbor 10.1.1.1 remote-as 65101
```

- A. neighbor 10.1.1.1 access-list 1 in
- access-list 1 permit 172.16.0.0/24
- B. neighbor 10.1.1.1 access-list 1 in

- access-list 1 permit 172.16.0.0 0.0.255.255
- C. neighbor 10.1.1.1 access-list 1 in
- access-list 1 permit 172.16.0.0 255.255.0.0
- D. neighbor 10.1.1.1 prefix-list test in
- ip prefix-list test seq 5 permit 172.16.0.0/17 le 24
- E. neighbor 10.1.1.1 prefix-list test in
- ip prefix-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 28:**

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 le 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 le 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 29:**

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 30:**

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 31:**

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 32:**

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 33:**

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

```
router bgp 65111
neighbor 10.1.1.1 remote-as 65112
neighbor 10.1.1.1 route-map test out
!
ip as-path access-list 1 permit _65222$
ip as-path access-list 2 permit.*
!
route-map test permit 10
match as-path 1
set metric 200
!
route-map test permit 20
match as-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 34:**

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 35:**

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 36:**

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 37:**

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 38:**

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 39:**

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 800
neighbor 10.2.2.2 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

Explanation:

The routes with better weight are always used first.

---

**QUESTION 40:**

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

Explanation:

The weight value is local to one router. It is not shared between routers. If one ISP connection terminates in one of the customer routers, and the other ISP connection terminate in another, the two customer routers must agree on which link to use. Using local-preference instead of wight can do this.

Reference: Configuring BGP on Cisco routers volume1 p.4-9

---

**QUESTION 41:**

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

---

**QUESTION 42:**

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

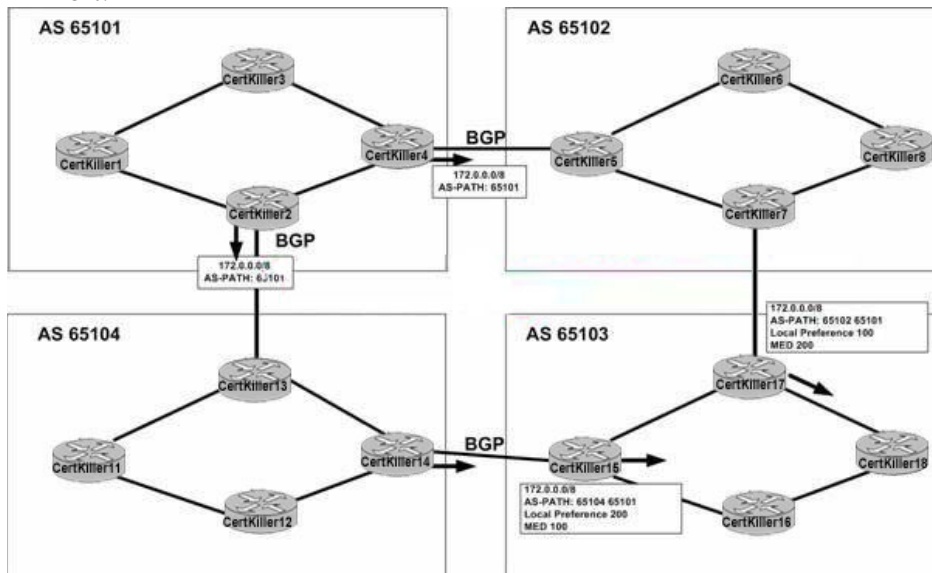
Explanation: Command requires the use of as path `_65327$` mean originating from 65237

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 43:**

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 44:**

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

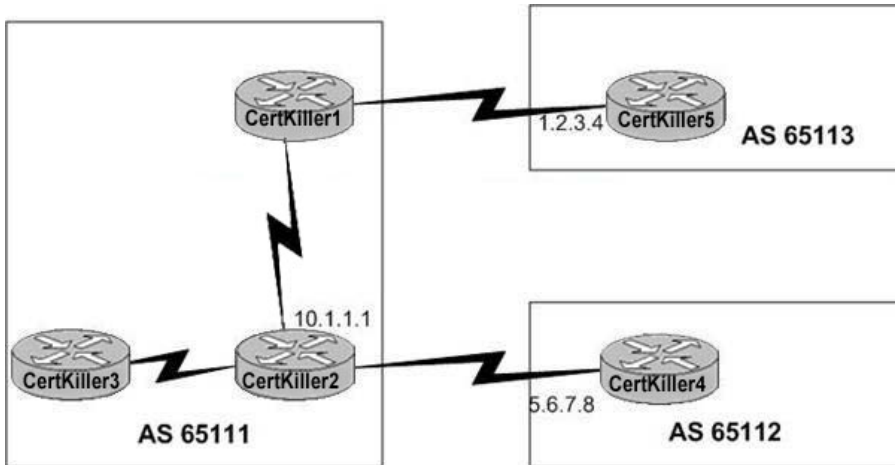
Wrong Answer

1. C Invalid command
2. D Invalid command

---

**QUESTION 45:**

Exhibit:



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

```
hostname Certkiller 1
!  
!output omitted  
!  
router bgp 65111  
neighbor 10.1.1.1 remote-as 65111  
neighbor 10.2.2.2 remote-as 65111  
neighbor 1.2.3.4 remote-as 65113  
neighbor 1.2.3.4 route-map L2M in  
!  
route-map L2M permit 10  
set local-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 46:**

Given the following:  
router bgp 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-map setlp permit 10
match community 1
set local-preference 150
!
route-map setlp permit 20
!
ip community-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.

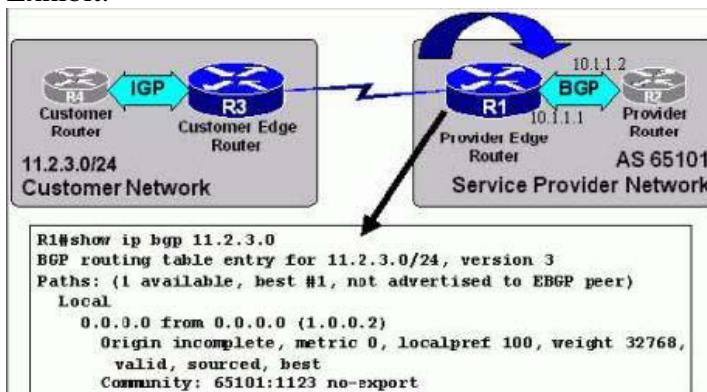
Answer: B

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 47:**

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

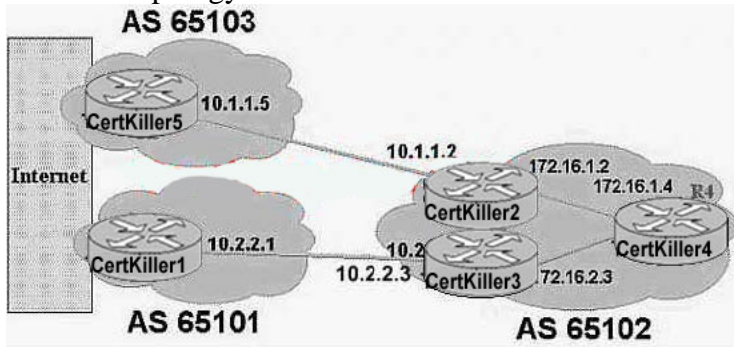
Wrong

- 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 48:**

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

Explanation:

Bgp default local-preference command can be used in BGP configuration mode to change that default value of local preference. The new default value only applies to locally originated routes and those received from external neighbors.

Setting a valued lower than the default 100 will result in the router preferring internal paths to external (normally a router would prefer external routes).

Setting a value higher than 100 will result in external paths being preferred to all internal paths (also those with a shorter AS path).

Reference: Configuring BGP on Cisco routers p.7-11

---

**QUESTION 49:**

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 50:**

DRAG DROP

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.



## 642-661

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 51:**

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

Explanation:

Really, MED is assigned using set metric command. The prefix 10.0.0/8 will match route-map 10 entry, will be assigned MED 100 and will not be further checked.

### QUESTION 52:

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 10.2.2.2 remote-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 setmed 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 53:**

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 54:**

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 55:**

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

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

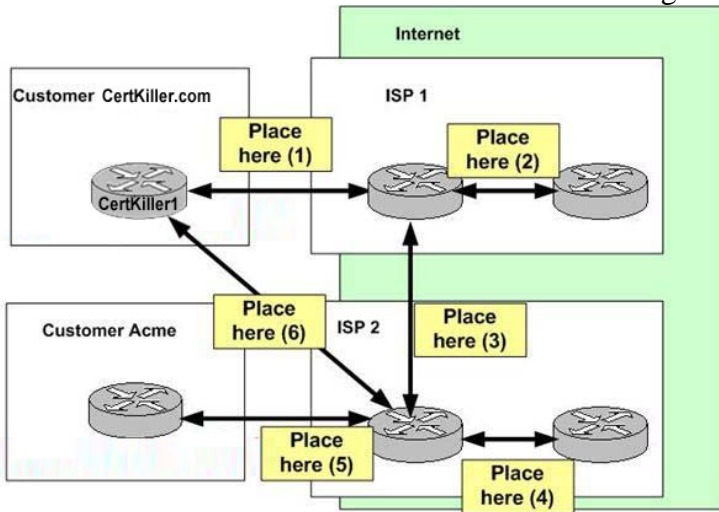
Answer: C

---

**QUESTION 56:**

DRAG DROP

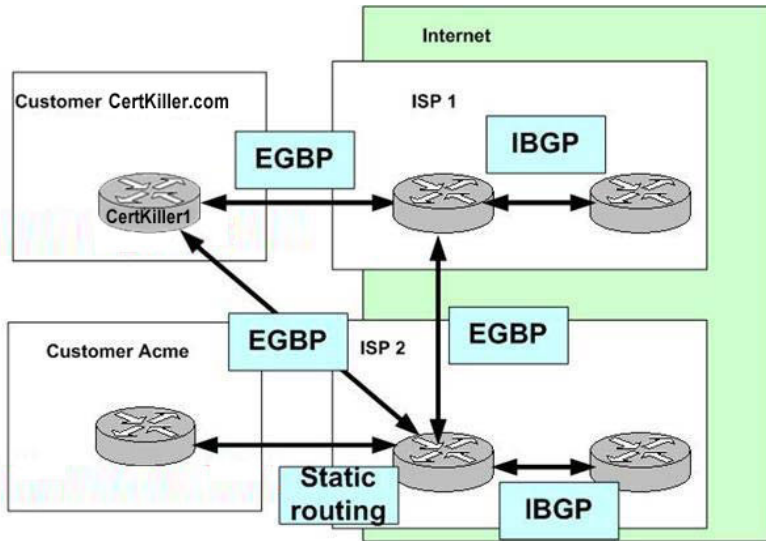
Identify the most appropriate method to connect the customers to the ISPs, to interconnect the two ISPs, and to interconnect the router within each ISP. Place the methods in their correct locations in the diagram.



Methods, Select from these

- EGBP
- IBGP
- Static routing

Answer:



Methods, Select from these

- EGBP
- IBGP
- Static routing

Explanation:

1. Box 1 EGBP Certkiller .com 1 to ISP1. This is a multi-homed Internet connection. Most appropriate
2. Box 2 IBGP Routers within ISP1
3. Box 3 EGBP 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 57:**

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

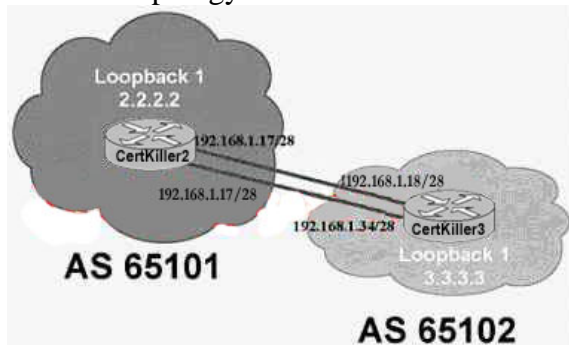
- A. The VRF will leak routes into the global routing table.
- B. The P router must run both IGP and BGP.
- C. The PE routers must perform Internet routing
- D. A limited set of Internet services will be implemented.
- E. A separate physical link or specific WAN encapsulation is required.

Answer: C, E

---

**QUESTION 58:**

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 59:**

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 60:**

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

Explanation:

Removal of a private AS-number from the AS-path is accomplished by using remove-private-as on the ISP's EBGP sessions to the rest of the Internet.

Reference: Configuring BGP on Cisco routers volume2 p.8-66

---

### QUESTION 61:

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:

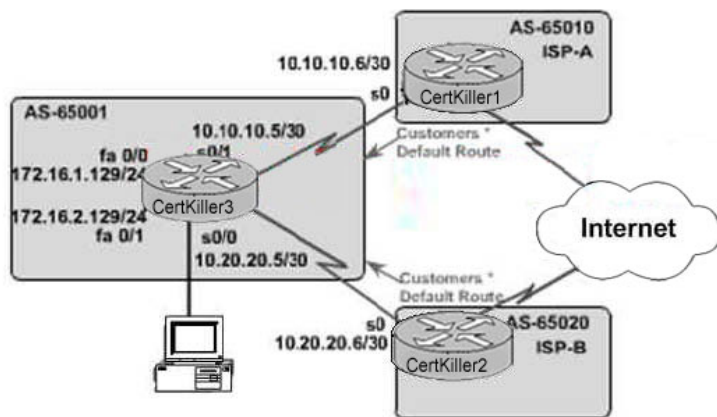
1. A Customers do not want to become a transit for ISP's
  2. E This a valid way to advertise routes to the ISP
- Wrong
1. B ISP essentials says most popular, does not say recommend
  2. C affects inbound traffic
  3. D local preference is for outbound traffic

---

## QUESTION 62:

### 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-patch ACL number 10 containing only one statement.

Start the simulation by clicking on the local host icon.

Answer:

To properly inject prefixes in a BGP routing table, it must exist in a IP routing table. network command and ip route with Null0 would announce a route to a BGP peers (ISP).

```
interface FastEthernet0/0
ip address 172.16.1.129 255.255.255.0
no shutdown
!
interface FastEthernet0/1
ip address 172.16.2.129 255.255.255.0
no shutdown
!
interface Serial0/0
ip address 10.20.20.5 255.255.255.252
no shutdown
!
interface Serial0/1
ip address 10.10.10.5 255.255.255.252
no shutdown
!
ip route 0.0.0.0 0.0.0.0 10.19.19.6
ip route 0.0.0.0 0.0.0.0 10.20.20.6
ip route 172.16.1.0 255.255.255.0 Null0
ip route 172.16.2.0 255.255.255.0 Null0
!
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
no synchronization
no auto-summary
!
ip as-path access-list 10 permit ^$
!
```

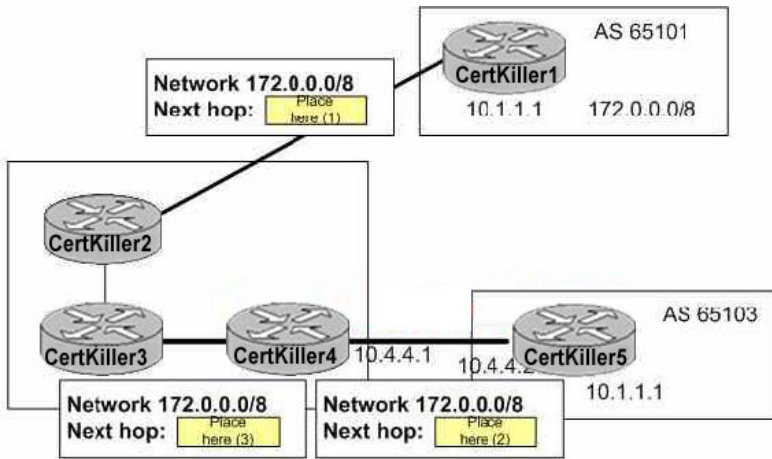
---

### QUESTION 63:

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?

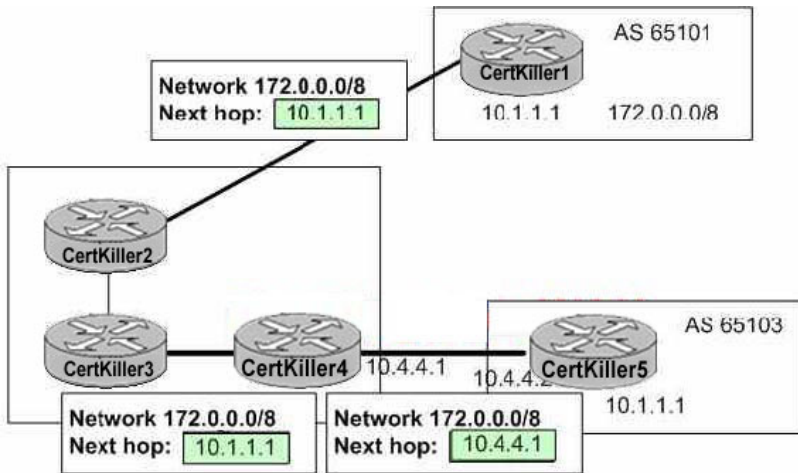
Place the correct next-hops in their locations in the diagram.



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 64:**

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

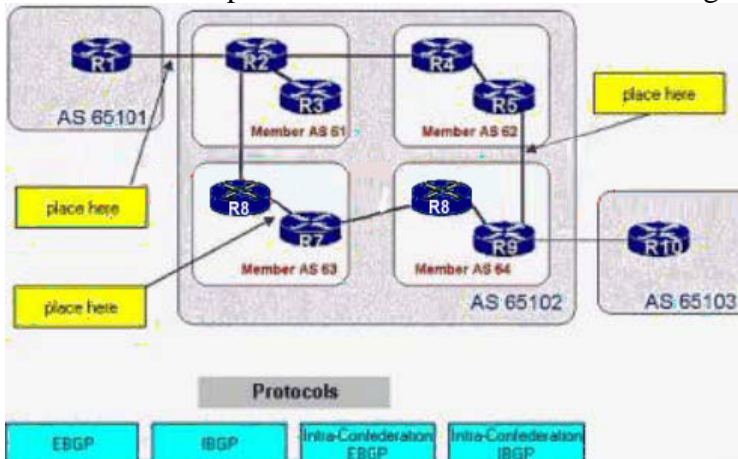
- 1. A Synchronization ensures that the routes of the EGP match the IGP.
- 2. B this has no bearing
- 3. C Generally this is insane
- 4. E You still need to be sure that you have a route back to the loopback interface

---

**QUESTION 65:**

DRAG DROP

Place the correct protocols in their locations on the diagram.



Answer:

Explanation:

From TOP Left going clockwise


- 1. Box 1 EIGRP R1-R2

- 2. Box 2 Intra confederation EBGP
- 3. Box 3 IBGP

---

**QUESTION 66:**

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 67:**

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 68:**

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. Cisco Express Forwarding (CEF) pre-builds complete IP forwarding table (called Forwarding Information Base or FIB) based on IP routing table. Once an entry is installed in the routing table, due to incoming routing information updates, the recursive lookup is done, and the outgoing interface and the actual physical next-hop are determined. MAC address resolution and MAC header generation are still traffic-driven and stored in the cache.

Reference: Configuring BGP on Cisco routers p.4-32

---

### **QUESTION 69:**

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

Explanation:

Core routers need to run BGP and IGP and:

- 1) BGP shall carry all external routes.
- 2) IGP shall only propagate BGP next-hops and other core subnets
- 3) All customer routes shall also be carried in BGP
- 4) Reduces IGP topology database
- 5) Removes customer-caused route flaps from IGP: IGP becomes more stable.

Reference: Configuring BGP on Cisco routers p.4-34

---

### **QUESTION 70:**

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

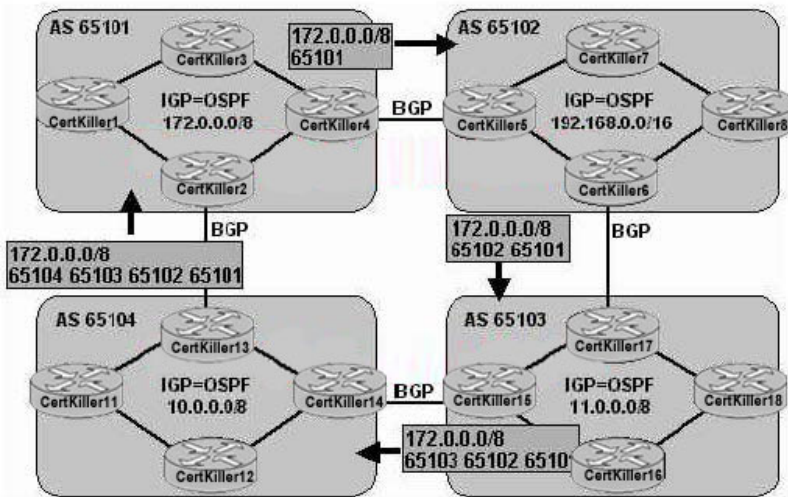
Explanation:

Synchronization was used in older transit autonomous system design that relied on BGP route redistribution into IGP. Modern autonomous system designs cannot rely on this any more as the number of routes carried in the Internet exceeds the scalability range of any known IGP. Redistribution into IGP is thus no longer applicable and the synchronization feature has to be disabled in order for you transit autonomous system to work.

Reference: Configuring BGP on Cisco Routers p.4-45

**QUESTION 71:**

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 72:**

Exhibit:



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
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
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
no shut

```



```
!  
interface e1  
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  
neighbor 2.2.2.2 update-source loopback0  
!  
router eigrp 101  
network 10.0.0.0  
!
```


- 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 73:**

Exhibit:



Given the following configuration for Certkiller 2:

```
! output omitted  
!  
hostname Certkiller 2  
!  
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

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

---

### QUESTION 74:

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 75:**

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 redistributed 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 76:**

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 77:**

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 78:**

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

Explanation:

If a client has IBGP sessions to route reflectors that belong to different clusters, the BGP update from the client will be forwarded into the full mesh with different cluster-IDs in the cluster-list attribute. When the BGP update enters the mesh it will reach the other route reflector, which will, unnecessarily, accept the route as valid and forward it into its cluster. This, in turn, causes unnecessarily duplication of updates to the clients.

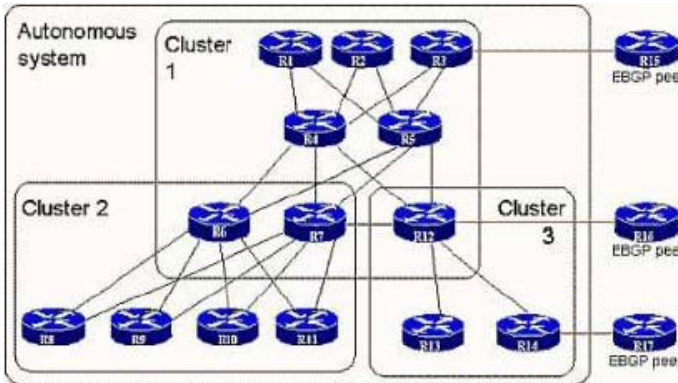
If a client has IBGP sessions to other clients in the same cluster, those clients will receive unnecessarily duplications of updates.

Reference: Reference: Configuring BGP on Cisco routers volume1 p.5-19

---

#### **QUESTION 79:**

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 80:

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 81:**

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 82:**

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 83:

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 84:

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 85:**

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 86:**

Network topology exhibit



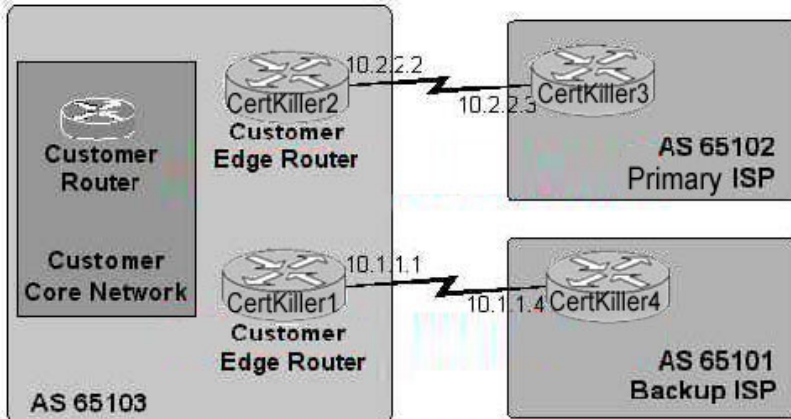


Exhibit #2

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

```
route-map rest 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

Explanation:

Manual AS-path manipulation is configured using a route-map with the set as-path prepend statement, which prepends the specified AS-numbers to the routes matched with the match condition. The route-map is then applied to the outgoing updates to an EBGP neighbor. The prepending is made first, then the route is subject to the normal AS-path modification procedures when it is sent over an EBGP session.

The route-map can also be used to select only a subset of routes that should have their AS-path manually manipulated. In the appropriate route-map permit statement, the set as-path prepend command is used.

### QUESTION 87:

Given the following:

```
router bgp 65123
 neighbor customers peer-group
 neighbor 10.1.1.2 remote-as 65213
```

```
neighbor 10.1.1.2 peer-group customers
neighbor 10.2.2.2 remote-as 65314
neighbor 10.2.2.2 peer-group customers
neighbor 10.2.2.2 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 88:**

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 89:**

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-mapcommand.
- E. Using the neighbor ip-address soft-configuration command.

Answer: A, B

Explanation:

1. A bgp dampening [half-life reuse supress max-supress-time][route-map map]
2. B bgp dampening [half-life reuse supress max-supress-time

Wrong

1. C used for max prefixes
2. D not a command
3. E control refresh

---

**QUESTION 90:**

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 91:**

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 then the half-time limit.
- E. When the penalty associated with a route drops below half the reuse limit.

Answer: E

Explanation:

A router stops tracking penalty points once they are below half of the reuse limit. After enabling RFD, the router never removes a route from the BGP table. A route that has been withdrawn by a BGP neighbor can still be seen in the BGP table marked with "h" (history state. The penalty is always applied to a path and not a prefix. If one of the paths is flapping, it does not mean that the destination is flapping.

Reference: Configuring BGP on Cisco routers p. 7-120

---

**QUESTION 92:**

DRAG DROP

Drag the VRF-aware protocols to the method used that redistributes routes from MP-BGP into the protocol VRF instance.

EBGP	EIGRP	IGRP	IS-IS
OSPF	RIPv1	RIPv2	CHAP

Automatic redistribution

Place here

Manually configured redistribution

Place here    Place here    Place here

Answer:

Automatic redistribution

EBGP


Manually configured redistribution

EIGRP	OSPF	RIPv2
-------	------	-------

---

**QUESTION 93:**

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
!  
router bgp 65101  
neighbor 10.1.1.2 remote-as 65102  
neighbor 10.2.2.2 remote-as 65102
```

- A. variance 2
- B. maximum-paths 2
- C. neighbor 10.1.1.2 weight 100  
neighbor 10.2.2.2 weight 100
- D. neighbor 10.1.1.2 ebgp-multihop 2  
neighbor 10.2.2.2 ebgp-multihop 2
- E. neighbor 10.1.1.2 local-preference 100  
neighbor 10.2.2.2 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 94:**

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

Explanation:

The BGP route is installed in the routing table only if the IP address in the next-hop attribute is reachable according to the information already in the routing table. The BGP route is installed with a reference to that next-hop address. So the network will be reachable via an IP address, which may or may not be directly connected. Since there is no clear reference to a physical interface, the BGP route is installed in the IP routing table without any information about outgoing interface.

Reference: Configuring BGP on Cisco routers. P.4-30

---

**QUESTION 95:**

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

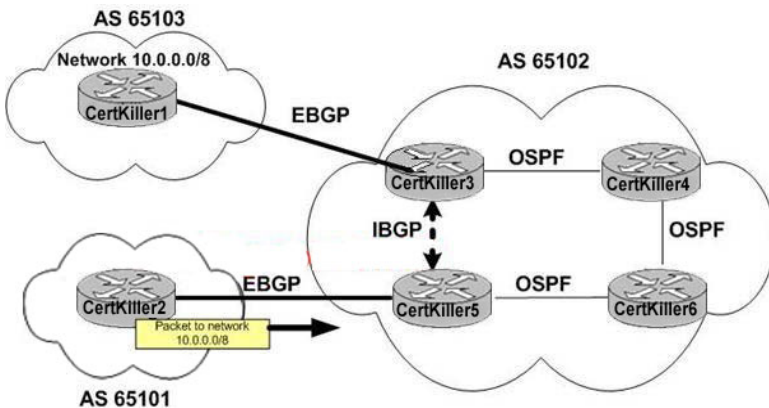
- A. half-time
- B. reuse-limit
- C. max-suppress
- D. suppress-limit
- E. per-flap-penalty

Answer: E

---

**QUESTION 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

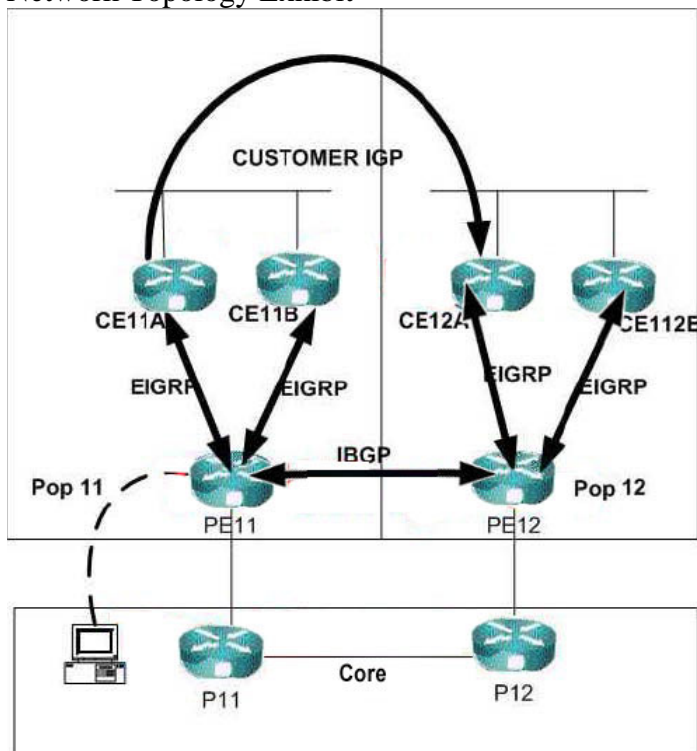
Explanation: Even with synchronization off testing 6 does not know how to route the packet because OSPF does not know about the route that BGP knows about. Synchronization being off assumes that there is a full mesh of BGP routers.

---

**QUESTION 97:**

**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 **\*\*MISSING\*\***.

The RD is 10:10.

The route target for import is 17:10.

The route target for export is 17:10.

Configure EIGRP between CE11A and PE11 using AS number 10.

[642-661](#)

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
description customer vrf
rd 10:10
route-target export 17:10
route-target import 17:10
!
interface FastEthernet0/0
description customer
ip vrf forwarding customer
ip address 10.1.11.1 255.255.255.0
no shutdown
!
router eigrp 10
no auto-summary
!
address-family ipv4 vrf customer
redistribute bgp
default-metric 10000 100 255 1 1500
network 10.1.11.0
network 10.1.12.0
no auto-summary
exit-address-family
!
router bgp 17
neighbor 10.1.12.1 remote-as 17
no auto-summary
!
address-family ipv4
no auto-summary
no synchronization
exit-address-family
!
address-family vpnv4
neighbor 10.1.12.1 activate
neighbor 10.1.12.1 send-community extended
exit-address-family
```



```
!  
address-family ipv4 vrf customer  
redistribute eigrp 10  
default-metric 1  
no auto-summary  
no synchronization  
exit-address-family  
!
```

---

**QUESTION 98:**

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

Explanation:

Route Refresh Capability:

- 1) Does not require the BGP session to close
  - 2) Does not require any additional memory
  - 3) No down-time
  - 4) Two versions:
    - Cisco proprietary (Capability Code 128; for all protocols)
    - Standard (RFC 2918; Capability Code 2; per-address-family)
- Reference: Configuring BGP on Cisco Routers p.0-29
- 

**QUESTION 99:**

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

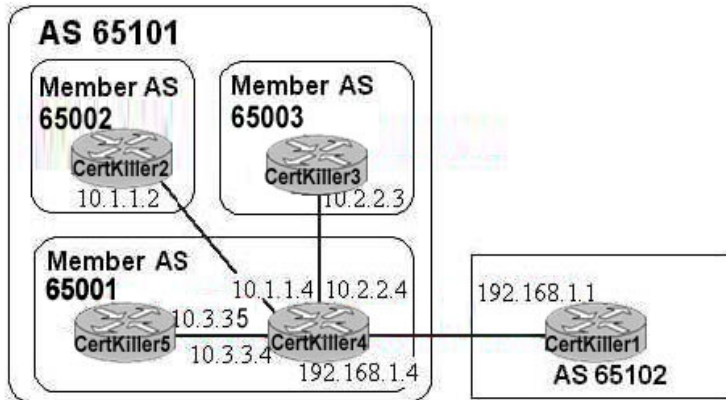
Answer: A, B, D

---

**QUESTION 100:**

Network topology exhibit

## Confederations



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

- A. router bgp 65001  
 bgp confederation indentifier 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 indentifier 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

bgpconfederation 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 101:**

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 102:**

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 103:**

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

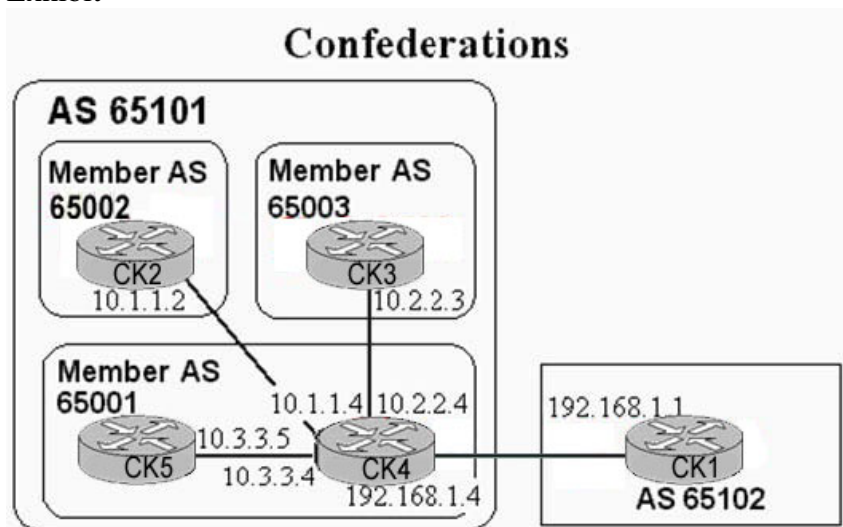
- A. All routers within the BGP confederation must support BGP confederations.
- B. The member AS numbers used with the confederation are visible from outside the confederation.
- C. An Intra-Confederation EBGP session behaves like an EBGP session when propagating routing updates.
- D. Intra-Confederation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

Answer: A, D

---

**QUESTION 104:**

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

Explanation:

Bgp confederation peers command defines all the other autonomous systems in the confederation.

---

**QUESTION 105:**

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 106:**

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 107:**

**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 108:**

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 109:**

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: E

---

**QUESTION 110:**

```
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 111:**

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 112:**

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 113:**

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 114:**

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 115:**

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

- A. The core subnets and loopback 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 116:**

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 117:**

```
PIR1#show route-map test
route-map test . permit, sequence 10
  Match clauses:
    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 118:**

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 119:**

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 120:**

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 121:**

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 122:**

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 123:**

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 124:**

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 125:**

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 126:**

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 Preference	3
Origin Code	4
As Path	5

Answer:

Weight
As Path
Local Preference
Origin Code
MED

---

**QUESTION 127:**

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 128:**

```
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 129:**

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 130:**

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 131:**

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 132:**

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 133:**

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 134:**

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 135:**

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 136:**

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 137:**

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 138:**

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?

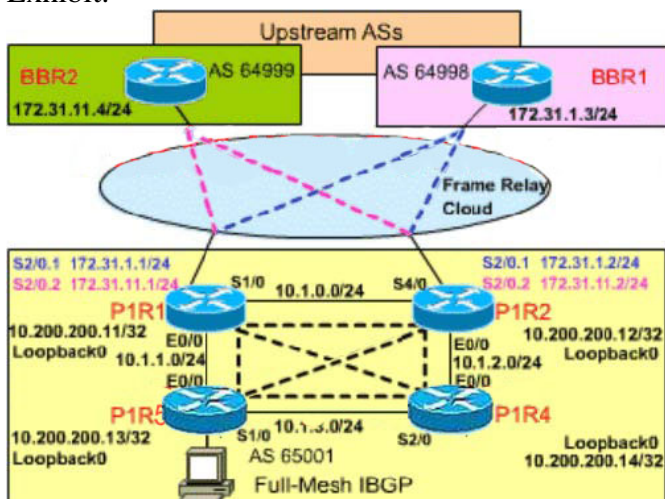
- A. CK1 only
- B. CK5 and CK6
- C. CK5 , CK6 and CK1
- D. CK4 , CK5 and CK6
- E. CK4 , CK5 , CK6 , and CK1
- F. To no other router

Answer: C

### QUESTION 139:

#### 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:

Explanation:

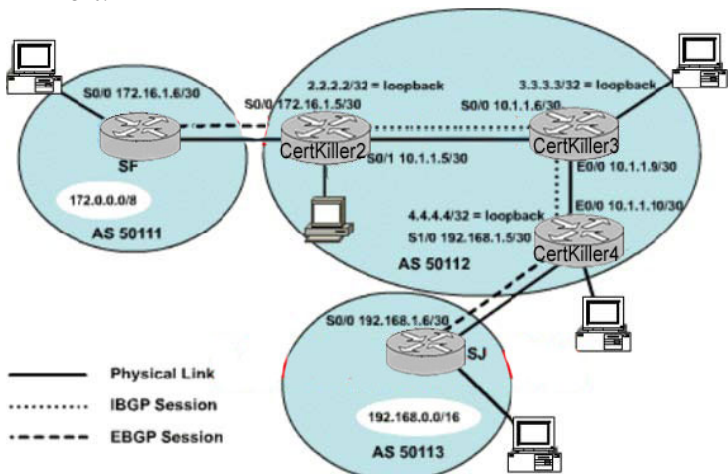
1. no synchronization
2. neighbor 10.200.200.12 weight 100  
clear ip bgp \* soft in

---

**QUESTION 140:**

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:

- ```
router bgp 50113
neighbor 4.4.4.4 remote-as 50112
neighbor 2.2.2.2 remote-as 50112
```

```
neighbor 4.4.4.4 route-reflector-client  
neighbor 2.2.2.2 route-reflector-client
```

---

**QUESTION 141:**

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 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 142:**

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 143:**

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 144:**

Which prefix-list will permit the 172.16.1.32 255.255.255.240 prefix?

- A. ip prefix-list Certkiller permit 172.16.1.0/28
- B. ip prefix-list Certkiller permit 172.0.0.0/8 le 30
- C. ip prefix-list Certkiller permit 172.16.0.0/16 eq 28
- D. ip prefix-list Certkiller permit 172.16.1.0/8 ge 29

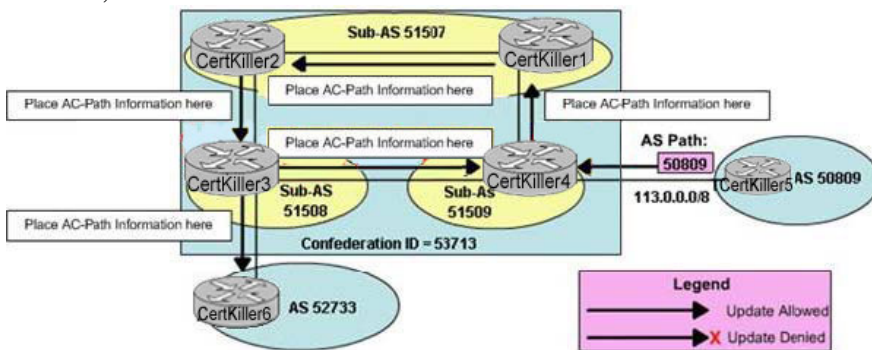
Answer: B

---

**QUESTION 145:**

**SIMULATION**

Exhibit, Simulation



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 146:**

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 147:**

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 148:**

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 149:**

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 150:**

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 151:**

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 152:**

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 153:**

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 154:**

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 155:**

**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**

|                                                                       |
|-----------------------------------------------------------------------|
| Display the BGP Table.                                                |
| Display BGP attributes attached to a prefix like the community value. |
| Verify the BGP neighbors state.                                       |
| Display the BGP routes in the routing table.                          |
| Verify AS-path access-list configuration                              |

**Show commands, place here**

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

Show commands, Select from these

---

**QUESTION 156:**

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 157:**

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 158:**

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 159:**

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 160:**

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 161:**

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 162:**

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 163:**

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 164:**

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 192.168.13.0/24 [20/0] via 172.31.11.4, 00:52:57
B 192.168.14.0/24 [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 165:**

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_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 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 166:**

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 167:**

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 168:**

**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 169:**

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 170:**

**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

Used to influence the outbound traffic flow

- Setting the weight
- Setting the local preference

Setting the admin distance

Using recursive route lookup

Used to influence the inbound traffic flow

- Setting the MED
- Using AS-path prepending

---

**QUESTION 171:**

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 172:**

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 173:**

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 174:**

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.

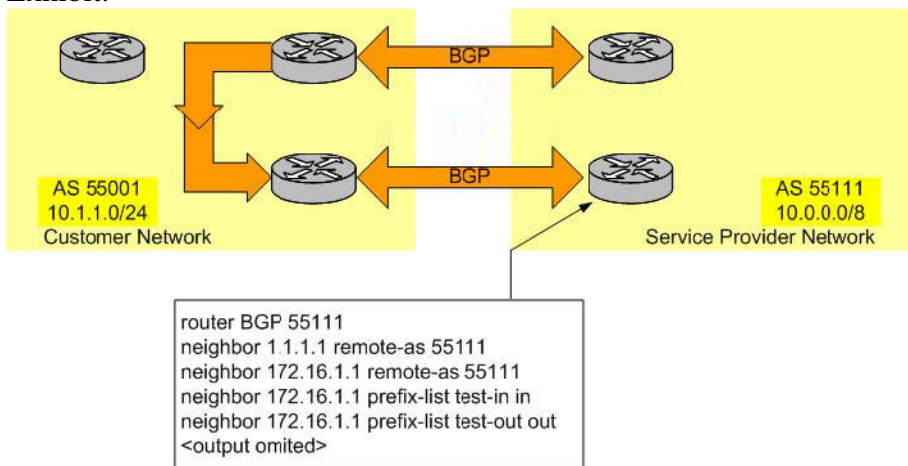
Answer: D

---

**QUESTION 175:**

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 176:**

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 177:**

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 178:**

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 179:**

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 180:**

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.

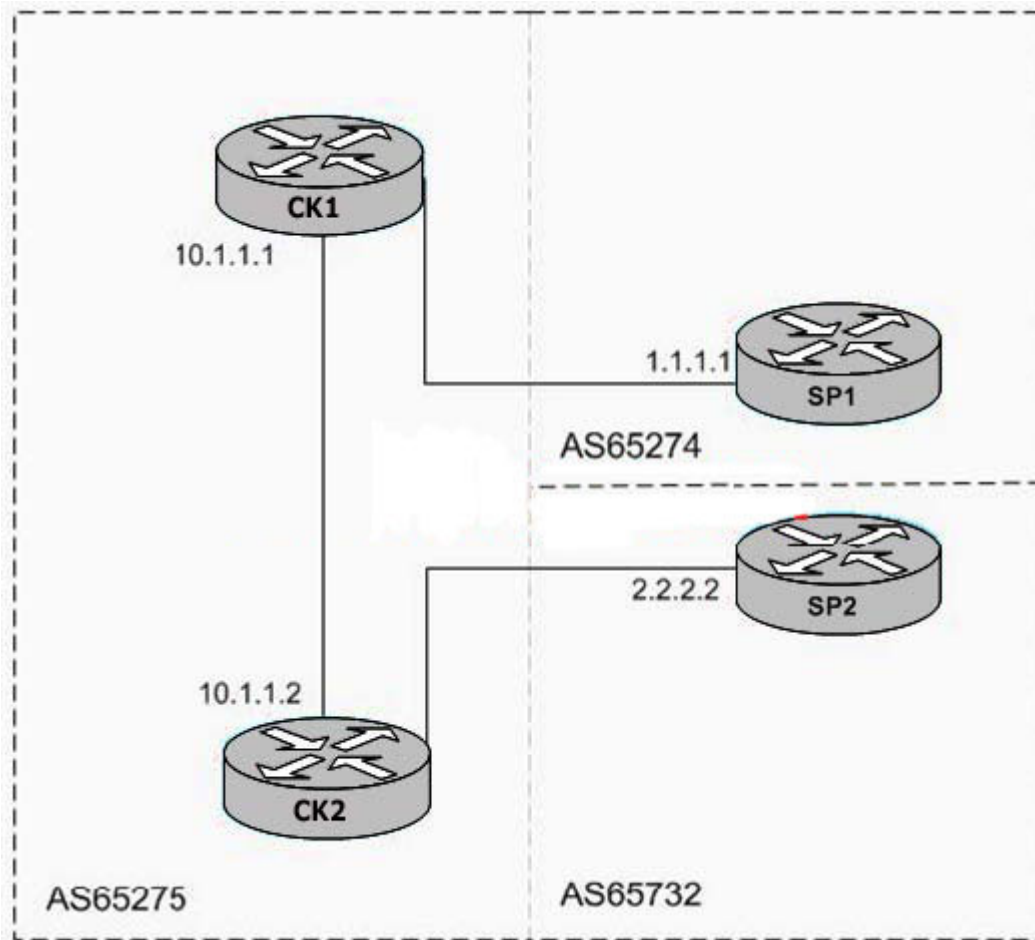
Answer: B, E

---

**QUESTION 181:**

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 182:**

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 183:**

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 184:**

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 185:**

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 186:**

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 187:**

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 188:**

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 189:**

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 190:**

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 191:**

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 192:**

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 193:**

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 194:**

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 195:**

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

Confederations

All routers within the AS must support it

Reduce IBGP sessions by using special EBGP sessions

Override the BGP synchronization rule

Start the BGP routing process using the cluster ID

Limit the prefix being received from a BGP peer

Route Refelctors

Modify the BGP split horizon rule

Can use a hierarchical approach

---

**QUESTION 196:**

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 197:**

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 198:**

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

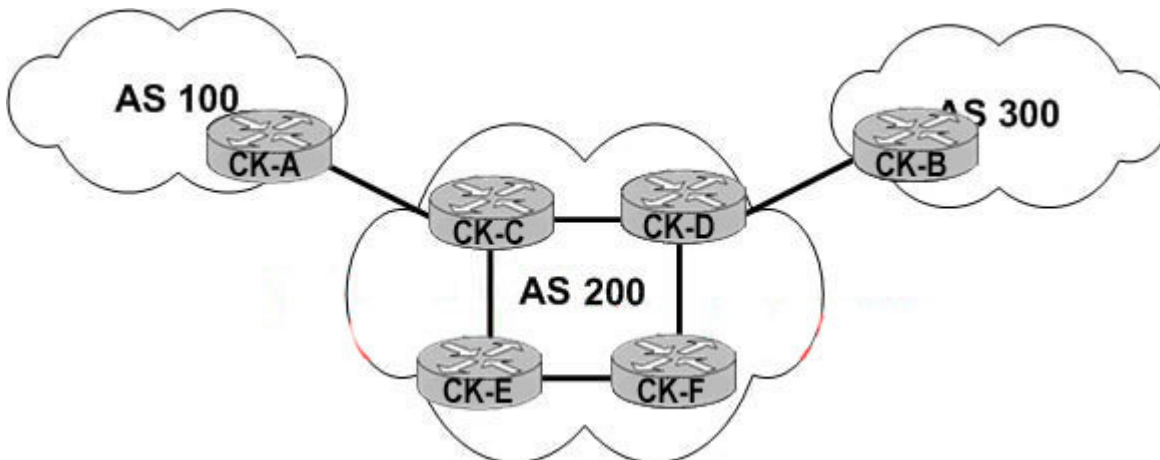
Answer: E

---

**QUESTION 199:**

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 200:**

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 201:**

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 202:**

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 203:**

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 autonomous system.     |  |

Place here

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

Answer:

Select from these

|                                                                                               |  |
|-----------------------------------------------------------------------------------------------|--|
| 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. |  |

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

Explanation:

question:

bgp deterministic med

answer:

Enforce the comparison of the MED between all paths received from the same autonomous system

explanation:

The bgp deterministic-med command can be configured to enforce deterministic comparison of the

MED value between all paths received from within the same autonomous system.

question:

bgp always-compare-med

answer:

Allow the comparison of MED for path from neighbors in different autonomous system.

explanation:

The MED is one of the parameters that is considered when selecting the best path among many alternative paths. The path with a lower MED is preferred over a path with a higher MED. During the

best-path selection process, MED comparison is done only among paths from the same autonomous

system. The bgp always-compare-med command is used to change this behavior by enforcing MED

comparison between all paths, regardless of the autonomous system from which the paths are received.

question:

bgp bestpath med-confed

answer:

Enable MED comparison among paths learned from confederation peers

explanation:

The MED comparison between confederation peers occurs only if no external autonomous systems are

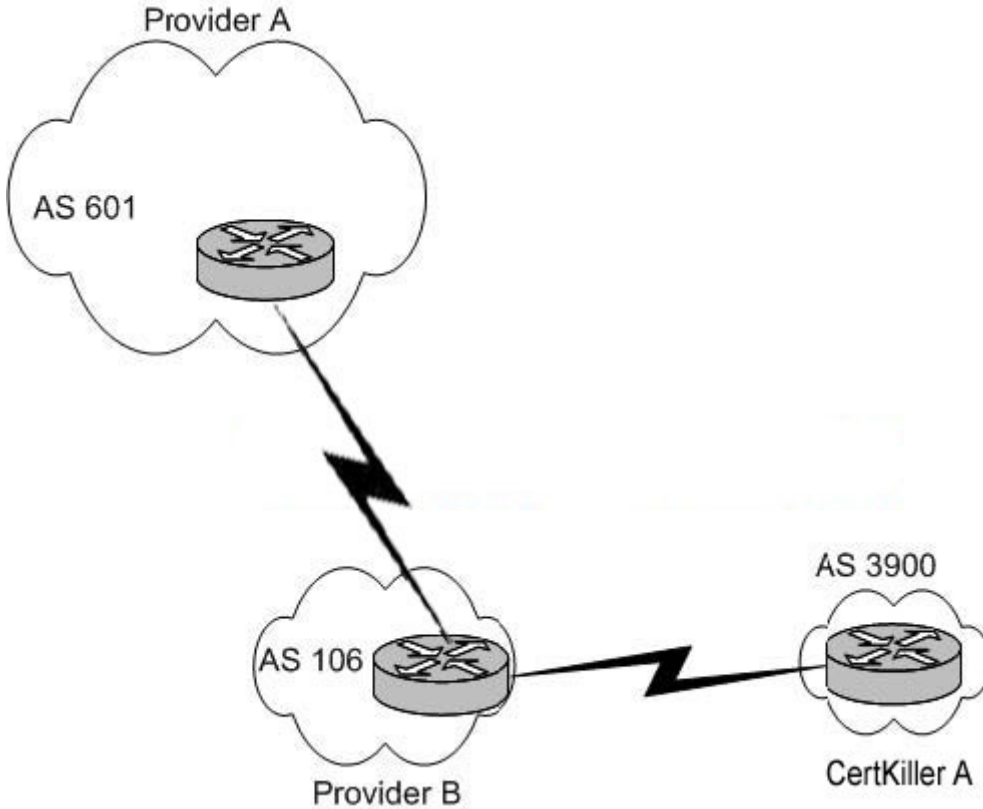
in the path (an external autonomous system is an autonomous system that is not within the confederation). If an external autonomous system is in the path, then the external MED is passed transparently through the confederation, and the comparison does not occur.

---

**QUESTION 204:**

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 205:**

DRAG DROP

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

642-661

Select from these

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.

A group of redundant route reflectors and their clients.

A BGP attribute configured on the route reflector and each route reflector cluster must have a unique value.

Occurs when the same packets are received by the router that sent them.

The route reflector set this attribute to the router ID of the peer from which it received the route.

Place here

Cluster ID

Route reflector cluster

Originator ID

Routing loop

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 206:**

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 207:**

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  
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 208:**

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 209:**

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:



**! CK1**

```
interface Loopback0
ip address 20.20.20.20 255.255.255.255
|
interface Ethernet0/0
ip address 172.16.2.1 255.255.255.0
|
interface Serial0/0
ip address 172.16.1.1 255.255.255.252
|
router bgp 55001
no synchronization
neighbor 172.16.1.2 remote-as 55002
```

**! CK2**

```
interface Loopback0
ip address 10.10.10.10 255.255.255.255
|
interface Ethernet0/0
ip address 172.16.3.1 255.255.255.0
|
interface Serial0/0
ip address 172.16.1.2 255.255.255.252
|
router bgp 55002
no synchronization
neighbor 172.16.1.1 remote-as 55001
```

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


Answer: A

---

**QUESTION 210:**

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 211:

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 212:

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 213:**

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

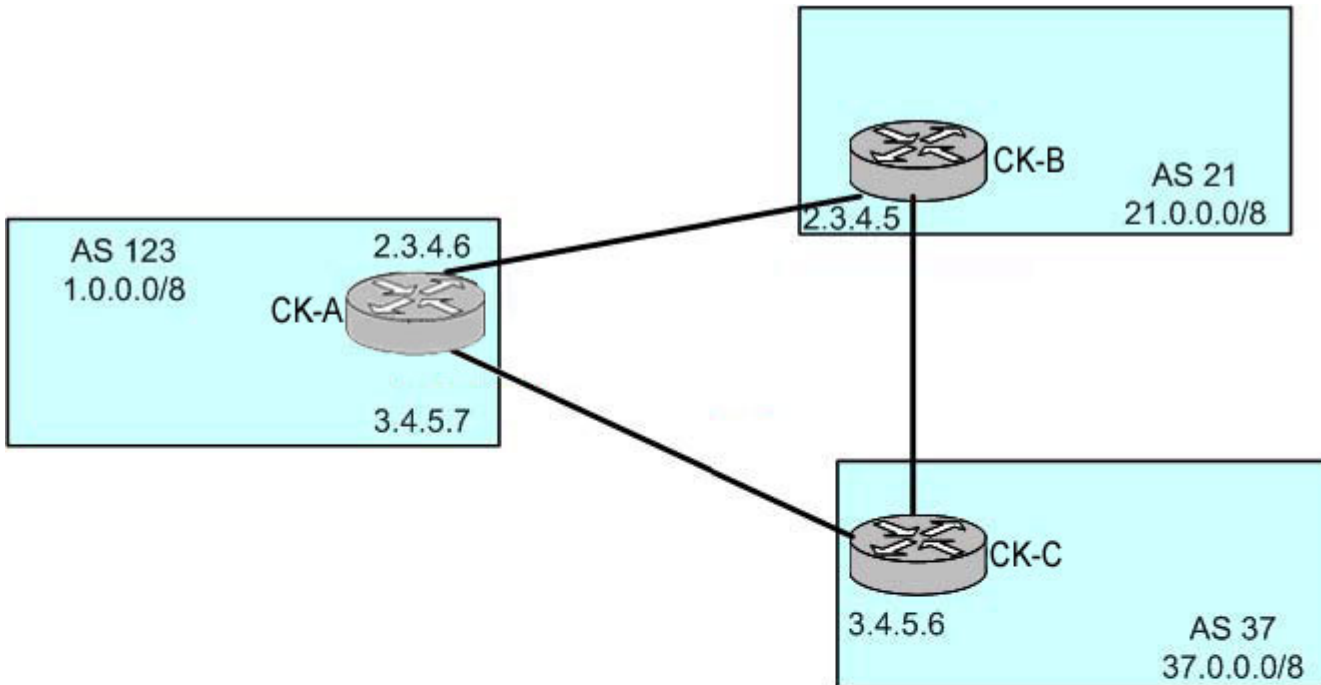
Answer: D

---

**QUESTION 214:**

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 215:

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 216:**

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 217:**

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.

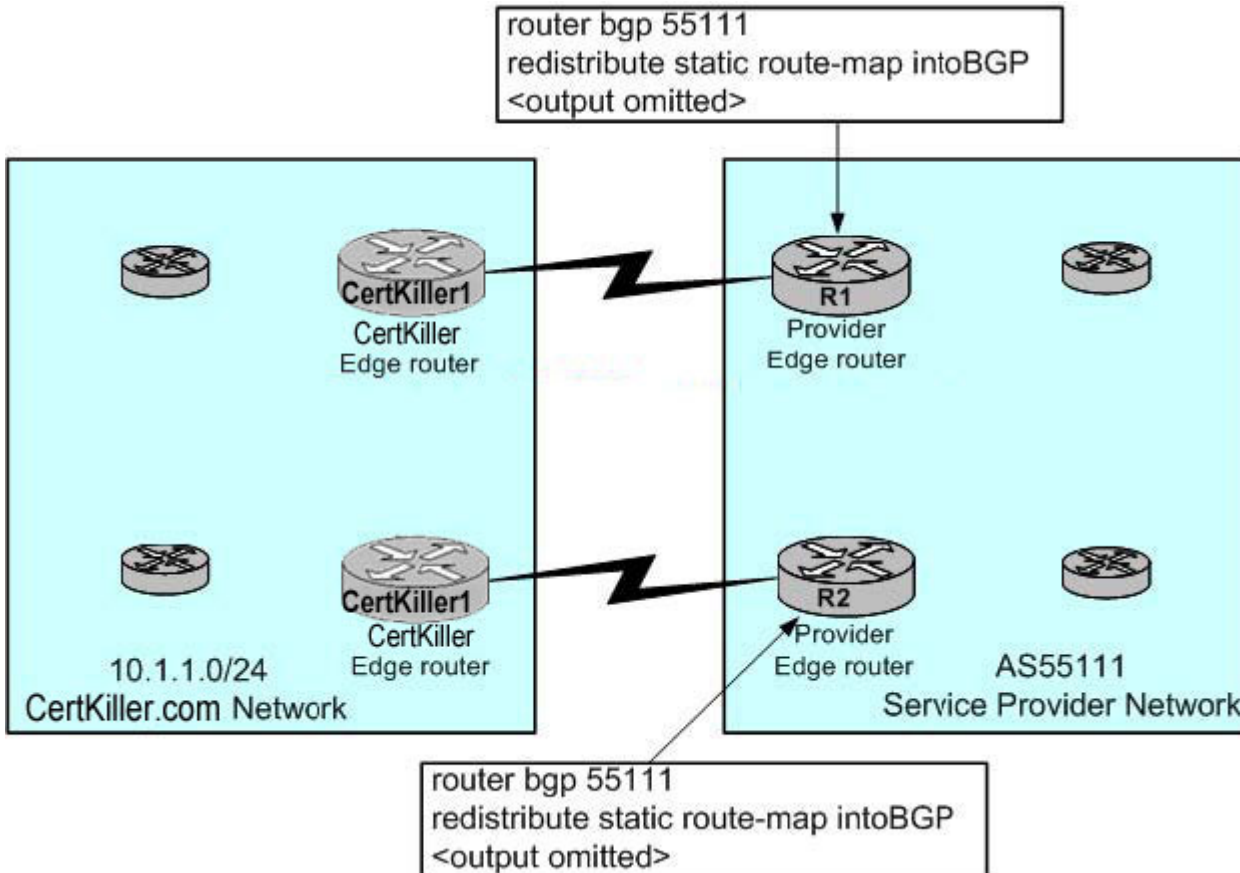
Answer: C

---

**QUESTION 218:**

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 219:**

**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 na external AS
- Propagates the external routes within the AS

| EBGP |
|------|
|      |
|      |

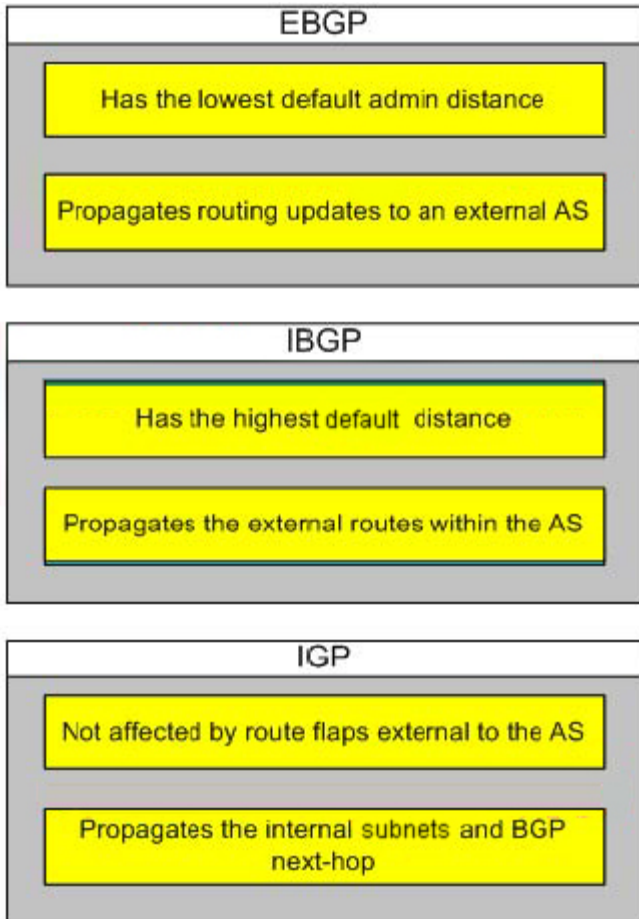
| IBGP |
|------|
|      |
|      |

| IGP |
|-----|
|     |
|     |

Answer:





Explanation:

EBP

1) Has lowest default admin distance:

Explanation: The Default admin distances are

EBGP 20

IGP 100

IBGP 200

See: [http://www.cisco.com/warp/public/105/admin\\_distance.pdf](http://www.cisco.com/warp/public/105/admin_distance.pdf)

2) Propagates routing updates to external neighbours

Explanation: eBGP is used to send routing updates to external neighbours

-----  
IBGP

1) Has the highest default admin distance:

Explanation: The Default admin distances are

EBGP 20

IGP 100

IBGP 200

See: [http://www.cisco.com/warp/public/105/admin\\_distance.pdf](http://www.cisco.com/warp/public/105/admin_distance.pdf)

2) Propagates the external routes within an as

Explanation: iBGP is used to propagate routes learned from external neighbour via eBGP to other iBGP neighbours.

IGP

1) Not affected by route flaps external to the AS.

Explanation: IGP is used for the interior routes, BGP handles the external routes.

2) Propagates the internal subnets and BGP next hops.

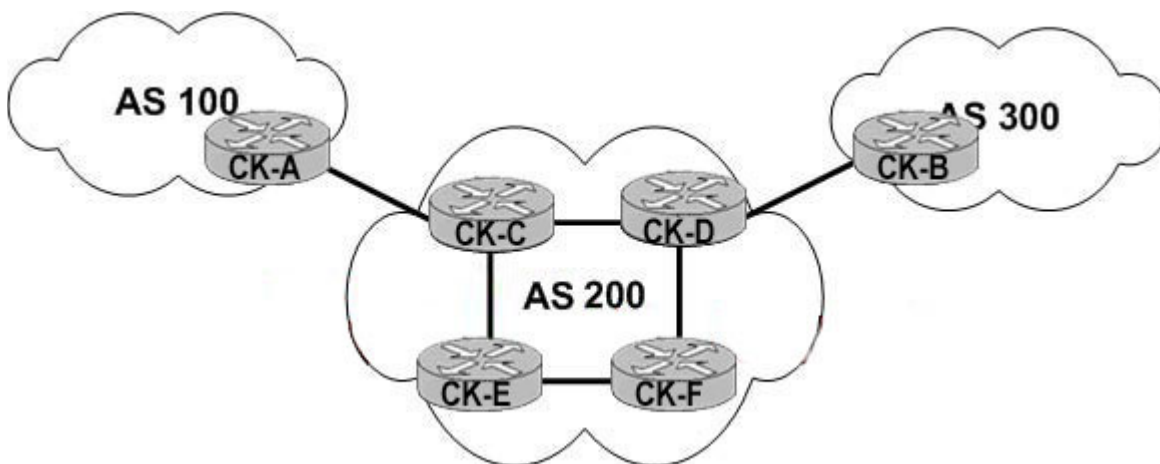
Explanation: IGP is used for the interior routes / internal subnets. BGP handles the external routes. The IGP propagates the next hops for BGP so IBGP neighbours can establish.

---

**QUESTION 220:**

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 221:**

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 222:**

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

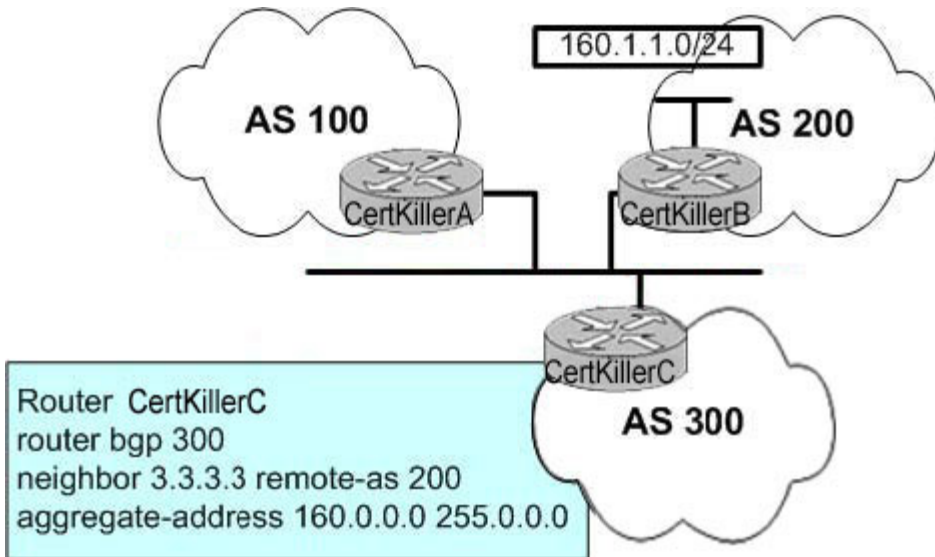
- A. BGP admin distance does not influence the path selection algorithm, but does influence if BGP-learned routes are placed in the IP routing table.
- B. External BGP routes have a default admin distance of 5.
- C. External BGP routes have a default admin distance of 20.
- D. Internal BGP routes have 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 223:**

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 224:**

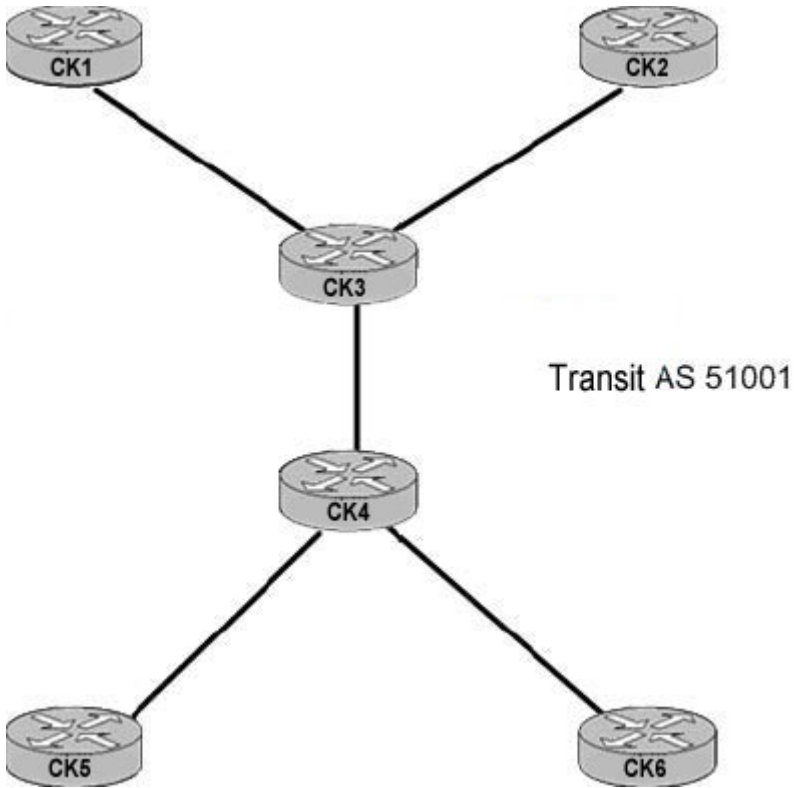
Which one of these statements regarding intraconfederation EBGP sessions is correct?

- 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 225:**

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 226:**

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)

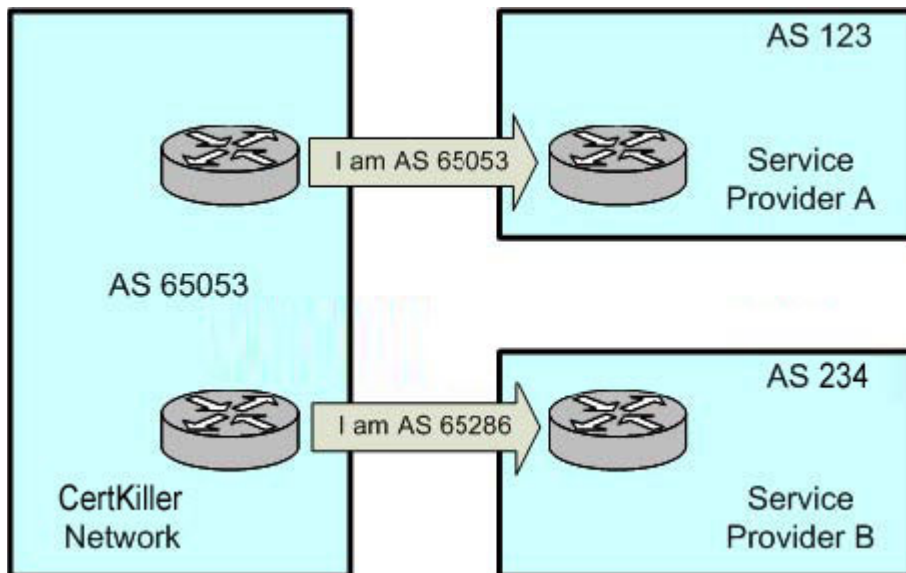
Answer: B

---

**QUESTION 227:**

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 228:**

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 - 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, 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 subnitted, 3 subnets
C 192.168.1.12 is directly connected, Serial2/0.2
C 192.168.1.0 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      | 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 229:**

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 230:**

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       0.0.0.0            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 231:**

Based on this configuration, which two peering router neighbor statements are correct? (Choose two.)

```
router bgp 50001
neighbor 192.168.1.1 remote-as 50001
neighbor 10.1.1.1 remote-as 50002
neighbor 10.1.1.1 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 and neighbor 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 232:**

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