


642-885 - SPADVROUTE Deploying Cisco Service Provider Advanced Routing (SPADVROUTE)

<https://www.certleader.com/642-885-dumps.html>



1. Referring to the topology diagram shown in the exhibit,

Which three statements are correct regarding the BGP routing updates? (Choose three.)

- A. The EBGP routing updates received by R1 from R5 will be propagated to the R2, R4, and R7 routers
- B. The EBGP routing updates received by R3 from R6 will be propagated to the R2 and R4 routers
- C. The EBGP routing updates received by R1 from R5 will be propagated to the R2 and R4 routers
- D. The IBGP routing updates received by R3 from R2 will be propagated to the R6 router
- E. The IBGP routing updates received by R2 from R1 will be propagated to the R3 router
- F. The IBGP routing updates received by R1 from R4 will be propagated to the R5, R7, and R2 routers

Answer: A,B,D

2. When a BGP route reflector receives an IBGP update from a non-client IBGP peer, the route reflector will then forward the IBGP updates to which other router(s)?

- A. To the other clients only
- B. To the EBGP peers only
- C. To the EBGP peers and other clients only
- D. To the EBGP peers and other clients and non-clients

Answer: C

3. Which two BGP mechanisms are used to prevent routing loops when using a design with redundant route reflectors? (Choose two.)

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

Answer: A,C

Explanation:

As the iBGP learned routes are reflected, routing information may loop.

The route reflector model has the following mechanisms to avoid routing loops:

- Originator ID is an optional, nontransitive BGP attribute. It is a 4-byte attributed created by a route reflector.

The attribute carries the router ID of the originator of the route in the local autonomous system. Therefore, if amisconfiguration causes routing information to come back to the originator, the information is ignored.

- Cluster-list is an optional, nontransitive BGP attribute. It is a sequence of cluster IDs that the route haspassed. When a route reflector reflects a route from its clients to nonclient peers, and vice versa, it appendsthe local cluster ID to the cluster-list. If the cluster-list is empty, a new cluster-list is created. Using this attribute, a route reflector can identify if routing information is looped back to the same cluster due to misconfiguration. If the local cluster ID is found in the cluster-list, the advertisement is ignored.

4.Which two statements correctly describe the BGP ttl-security feature? (Choose two.)

- A. This feature protects the BGP processes from CPU utilization-based attacks from EBGP neighbors which can be multiple hops away
- B. This feature prevents IBGP sessions with non-directly connected IBGP neighbors
- C. This feature will cause the EBGP updates from the router to be sent using a TTL of 1
- D. This feature needs to be configured on each participating BGP router
- E. This feature is used together with the ebgp-multihop command

Answer: A,D

5.When implementing source-based remote-triggered black hole filtering, which two configurations are required on the edge routers that are not the signaling router? (Choose two.)

- A. A static route to a prefix that is not used in the network with a next hop set to the Null0 interface
- B. A static route pointing to the IP address of the attacker

- C. uRPF on all external facing interfaces at the edge routers
- D. Redistribution into BGP of the static route that points to the IP address of the attacker
- E. A route policy to set the redistributed static routes with the no-export BGP community

Answer: A,C

Explanation:

Source-Based RTBH Filtering

With destination-based black holing, all traffic to a specific destination is dropped after the black hole has been activated, regardless of where it is coming from. Obviously, this could include legitimate traffic destined for the target. Source-based black holes provide the ability to drop traffic at the network edge based on a specific source address or range of source addresses.

If the source address (or range of addresses) of the attack can be identified (spoofed or not), it would be better to drop all traffic at the edge based on the source address, regardless of the destination address.

This would permit legitimate traffic from other sources to reach the target.

Implementation of source-based black hole filtering depends on Unicast Reverse Path Forwarding (uRPF), most often loose mode uRPF.

Loose mode uRPF checks the packet and forwards it if there is a route entry for the source IP of the incoming packet in the router forwarding information base (FIB). If the router does not have an FIB entry for the source IP address, or if the entry points to a null interface, the Reverse Path Forwarding (RPF) check fails and the packet is dropped, as shown in Figure

2. Because uRPF validates a source IP

address against its FIB entry, dropping traffic from specific source addresses is accomplished by configuring loose mode uRPF on the external interface and ensuring the RPF check fails by inserting a route to the source with a next hop of Null0.

This can be done by using a trigger device to send IBGP updates. These updates set the next hop for the source IP to an unused IP address that has a static entry at the edge, setting it to null as shown in Figure

2.

6. Refer to the topology diagram shown in the exhibit and the partial configurations shown below.

Once the attack from 209.165.201.144/28 to 209.165.202.128/28 has been detected, which additional configurations are required on the P1 IOS-XR router to implement source-based remote-triggered black hole filtering?

```
!  
router bgp 123  
address-family ipv4 unicast redistribute static route-policy test
```

```
!
```

A. router static

```
address-family ipv4 unicast  
209.165.202.128/28 null0 tag 666  
192.0.2.1/32 null0 tag 667
```

```
!  
route-policy test if tag is 666 then  
set next-hop 192.0.2.1 endif  
if tag is 667 then  
set community (no-export)  
endif  
end-policy
```

```
!  
B. router static  
address-family ipv4 unicast  
209.165.201.144/28 null0 tag 666  
192.0.2.1/32 null0 tag 667  
!  
route-policy test if tag is 666 then
```

```
set next-hop 192.0.2.1 endif
```

```
if tag is 667 then
```

```
set community (no-export)
```

```
endif
```

```
end-policy
```

```
!
```

```
C. router static
```

```
address-family ipv4 unicast
```

```
209.165.201.144/28 null0 tag 666
```

```
192.0.2.1/32 null0
```

```
!
```

```
route-policy test if tag is 666 then
```

```
set next-hop 192.0.2.1
```

```
set community (no-export)
```

```
endif
```

```
end-policy
```

```
D. router static
```

```
address-family ipv4 unicast
```

```
209.165.202.128/28 null0 tag 666
```

```
192.0.2.1/32 null0
```

```
!
```

```
route-policy test if tag is 666 then
```

```
set next-hop 192.0.2.1
```

```
set community (no-export)
```

```
endif
```

```
end-policy
```

!

Answer: C

Explanation:

Source-Based RTBH Filtering

With destination-based black holing, all traffic to a specific destination is dropped after the black hole has been activated, regardless of where it is coming from. Obviously, this could include legitimate traffic destined for the target. Source-based black holes provide the ability to drop traffic at the network edge based on a specific source address or range of source addresses.

If the source address (or range of addresses) of the attack can be identified (spoofed or not), it would be better to drop all traffic at the edge based on the source address, regardless of the destination address.

This would permit legitimate traffic from other sources to reach the target. Implementation of source-based black hole filtering depends on Unicast Reverse Path Forwarding

(uRPF), most often loose mode uRPF. Loose mode uRPF checks the packet and forwards it if there is a route entry for the source IP of the incoming packet in the router forwarding information base (FIB). If the router does not have an FIB entry for the source

IP address, or if the entry points to a null interface, the Reverse Path Forwarding (RPF) check fails and the packet is dropped, as shown in Figure

2. Because uRPF validates a source IP address against its

FIB entry, dropping traffic from specific source addresses is accomplished by configuring loose mode uRPF on the external interface and ensuring the RPF check fails by inserting a route to the source with a next hop of Null0. This can be done by using a trigger device to send IBGP updates. These updates set the next hop for the source IP to an unused IP address that has a static entry at the edge, setting it to null as shown in Figure 2.

7. In Cisco IOS-XR, the maximum-prefix command, to control the number of prefixes that can be installed from a BGP neighbor, is configured under which configuration mode?

A. RP/0/RSP0/CPU0:P2(config-bgp)#

- B. RP/0/RSP0/CPU0:P2(config-bgp-af)#
- C. RP/0/RSP0/CPU0:P2(config-bgp-nbr)#
- D. RP/0/RSP0/CPU0:P2(config-bgp-nbr-af)#

Answer: D

8. In Cisco IOS-XR, the ttl-security command is configured under which configuration mode?

- A. RP/0/RSP0/CPU0:P2(config)#
- B. RP/0/RSP0/CPU0:P2(config-bgp)#
- C. RP/0/RSP0/CPU0:P2(config-bgp-nbr)#
- D. RP/0/RSP0/CPU0:P2(config-bgp-af)#
- E. RP/0/RSP0/CPU0:P2(config-bgp-nbr-af)#

Answer: C

9. Refer to the exhibit.

Given the partial BGP configuration, which configuration correctly completes the Cisco IOS-XR route reflector configuration where both the 1.1.1.1 and 2.2.2.2 routers are the clients and the 3.3.3.3 router is a non-client IBGP peer?

A. neighbor 1.1.1.1 remote-as 65123 route-reflector-client neighbor 2.2.2.2 remote-as 65123 route-reflector-client

neighbor 3.3.3.3 remote-as 65123

B. neighbor 1.1.1.1

address-family ipv4 unicast remote-as 65123

route-reflector-client neighbor 2.2.2.2

address-family ipv4 unicast remote-as 65123

route-reflector-client neighbor 3.3.3.3

address-family ipv4 unicast remote-as 65123

C. neighbor 1.1.1.1 remote-as 65123

address-family ipv4 unicast route-reflector-client neighbor 2.2.2.2

remote-as 65123

address-family ipv4 unicast route-reflector-client neighbor 3.3.3.3

remote-as 65123

D. neighbor 1.1.1.1 remote-as 65123 neighbor 1.1.1.1 route-reflector-client neighbor 2.2.2.2 remote-as

65123 neighbor 2.2.2.2 route-reflector-client neighbor 3.3.3.3 remote-as 65123

Answer: C

10. Which three methods can be used to reduce the full-mesh IBGP requirement in a service provider core network? (Choose three.)

A. Implement route reflectors

B. Enable multi-protocol BGP sessions between all the PE routers

C. Implement confederations

D. Implement MPLS (LDP) in the core network on all the PE and P routers

E. Enable BGP synchronization

F. Disable the IBGP split-horizon rule

Answer: A,C,D

11. Which type of BGP session behaves like an EBGP session during session establishment but behaves like an IBGP session when propagating routing updates where the local preference, multi-exit discriminator, and next-hop attributes are not changed?

A. BGP sessions between a route reflector and its clients

B. BGP sessions between a route reflector and its non-client IBGP peers

C. BGP sessions between a route reflector and another route reflector

D. Intra-confederation IBGP sessions

E. Intra-confederation EBGP sessions

Answer: E

Explanation:

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.7/routing/configuration/guide/rc37bgp.html#wp1191371

BGP Routing Domain Confederation

One way to reduce the iBGP mesh is to divide an autonomous system into multiple subautonomous systems and group them into a single confederation. To the outside world, the confederation looks like a single autonomous system. Each autonomous system is fully meshed within itself and has a few connections to other autonomous systems in the same confederation. Although the peers in different autonomous systems have EBGP sessions, they exchange routing information as if they were iBGP peers. Specifically, the next hop, MED, and local preference information is preserved. This feature allows you to retain a single IGP for all of the autonomous systems.

12. You noticed a recent change to the BGP configuration on a PE router, the bgp scan time has been changed from the default value to 30s. Which three effects will this change have? (Choose three.)

- A. The BGP table will be examined and verified more frequently
- B. The BGP keepalive messages will be sent to the BGP peers at a faster rate
- C. The BGP table will be modified more quickly in the event that a next-hop address becomes unreachable
- D. The CPU load of the router will increase
- E. The minimum time interval between sending EBGP and IBGP routing updates will decrease
- F. The BGP convergence time will increase

Answer: A,C,D

Explanation:

13. On Cisco IOS-XR, which BGP process can be distributed into multiple instances?

- A. BGP process manager
- B. BGP RIB process
- C. BGP speaker process
- D. BGP scanner process
- E. BGP dampening process

Answer: C

Explanation:

Cisco IOS XR allows you to control the configuration of the number of distributed speakers and enables you to selectively assign neighbors to specific speakers. On the CRS-1 platform, multiple speaker processes up to 15 may be configured. However, configuring all the different speakers on the primary route processor simply adds to the load on the single RP.

Distributed speaker functionality is useful if Distributed Route Processor (DRP) hardware is available to take advantage of process placement. Later sections in this chapter depict distributed BGP and placement of BGP process speakers on DRPs on a CRS-1 router.

In addition to the speaker process, BPM starts the bRIB process once BGP is configured.

bRIB process is responsible for performing the best-path calculation based on partial best paths received from the speaker processes. The best route is installed into the bRIB and is advertised back to all speakers. The bRIB process is also responsible for installing routes

14. Refer to the configuration exhibit, taken from a Cisco IOS-XR router.

Which configuration change is required to properly enable this router as the signaling router for implementing source-based RTBH filtering?

- A. Set community (no-export) in the route policy
- B. Pass in the route policy
- C. Set local-preference 1000 in the route policy

D. The 192.0.2.1/32 static route should be tagged as 666 (tag 666)

Answer: A

15. Refer to the Cisco IOS-XR BGP configuration exhibit.

Identify two configuration errors. (Choose two.)

A. The neighbor-group efg is missing the ebgp-multihop 2 configuration

B. The ttl-security configuration command is missing the option to set the number of hops

C. The passall route policy is wrong

D. The route-policy passall in and route-policy passall out commands should be configured under the neighbor-group efg instead of the af-group abc

E. The maximum-prefix 10 configuration should be configured under the af-group abc instead of the neighbor-group efg

Answer: C,E

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 642-885 Exam with Our Prep Materials Via below:

<https://www.certleader.com/642-885-dumps.html>