TestpassportQ&A

Higher Quality Better Service!

We offer free update service for one year

Http://www.testpassport.com

Exam : **350-501**

Title : Implementing and Operating Cisco Service Provider Network Core Technologies

Version : DEMO

1.Egress PE NAT is being used via a single centralized router to provide Internet access to L3VPN customers.
Which description of the NAT operation is true?
A. Users m different VRFs cannot share the same outside global IP address
B. The NAT table contains a field to identify the inside VRF of a translation
C. Multiple address pools are needed for the same L3VPN because each site has a separate NAT
D. The different L3VPNs using the Internet access must not have IP overlaps internally
Answer: B
Explanation:
Egress PE NAT via a single centralized router allows L3VPN customers to access the internet. The NAT operation maintains a table that includes a field to identify the inside VRF of a translation. This is crucial for distinguishing between different customer VPNs that may use overlapping IP address spaces. By having this field, the NAT process can correctly associate each internal IP with its respective VRF and outside global IP address1.
Reference: Centralized egress to internet, Egress PE NAT single centralized router.

2.How much must the MTU be increased when configuring the 802.1q VLAN tag?
A. 2 bytes
B. 4 bytes
C. 8 bytes
D. 12 bytes
Answer: B
Explanation:
When configuring the 802.1q VLAN tag, the MTU must be increased by 4 bytes. This is because the 802.1Q tag adds an additional header to the Ethernet frame, which includes information such as the VLAN ID and priority. To accommodate this extra header without fragmenting the packet, the MTU must be increased accordingly2.
Reference: Configure 802.1Q VLAN Tagging, 802.1q and MTU.

3.Refer to the exhibit:

```
ip flow-export source loopback 0
ip-flow-export destination 192.168.1.1
ip-flow-export version 9 origin-as
```

Export statistics received do not include the BGP next hop.
Which statement about the NetFlow export statistics is true?
A. Only the origin AS of the source router will be included in the export statistics.
B. Loopback 0 must be participating in BGP for it to be included in the export statistics.
C. The origin AS and the peer-as will be included in the export statistics.
D. To include the BGP next hop in the export statistics, those keywords must be included with the version 9 entry.
Answer: D
Explanation:
In NetFlow export statistics, the BGP next hop information is not included by default. To include the BGP

next hop in the export statistics, specific keywords related to BGP next hop must be configured with the version 9 NetFlow entry. This configuration allows for the measurement of network traffic on a per-BGP next hop basis, which is essential for detailed traffic analysis and accounting1.
Reference: Configuring NetFlow BGP Next Hop Support for Accounting and Analysis - Cisco Systems1.

4.Refer to the exhibit:

```
PE-A#show ip bgp vpnv4 vrf Customer-A neighbors 10.10.10.2 routes
  BGP table version is 13148019, local router ID is 10.10.10.10
  Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
                r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
                x best-external, a additional-path, c RIB-compressed,
  Origin codes: i - IGP, e - EGP, ? - incomplete
  RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 Route Distinguisher: 65000:1111 (default for vrf Customer-A)
 *>    192.168.0.0/19    10.10.10.2               0         0 4282 65001 ?
 *>    192.168.0.0/17    10.10.10.2               0         0 4282 65001 ?
 *>    192.168.0.0/16    10.10.10.2               0         0 4282 65001 ?

Total number of prefixes 5

PE-A#config t
  Enter configuration commands, one per line. End with CNTL/Z.
  PE-A(config)#ip prefix-list ALLOW permit 192.168.0.0/16 ge 17 le 19
  PE-A(config)#router bgp 65000
  PE-A(config-router)#address-family ipv4 vrf Customer-A
  PE-A(config-router-af)#neighbor 10.10.10.2 prefix-list ALLOW in
```

Which three outcomes occur if the prefix list is added to the neighbor? (Choose three)
A. 192.168 0.0/19 is denied.
B. 192.168 0.0/17 is denied.
C. 192.168 0.0/17 is permitted
D. 192.168.0.0/16 is denied
E. 192.168 0.0/16 is permitted
F. 192.168 0.0/19 is permitted
Answer: A, D, E
Explanation:
When a prefix list is applied to a BGP neighbor, it filters routes based on the specified conditions. In this scenario, the prefix list is designed to permit prefixes that are equal to or longer than /17 but shorter than or equal to /19.
Therefore:
A: 192.168.0.0/19 is denied because it falls outside the specified range.
D: 192.168.0.0/16 is denied as it is shorter than the minimum length of /17.
E: 192.168.0.0/16 is permitted because it is the exact prefix specified in the prefix list, and the ge and le modifiers apply to more specific routes derived from this prefix.
Reference: Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR) - Cisco official courseware.

5.Which statement about segment routing prefix segments is true?

A. It is linked to a prefix SID that is globally unique within segment routing domain.
B. It is the longest path to a node.
C. It is linked to an adjacency SID that is globally unique within the router.
D. It requires using EIGRP to operate.

Answer: A

Explanation:

Segment routing prefix segments are linked to a prefix SID that is globally unique within the segment routing domain. This means that each prefix SID identifies a specific prefix in the network and is unique across the entire segment routing-enabled domain, ensuring that the path to the prefix can be identified and used by any node within the domain1.

Reference: Introduction to Segment Routing - Cisco Learning Network2, ASR9000/XR Introduction to Segment Routing - Cisco Community3, About Segment Routing - Cisco Content Hub