

Juniper

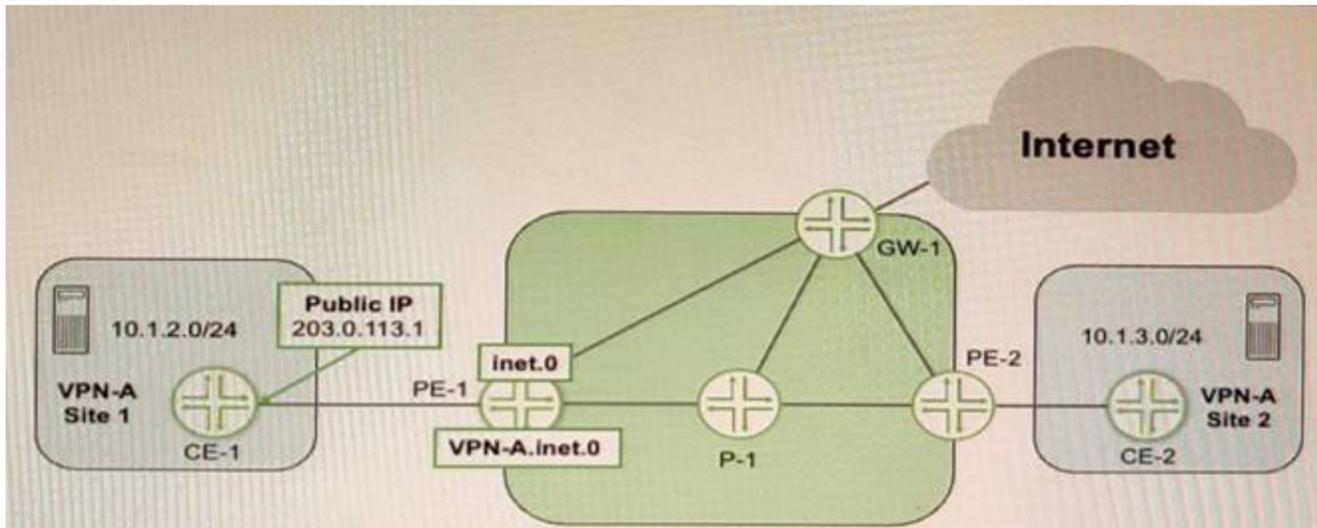
Exam Questions JN0-664

Service Provider - Professional (JNCIP-SP)



NEW QUESTION 1

Exhibit



Referring to the exhibit, CE-1 is providing NAT services for the hosts at Site 1 and you must provide Internet access for those hosts. Which two statements are correct in this scenario? (Choose two.)

- A. You must configure a static route in the main routing instance for the 10.1.2.0/24 prefix that uses the VPN-A.inet.0 table as the next hop.
- B. You must configure a static route in the main routing instance for the 203.0.113.1/32 prefix that uses the VPN-A.inet.0 table as the next hop.
- C. You must configure a RIB group on PE-1 to leak a default route from the inet.0 table to the VPN-A.inet.0 table.
- D. You must configure a RIB group on PE-1 to leak the 10.1.2.0/24 prefix from the VPN-A.inet.0 table to the inet.0 table.

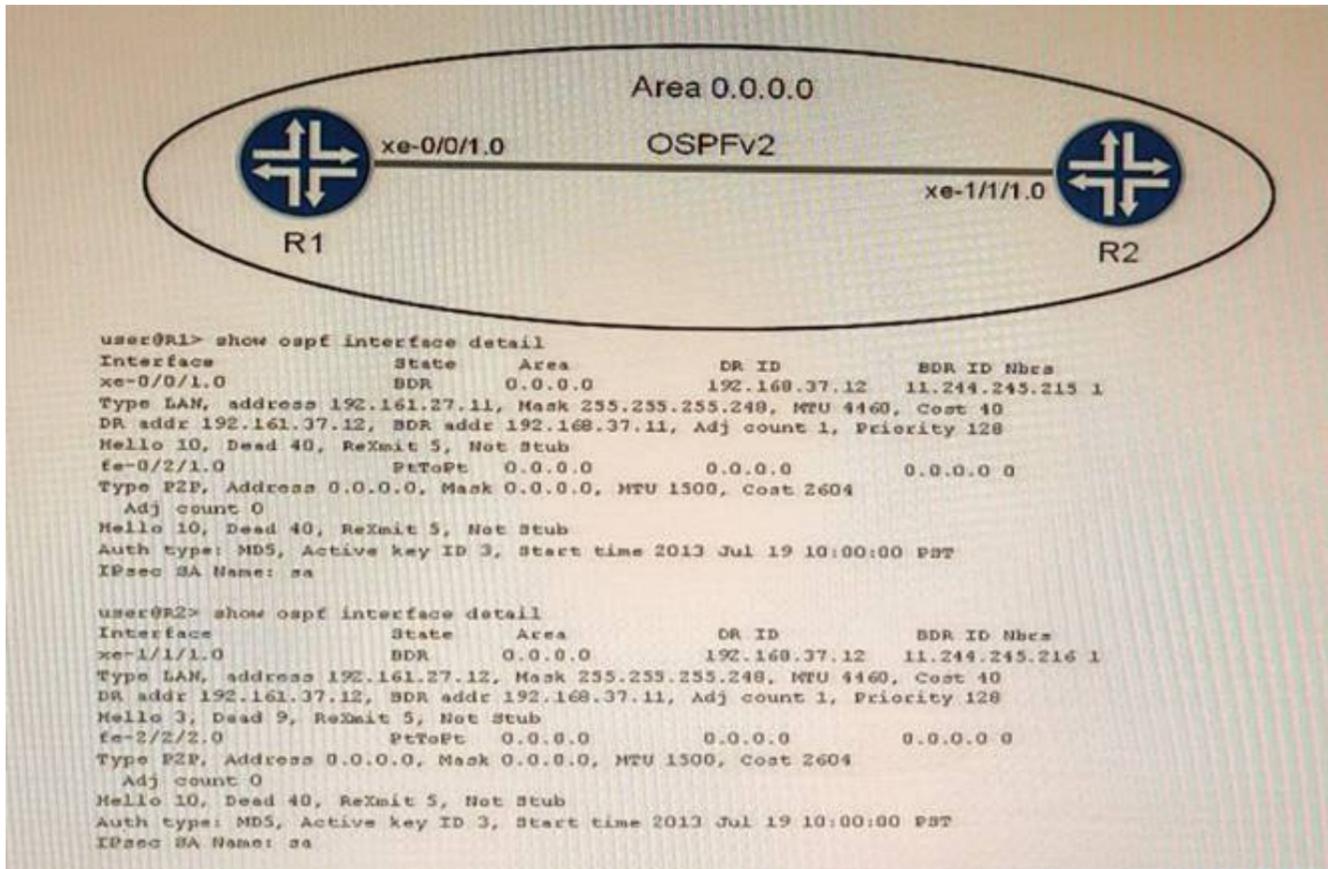
Answer: AB

Explanation:

To provide Internet access for the hosts at Site 1, you need to configure static routes in the main routing instance on PE-1 that point to the VPN-A.inet.0 table as the next hop. This allows PE-1 to forward traffic from the Internet to CE-1 using MPLS labels and vice versa. You need to configure two static routes: one for the 10.1.2.0/24 prefix that represents the private network of Site 1, and one for the 203.0.113.1/32 prefix that represents the public IP address of CE-1.

NEW QUESTION 2

Exhibit



Which two statements are true about the OSPF adjacency displayed in the exhibit? (Choose two.)

- A. There is a mismatch in the hello interval parameter between routers R1 and R2.
- B. There is a mismatch in the dead interval parameter between routers R1 and R2.
- C. There is a mismatch in the OSPF hold timer parameter between routers R1 and R2.
- D. There is a mismatch in the poll interval parameter between routers R1 and R2.

Answer: AB

Explanation:

The hello interval is the time interval between two consecutive hello packets sent by an OSPF router on an interface. The dead interval is the time interval after which a neighbor is declared down if no hello packets are received from it. These parameters must match between two OSPF routers for them to form an adjacency. In the exhibit, router R1 has a hello interval of 10 seconds and a dead interval of 40 seconds, while router R2 has a hello interval of 30 seconds and a dead interval of 120 seconds. This causes a mismatch and prevents them from becoming neighbors.

NEW QUESTION 3

Which two statements describe PIM-SM? (Choose two)

- A. Routers with receivers send join messages to their upstream neighbors.
- B. Routers without receivers must periodically prune themselves from the SPT.
- C. Traffic is initially flooded to all routers and an S,G is maintained for each group
- D. Traffic is only forwarded to routers that request to join the distribution tree.

Answer: AD

Explanation:

PIM sparse mode (PIM-SM) is a multicast routing protocol that uses a pull model to deliver multicast traffic. In PIM-SM, routers with receivers send join messages to their upstream neighbors toward a rendezvous point (RP) or a source-specific tree (SPT). The RP or SPT acts as the root of a shared distribution tree for a multicast group. Traffic is only forwarded to routers that request to join the distribution tree by sending join messages. PIM-SM does not flood traffic to all routers or prune routers without receivers, as PIM dense mode does.

NEW QUESTION 4

Which two statements are correct about VPLS tunnels? (Choose two.)

- A. LDP-signaled VPLS tunnels only support control bit 0.
- B. LDP-signaled VPLS tunnels use auto-discovery to provision sites
- C. BGP-signaled VPLS tunnels can use either RSVP or LDP between the PE routers.
- D. BGP-signaled VPLS tunnels require manual provisioning of sites.

Answer: BC

Explanation:

VPLS is a Layer 2 VPN technology that allows multiple sites to connect over a shared IP/MPLS network as if they were on the same LAN. VPLS tunnels can be signaled using either Label Distribution Protocol (LDP) or Border Gateway Protocol (BGP). LDP-signaled VPLS tunnels use auto-discovery to provision sites, meaning that PE routers can automatically discover other PE routers that belong to the same VPLS instance

NEW QUESTION 5

Exhibit

```
user@RI show configuration interpolated-profile { interpolate {
fill-level [ 50 75 drop—probability [ > }
class-of-service drop-profiles
];
20 60 ];
```

Which two statements are correct about the class-of-service configuration shown in the exhibit? (Choose two.)

- A. The drop probability jumps immediately from 20% to 60% when the queue level reaches 75% full.
- B. The drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full
- C. To use this drop profile, you reference it in a scheduler.
- D. To use this drop profile, you apply it directly to an interface.

Answer: BC

Explanation:

class-of-service (CoS) is a feature that allows you to prioritize and manage network traffic based on various criteria, such as application type, user group, or packet loss priority. CoS uses different components to classify, mark, queue, schedule, shape, and drop traffic according to the configured policies. One of the components of CoS is drop profiles, which define how packets are dropped when a queue is congested. Drop profiles use random early detection (RED) algorithm to drop packets randomly before the queue is full, which helps to avoid global synchronization and improve network performance. Drop profiles can be discrete or interpolated. A discrete drop profile maps a specific fill level of a queue to a specific drop probability. An interpolated drop profile maps a range of fill levels of a queue to a range of drop probabilities and interpolates the values in between.

In the exhibit, we can see that the class-of-service configuration shows an interpolated drop profile with two fill levels (50 and 75) and two drop probabilities (20 and 60). Based on this configuration, we can infer the following statements:

? The drop probability jumps immediately from 20% to 60% when the queue level reaches 75% full. This is not correct because the drop profile is interpolated, not discrete. This means that the drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full. The drop probability for any fill level between 50% and 75% can be calculated by using linear interpolation formula.

? The drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full. This is correct because the drop profile is interpolated and uses linear interpolation formula to calculate the drop probability for any fill level between 50% and 75%. For example, if the fill level is 60%, the drop probability is 28%, which is calculated by using the formula: $(60 - 50) / (75 - 50) * (60 - 20) + 20 = 28$.

? To use this drop profile, you reference it in a scheduler. This is correct because a scheduler is a component of CoS that determines how packets are dequeued from different queues and transmitted on an interface. A scheduler can reference a drop profile by using the random-detect statement under the [edit class-of-service schedulers] hierarchy level. For example: scheduler test { transmit-rate percent 10; buffer-size percent 10; random-detect test-profile; }

? To use this drop profile, you apply it directly to an interface. This is not correct because a drop profile cannot be applied directly to an interface. A drop profile can only be referenced by a scheduler, which can be applied to an interface by using the scheduler-map statement under the [edit class-of-service interfaces] hierarchy level. For example: interfaces ge-0/0/0 { unit 0 { scheduler-map test-map; } }

NEW QUESTION 6

Exhibit

```
[edit routing-instances CE-1]
user@R1# show
protocols {
  bgp {
    group CE-1 {
      type external;
      peer-as 65555;
      neighbor 10.1.1.100;
    }
  }
}
instance-type vrf;
interface ge-0/0/2.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
[edit routing-instances CE-2]
user@R2# show
protocols {
  bgp {
    group CE-2 {
      type external;
      peer-as 64444;
      neighbor 10.1.5.100;
    }
  }
}
instance-type vrf;
interface ge-0/0/3.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
```

Referring to the exhibit, which statement is correct?

- A. The vrf-target configuration will allow routes to be shared between CE-1 and CE-2.
- B. The vrf-target configuration will stop routes from being shared between CE-1 and CE-2.
- C. The route-distinguisher configuration will allow overlapping routes to be shared between CE-1 and CE-2.
- D. The route-distinguisher configuration will stop routes from being shared between CE-1 and CE-2.

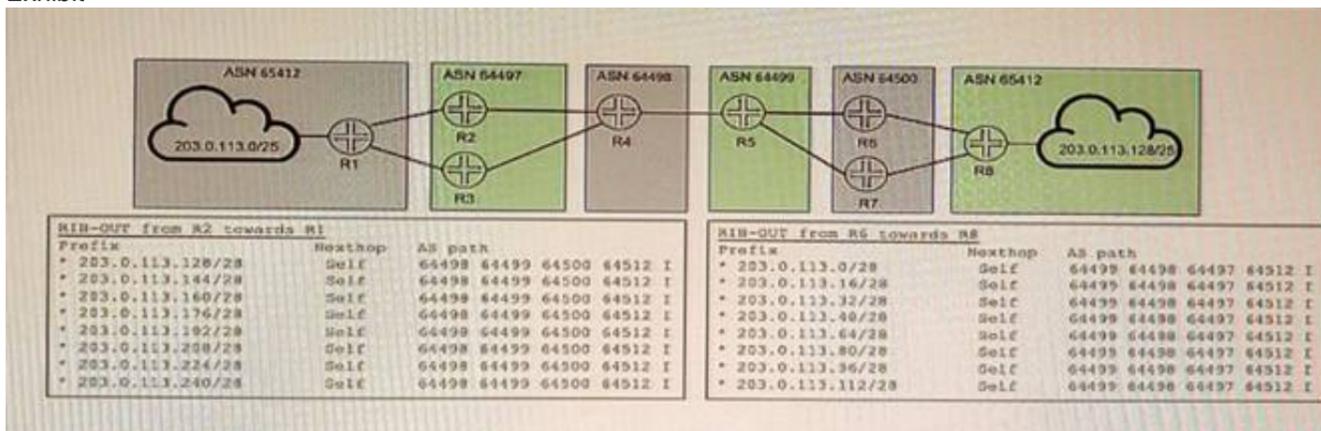
Answer: C

Explanation:

The route distinguisher (RD) is a BGP attribute that is used to create unique VPN IPv4 prefixes for each VPN in an MPLS network. The RD is a 64-bit value that consists of two parts: an administrator field and an assigned number field. The administrator field can be an AS number or an IP address, and the assigned number field can be any arbitrary value chosen by the administrator. The RD is prepended to the IPv4 prefix to create a VPN IPv4 prefix that can be advertised across the MPLS network without causing any overlap or conflict with other VPNs. In this question, we have two PE routers (PE-1 and PE-2) that are connected to two CE devices (CE-1 and CE-2) respectively. PE-1 and PE-2 are configured with VRFs named Customer-A and Customer-B respectively.

NEW QUESTION 7

Exhibit



R1 and R8 are not receiving each other's routes

Referring to the exhibit, what are three configuration commands that would solve this problem? (Choose three.)

- A. Configure loops and advertise-peer-as on routers in AS 64497 and AS 64450.
- B. Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498.

- C. Configure as-override on advertisement from AS 64500 toward AS 64512.
- D. Configure remove-private on advertisements from AS 64497 toward AS 64498
- E. Configure remove-private on advertisements from AS 64500 toward AS 64499

Answer: BDE

Explanation:

The problem in this scenario is that R1 and R8 are not receiving each other's routes because of private AS numbers in the AS path. Private AS numbers are not globally unique and are not advertised to external BGP peers. To solve this problem, you need to do the following:

? Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498. This allows R5 and R6 to advertise their own AS number (65412) instead of their peer's AS number (64498) when sending updates to R7 and R8. This prevents a loop detection issue that would cause R7 and R8 to reject the routes from R5 and R62.

? Configure remove-private on advertisements from AS 64497 toward AS 64498 and from AS 64500 toward AS 64499. This removes any private AS numbers from the AS path before sending updates to external BGP peers. This allows R2 and R3 to receive the routes from R1 and R4, respectively3.

NEW QUESTION 8

Exhibit

```

user@PE1# show routing-instances
VPN-A {
  instance-type vrf;
  interface ge-0/0/1.0;
  vrf-target target:64512:1234;
  protocols {
    bgp {
      group CE {
        type external;
        family inet {
          unicast;
        }
        neighbor 10.0.0.1 {
          peer-as 64512;
          as-override;
        }
      }
    }
  }
}

```

Which two statements about the configuration shown in the exhibit are correct? (Choose two.)

- A. This VPN connects customer sites that use different AS numbers.
- B. This VPN connects customer sites that use the same AS number
- C. A Layer 2 VPN is configured.
- D. A Layer 3 VPN is configured.

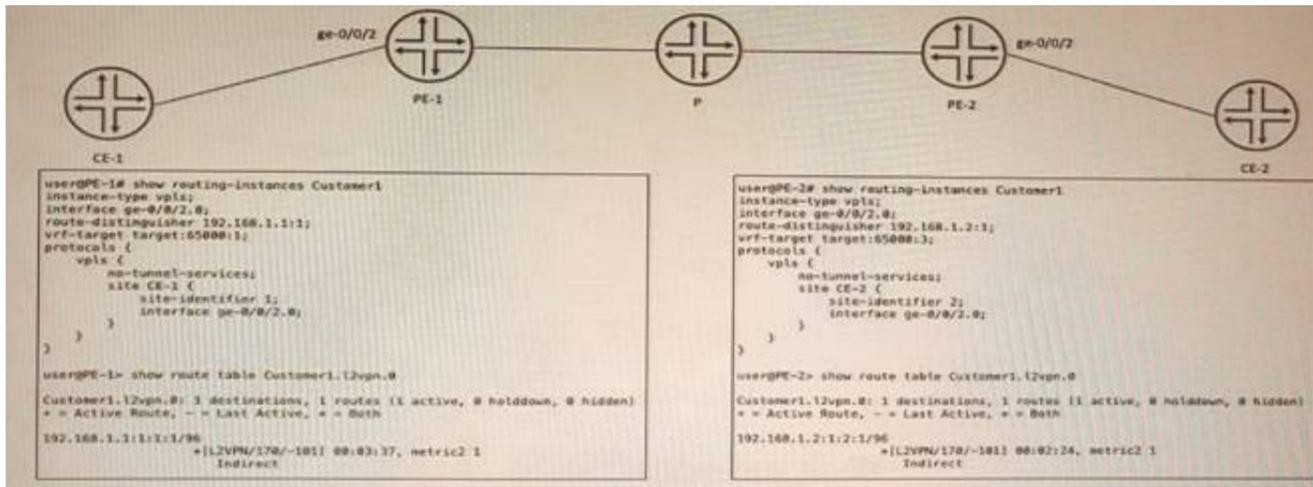
Answer: AD

Explanation:

The configuration shown in the exhibit is for a Layer 3 VPN that connects customer sites that use different AS numbers. A Layer 3 VPN is a type of VPN that uses MPLS labels to forward packets across a provider network and BGP to exchange routing information between PE routers and CE routers. A Layer 3 VPN allows customers to use different routing protocols and AS numbers at their sites, as long as they can peer with BGP at the PE-CE interface. In this example, CE-1 is using AS 65530 and CE-2 is using AS 65531, but they can still communicate through the VPN because they have BGP sessions with PE-1 and PE-2, respectively.

NEW QUESTION 9

Exhibit



CE-1 and CE-2 are part of a VPLS called Customer1. No connectivity exists between CE-1 and CE-2. In the process of troubleshooting, you notice PE-1 is not learning any routes for this VPLS from PE-2, and PE-2 is not learning any routes for this VPLS from PE-1.

- A. The route target must match on PE-1 and PE-2.
- B. The route distinguisher must match on PE-1 and PE-2.
- C. The instance type should be changed to l2vpn.
- D. The no-tunnel-services statement should be deleted on both PEs.

Answer: A

Explanation:

VPLS is a technology that provides Layer 2 VPN services over an MPLS network. VPLS uses BGP as its control protocol to exchange VPN membership information between PE routers. The route target is a BGP extended community attribute that identifies which VPN a route belongs to. The route target must match on PE routers that participate in the same VPLS instance, otherwise they will not accept or advertise routes for that VPLS.

NEW QUESTION 10

Exhibit

```
[edit routing-instances CE-1]
user@router# show
routing-options {
  static {
    route 10.101.1.0/24 next-hop 10.1.1.100;
  }
}
instance-type vrf;
interface ge-0/0/2.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
```

Referring to the exhibit, which statement is true?

- A. The 10.101.1.0/24 route will be shared if the vrf-table-label parameter is configured.
- B. The 10.101.1.0/24 route will only be shared if BGP is configured in the routing instance
- C. The 10.101.1.0/24 route will be shared if there are other VRFs that use the same route target community
- D. The 10.101.1.0/24 route will be shared if the auto-export parameter is configured

Answer: D

Explanation:

The auto-export parameter is a routing option that allows a routing instance to share routes with other routing instances or the master routing table. The auto-export parameter automatically exports routes from one routing instance to another based on the route target communities attached to the routes. In this scenario, the 10.101.1.0/24 route will be shared if the auto-export parameter is configured under [edit routing-options] hierarchy level.

NEW QUESTION 10

When using OSPFv3 for an IPv4 environment, which statement is correct?

- A. OSPFv3 only supports IPv4.
- B. OSPFv3 supports both IPv6 and IPv4, but not in the same routing instance.
- C. OSPFv3 is not backward compatible with IPv4
- D. OSPFv3 supports IPv4 only on interfaces with family inet6 defined

Answer: C

Explanation:

OSPFv3 is an extension of OSPFv2 that supports IPv6 routing and addressing. OSPFv3 is not backward compatible with IPv4 because it uses a different packet format and a different link-state advertisement (LSA) structure than OSPFv2. OSPFv3 also uses IPv6 link-local addresses as router IDs and neighbor addresses, instead of IPv4 addresses. To use OSPFv3 for an IPv4 environment, you need to enable the IPv4 unicast address family under [edit protocols ospf3] hierarchy level and configure IPv4 addresses on the interfaces.

NEW QUESTION 12

Your organization manages a Layer 3 VPN for multiple customers. To support advanced route than one BGP community on advertised VPN routes to remote PE routers.

Which routing-instance configuration parameter would support this requirement?

- A. vrf-export
- B. vrf-import
- C. vrf-target export
- D. vrf-target import

Answer: C

Explanation:

The vrf-target export parameter is used to specify one or more BGP extended community attributes that are attached to VPN routes when they are exported from a VRF routing instance to remote PE routers. This parameter allows you to control which VPN routes are accepted by remote PE routers based on their import policies. You can specify more than one vrf-target export value for a VRF routing instance to support advanced route filtering or route leaking scenarios.

NEW QUESTION 14

After a recent power outage, your manager asks you to investigate ways to automatically reduce the impact caused by suboptimal routing in your OSPF and OSPFv3 network after devices reboot.

Which three configuration statements accomplish this task? (Choose three.)

- A. set protocols ospf overload timeout 900
- B. set protocols ospf3 realm ipv4-unicast overload timeout 900
- C. set protocols ospf overload
- D. set protocols ospf3 overload timeout 900
- E. set protocols ospf3 overload

Answer: AE

Explanation:

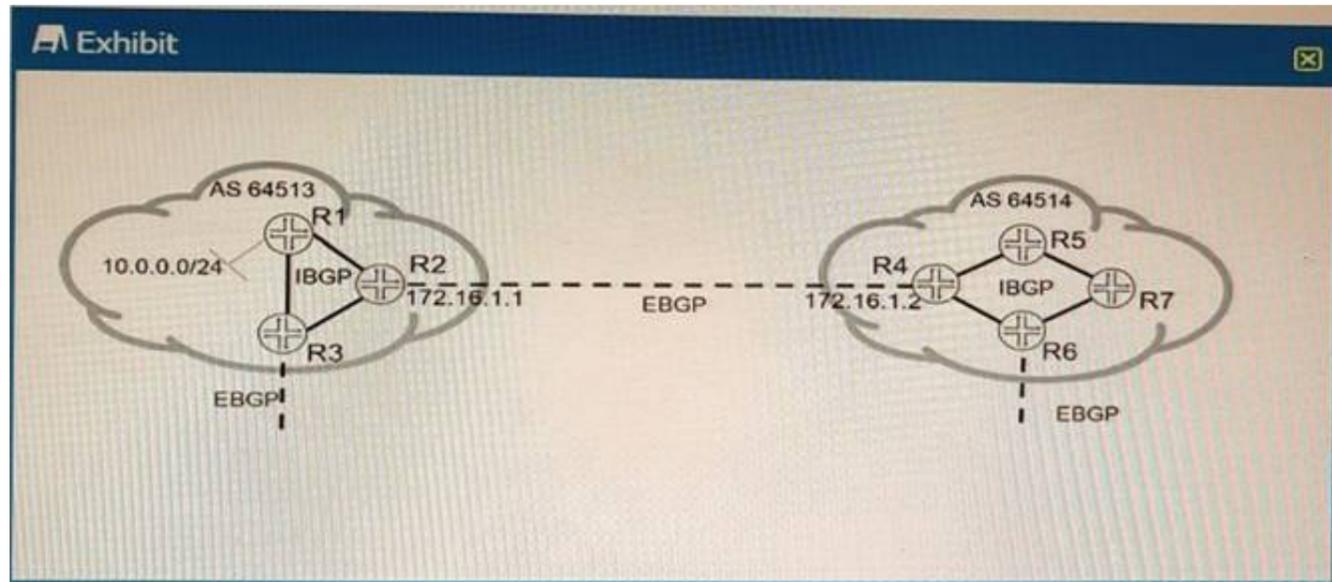
To reduce the impact of suboptimal routing in OSPF and OSPFv3 after devices reboot, you can use the overload feature to prevent a router from being used as a transit router for a specified period of time. This allows the router to stabilize its routing table before forwarding traffic for other routers. To enable the overload feature, you need to do the following:

? For OSPF, configure the overload statement under [edit protocols ospf] hierarchy level. You can also specify a timeout value in seconds to indicate how long the router should remain in overload state after it boots up. For example, set protocols ospf overload timeout 900 means that the router will be in overload state for 15 minutes after it boots up.

? For OSPFv3, configure the overload statement under [edit protocols ospf3] hierarchy level. You can also specify a realm (ipv4-unicast or ipv6-unicast) and a timeout value in seconds to indicate how long the router should remain in overload state after it boots up for each realm. For example, set protocols ospf3 realm ipv4-unicast overload timeout 900 means that the router will be in overload state for 15 minutes after it boots up for IPv4 unicast routing.

NEW QUESTION 17

Exhibit.



Referring to the exhibit; the 10.0.0.0/24 EBGP route is received on R5; however, the route is being hidden. What are two solutions that will solve this problem? (Choose two.)

- A. On R4, create a policy to change the BGP next hop to itself and apply it to IBGP as an export policy
- B. Add the external interface prefix to the IGP routing tables
- C. Add the internal interface prefix to the BGP routing tables.
- D. On R4, create a policy to change the BGP next hop to 172.16.1.1 and apply it to IBGP as an export policy

Answer: AB

Explanation:

the default behavior for iBGP is to propagate EBGP-learned prefixes without changing the next-hop. This can cause issues if the next-hop is not reachable via the IGP. One solution is to use the next-hop self command on R4, which will change the next-hop attribute to its own loopback address. This way, R5 can reach the next-hop via the IGP and install the route in its routing table.

Another solution is to add the external interface prefix (120.0.4.16/30) to the IGP routing tables of R4 and R5. This will also make the next-hop reachable via the IGP and allow R5 to use the route. According to 2, this is a possible workaround for a pure IP network, but it may not work well for an MPLS network.

NEW QUESTION 19

A packet is received on an interface configured with transmission scheduling. One of the configured queues in this scenario, which two actions will be taken by default on a Junos device? (Choose two.)

- A. The excess traffic will be discarded
- B. The exceeding queue will be considered to have negative bandwidth credit.
- C. The excess traffic will use bandwidth available from other queues
- D. The exceeding queue will be considered to have positive bandwidth credit

Answer: AB

Explanation:

Transmission scheduling is a CoS feature that allows you to allocate bandwidth among different queues on an interface. Each queue has a configured bandwidth percentage that determines how much of the available bandwidth it can use. If a queue exceeds its allocated bandwidth, it is considered to have negative bandwidth credit and its excess traffic will be discarded by default. If a queue does not use all of its allocated bandwidth, it is considered to have positive bandwidth credit and its unused bandwidth can be shared by other queues.

NEW QUESTION 21

Exhibit

R8 Routing Table

203.0.113.128/28	* [BGP/170]
203.0.113.144/28	* [BGP/170]
203.0.113.160/28	* [BGP/170]
203.0.113.176/28	* [BGP/170]
203.0.113.192/28	* [BGP/170]
203.0.113.208/28	* [BGP/170]
203.0.113.224/28	* [BGP/170]
203.0.113.240/28	* [BGP/170]

```

user@R8> show configuration policy-options policy-statement adv-routes
term 10 {
  from {
    protocol bgp;
    route-filter 203.0.113.128/25 exact;
  }
  then accept;
}
term last {
  then reject;
}
    
```

You are attempting to summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500. You implement the export policy shown in the exhibit and all routes from the routing table stop being advertised.

In this scenario, which two steps would you take to summarize the route in BGP? (Choose two.)

- A. Remove the from protocol bgp command from the export policy.
- B. Add the set protocols bgp family inet unicast add-path command to allow additional routes to the RIB table
- C. -
- D. Add the set routing-options static route 203.0.113.123/25 discard command.
- E. Replace exact in the export policy with orlonger.

Answer: CD

Explanation:

To summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500, you need to do the following:

- ? Add the set routing-options static route 203.0.113.128/25 discard command. This creates a static route for the summary prefix and discards any traffic destined to it. This is necessary because BGP can only advertise routes that are present in the routing table.
- ? Replace exact in the export policy with orlonger. This allows R8 to match and advertise any route that is equal or more specific than the summary prefix. The exact term only matches routes that are exactly equal to the summary prefix, which is not present in the routing table.

NEW QUESTION 26

You are a network architect for a service provider and want to offer Layer 2 services to your customers. You want to use EVPN for Layer 2 services in your existing MPLS network.

Which two statements are correct in this scenario? (Choose two.)

- A. Segment routing must be configured on all PE routers.
- B. VXLAN must be configured on all PE routers.
- C. EVPN uses Type 2 routes to advertise MAC address and IP address pairs learned using ARP snooping
- D. EVPN uses Type 3 routes to join a multicast tree to flood traffic.

Answer: CD

Explanation:

EVPN is a technology that connects L2 network segments separated by an L3 network using a virtual Layer 2 network overlay over the Layer 3 network. EVPN uses BGP as its control protocol to exchange different types of routes for different purposes. Type 2 routes are used to advertise MAC address and IP address pairs learned using ARP snooping from the local CE devices. Type 3 routes are used to join a multicast tree to flood traffic such as broadcast, unknown unicast, and multicast (BUM) traffic.

NEW QUESTION 27

Exhibit

```

user@router> show l2vpn connections
Layer-2 VPN connections:
Legend for connection status (St)
EI -- encapsulation invalid          NC -- interface encapsulation not
CCC/TCC/VPLS                         interface and instance encaps not same
EM -- encapsulation mismatch         WE -- interface hardware not present
VC-Dn -- Virtual circuit down        -> -- only outbound connection is up
CM -- control-word mismatch          <- -- only inbound connection is up
CN -- circuit not provisioned        Up -- operational
OR -- out of range                   Dn -- down
OL -- no outgoing label              CF -- call admission control failure
LD -- local site signaled down       SC -- local and remote site ID collision
RD -- remote site signaled down      LM -- local site ID not minimum designated
LN -- local site not designated      RM -- remote site ID not minimum designated
RN -- remote site not designated     IL -- no incoming label
XX -- unknown connection status      MI -- Mesh-Group ID not available
MM -- MTU mismatch                   ST -- Standby connection
BK -- Backup connection              PB -- Profile busy
PF -- Profile parse failure           SN -- Static Neighbor
RS -- remote site standby            RB -- Remote site not best-site
LB -- Local site not best-site       HS -- Hot-standby Connection
VM -- VLAN ID mismatch
Legend for interface status
Up -- operational
Dn -- down
Instance: vpn-A
Edge protection: Not-Primary
Local site: CE1-2 (2)
connection-site Type St      Time last up          # Up trans
1               rmt  Up      Apr 11 14:35:27 2020 1
Remote PE: 172.17.20.1, Negotiated control-word: Yes (Null)
Incoming label: 21, Outgoing label: 22
Local interface: ge-0/0/6.610, Status: Up, Encapsulation: VLAN
Flow Label Transmit: No, Flow Label Receive: No
    
```

Which two statements about the output shown in the exhibit are correct? (Choose two.)

- A. The PE is attached to a single local site.
- B. The connection has not flapped since it was initiated.
- C. There has been a VLAN ID mismatch.
- D. The PE router has the capability to pop flow labels

Answer: AD

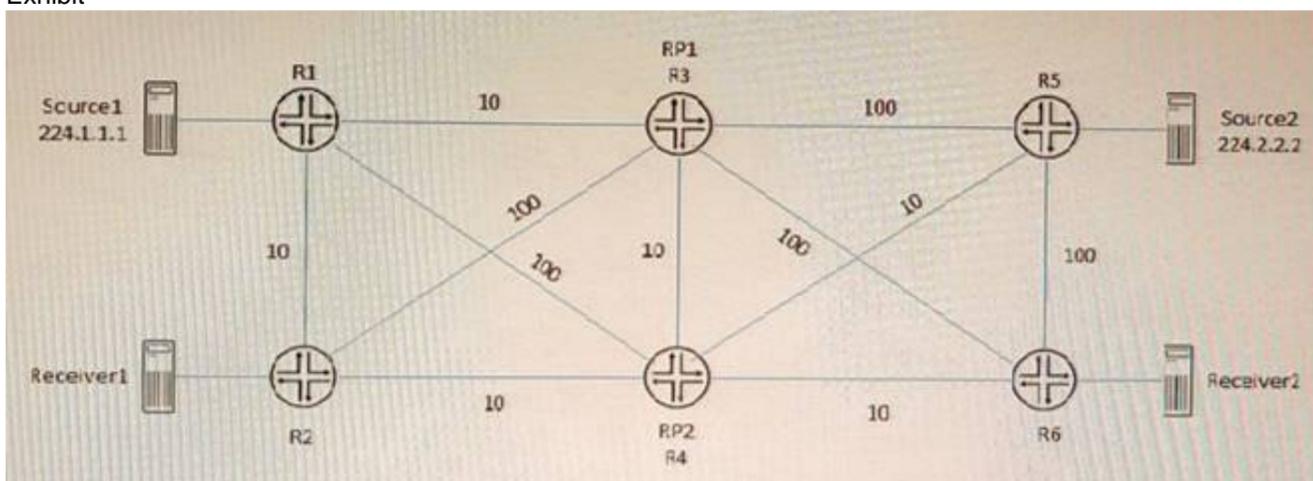
Explanation:

According to 1 and 2, BGP Layer 2 VPNs use BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. BGP uses the Layer 2 VPN (L2VPN) Routing Information Base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. The prefix and path information is stored in the L2VPN database, which allows BGP to make decisions about the best path. In the output shown in the exhibit, we can see some information about the L2VPN RIB and the pseudowire state. Based on this information, we can infer the following statements:

- ? The PE is attached to a single local site. This is correct because the output shows only one local site ID (1) under the L2VPN RIB section. A local site ID is a unique identifier for a site within a VPLS domain. If there were multiple local sites attached to the PE, we would see multiple local site IDs with different prefixes.
- ? The connection has not flapped since it was initiated. This is correct because the output shows that the uptime of the pseudowire is equal to its total uptime (1w6d). This means that the pseudowire has been up for one week and six days without any interruption or flap.
- ? There has been a VLAN ID mismatch. This is not correct because the output shows that the remote and local VLAN IDs are both 0 under the pseudowire state section. A VLAN ID mismatch occurs when the remote and local VLAN IDs are different, which can cause traffic loss or misdelivery. If there was a VLAN ID mismatch, we would see different values for the remote and local VLAN IDs.
- ? The PE router has the capability to pop flow labels. This is correct because the output shows that the flow label pop bit is set under the pseudowire state section. The flow label pop bit indicates that the PE router can pop (remove) the MPLS flow label from the packet before forwarding it to the CE device. The flow label is an optional MPLS label that can be used for load balancing or traffic engineering purposes.

NEW QUESTION 30

Exhibit



Referring to the exhibit, PIM-SM is configured on all routers, and Anycast-RP with Anycast- PIM is used for the discovery mechanism on RP1 and RP2. The interface metric values are shown for the OSPF area.

In this scenario, which two statements are correct about which RP is used? (Choose two.)

- A. Source2 will use RP2 and Receiver will use RP2 for group 224.2.2.2.

- B. Source2 will use RP1 and Receiver2 will use RP1 for group 224.2.2.2.
- C. Source1 will use RP1 and Receiver1 will use RP1 for group 224.1.1.1.
- D. Source1 will use RP1 and Receiver1 will use RP2 for group 224.1 1 1

Answer: AC

Explanation:

A sham link is a logical link between two PE routers that belong to the same OSPF area but are connected through an L3VPN. A sham link makes the PE routers appear as if they are directly connected, and prevents OSPF from preferring an intra-area back door link over the VPN backbone. A sham link creates an OSPF multihop neighborship between the PE routers using TCP port 646. The PEs exchange Type 1 OSPF LSAs instead of Type 3 OSPF LSAs for the L3VPN routes, which allows OSPF to use the correct metric for route selection.

NEW QUESTION 35

Exhibit

```

user@router> show route advertising-protocol bgp 10.0.0.43 extensive 10.0.0.188
inet.0: 23 destinations, 41 routes (23 active, 0 holddown, 0 hidden)
+ 10.0.0.188/32 (2 entries, 1 announced)
  BGP group underlay type External
    AS path: [65189] 65170 65188 I
    
```

Referring to the exhibit, what do the brackets [] in the AS path identify?

- A. They identify the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured
- B. They identify an AS set, which are groups of AS numbers in which the order does not matter
- C. They identify that the autonomous system number is incomplete and awaiting more information from the BGP protocol.
- D. They identify that a BGP confederation is being used to ensure that there are no routing loops.

Answer: B

Explanation:

The brackets [] in the AS path identify an AS set, which are groups of AS numbers in which the order does not matter. An AS set is used when BGP aggregates routes from different ASs into a single prefix. For example, if BGP aggregates routes 10.0.0.0/16 and 10.1.0.0/16 from AS 100 and AS 200, respectively, into a single prefix 10.0.0.0/15, then the AS path for this prefix will be [100 200]. An AS set reduces the length of the AS path and prevents routing loops.

NEW QUESTION 40

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

JN0-664 Practice Exam Features:

- * JN0-664 Questions and Answers Updated Frequently
- * JN0-664 Practice Questions Verified by Expert Senior Certified Staff
- * JN0-664 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * JN0-664 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The JN0-664 Practice Test Here](#)