Q&A!

Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt.
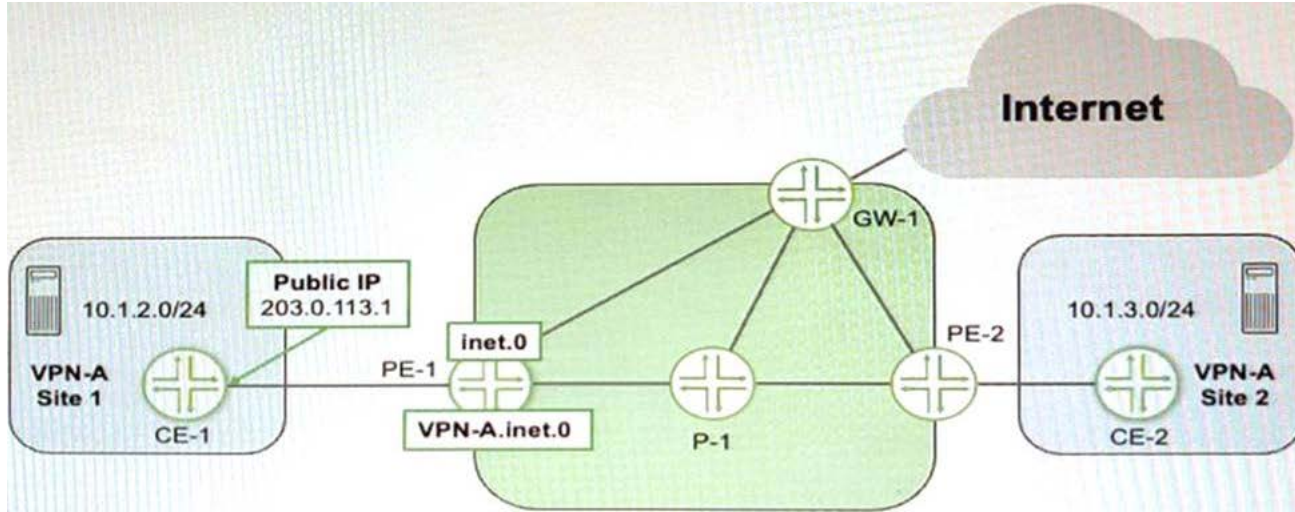----- Guaranteed.

PASS

Juniper

JN0-664

Service Provider Routing and Switching, Professional (JNCIP-SP)

Question: 119

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-
B. inet.0 table as the next hop
C. You must configure a static route in the main routing instance for the 203.0.113.1/32 prefix that uses the VPN-
D. inet.0 table as the next hop.
E. You must configure a RIB group on PE-1 to leak a default route from the inet.0 table to the VPN-
F. inet.0 table.
G. You must configure a RIB group on PE-1 to leak the 10 1 2.0/24 prefix from the VPN-
H. inet.0 table to the inet.0 table.

Answer: A,B

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.

Question: 120

Which three mechanisms are used by Junos platforms to evaluate incoming traffic for CoS purposes? (Choose three)

A. rewrite rules
B. behavior aggregate classifiers
C. traffic shapers
D. fixed classifiers
E. multifield classifiers

Answer: A,B,D,E

Explanation:

Junos platforms use different mechanisms to evaluate incoming traffic for CoS purposes, such as:

â Behavior aggregate classifiers: These classifiers use a single field in a packet header to classify traffic into different forwarding classes and loss priorities based on predefined or user-defined values.

â Fixed classifiers: These classifiers use a fixed field in a packet header to classify traffic into different forwarding classes and loss priorities based on predefined values.

â Multifield classifiers: These classifiers use multiple fields in a packet header to classify traffic into different forwarding classes and loss priorities based on user-defined values and filters.

Rewrite rules and traffic shapers are not used to evaluate incoming traffic for CoS purposes, but rather to modify or shape outgoing traffic based on CoS policies.

Question: 121

Exhibit
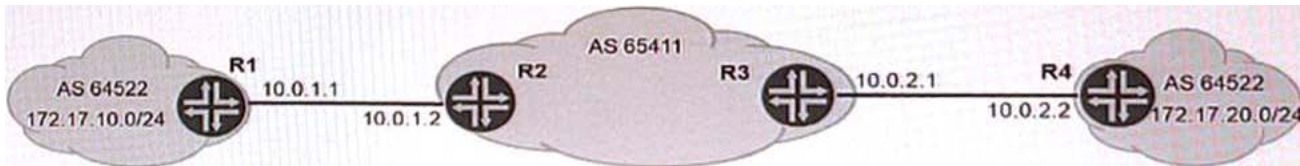
You are asked to exchange routes between R1 and R4 as shown in the exhibit. These two routers use the same AS number.

Which two steps will accomplish this task? (Choose two.)

A. Configure the BGP group with the advertise-peer-as parameter on R1 and R4.
B. Configure the BGP group with the as-override parameter on R2 and R3
C. Configure the BGP group with the advertise-peer-as parameter on R2 and R3.
D. Configure the BGP group with the as-override parameter on R1 and R4

Answer: A,B

Explanation:

The advertise-peer-as parameter allows a router to advertise its peerâs AS number as part of the AS path attribute when sending BGP updates to other peers. This parameter is useful when two routers in the same AS need to exchange routes through another AS, such as in the case of R1 and R4. By configuring this parameter on R1 and R4, they can advertise each otherâs AS number to R2 and R3, respectively.

The as-override parameter allows a router to replace the AS number of its peer with its own AS number when receiving BGP updates from that peer. This parameter is useful when two routers in different ASes need to exchange routes through another AS that has the same AS number as one of them, such as in the case of R2 and R3. By configuring this parameter on R2 and R3, they can override the AS number of R1 and R4 with their own AS number when sending BGP updates to each other.

Question: 122

You want to ensure that L1 IS-IS routers have only the most specific routes available from L2 IS-IS routers.

Which action accomplishes this task?

A. Configure the ignore-attached-bit parameter on all L2 routers.
B. Configure all routers to allow wide metrics.
C. Configure all routers to be L1.
D. Configure the ignore-attached-bit parameter on all L1 routers
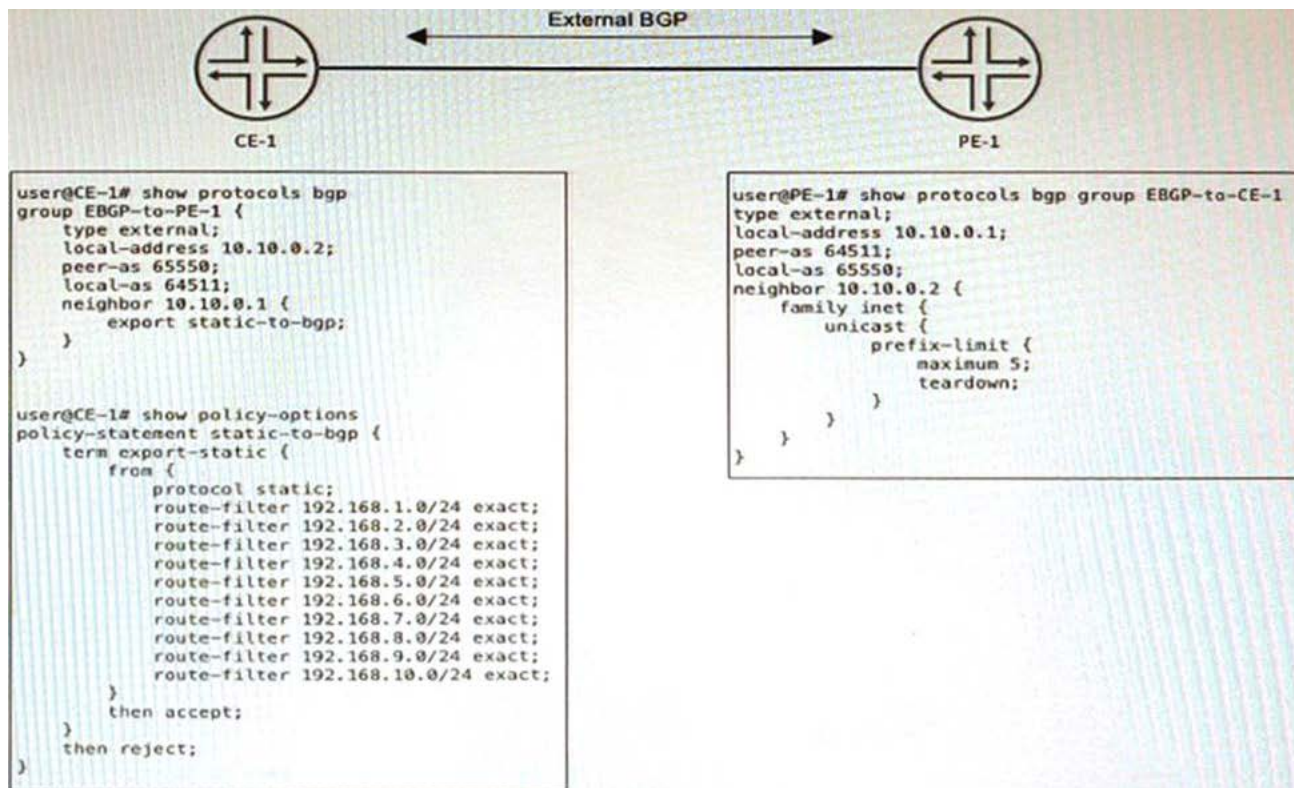
Answer: D

Explanation:

The attached bit is a flag in an IS-IS LSP that indicates whether a router is connected to another area or level (L2) of the network. By default, L2 routers set this bit when they advertise their LSPs to L1 routers, and L1 routers use this bit to select a default route to reach other areas or levels through L2 routers. However, this may result in suboptimal routing if there are multiple L2 routers with different paths to other areas or levels. To ensure that L1 routers have only the most specific routes available from L2 routers, you can configure the ignore-attached-bit parameter on all L1 routers. This makes L1 routers ignore the attached bit and install all interarea routes learned from L2 routers in their routing tables.

Question: 123

Exhibit

External BGP

CE-1   PE-1

```
user@CE-1# show protocols bgp
group EBGP-to-PE-1 {
    type external;
    local-address 10.10.0.2;
    peer-as 65550;
    local-as 64511;
    neighbor 10.10.0.1 {
        export static-to-bgp;
    }
}

user@CE-1# show policy-options
policy-statement static-to-bgp {
    term export-static {
        from {
            protocol static;
            route-filter 192.168.1.0/24 exact;
            route-filter 192.168.2.0/24 exact;
            route-filter 192.168.3.0/24 exact;
            route-filter 192.168.4.0/24 exact;
            route-filter 192.168.5.0/24 exact;
            route-filter 192.168.6.0/24 exact;
            route-filter 192.168.7.0/24 exact;
            route-filter 192.168.8.0/24 exact;
            route-filter 192.168.9.0/24 exact;
            route-filter 192.168.10.0/24 exact;
        }
        then accept;
    }
    then reject;
}
```

```
user@PE-1# show protocols bgp group EBGP-to-CE-1
type external;
local-address 10.10.0.1;
peer-as 64511;
local-as 65550;
neighbor 10.10.0.2 {
    family inet {
        unicast {
            prefix-limit {
                maximum 5;
                teardown;
            }
        }
    }
}
```

CE-1 must advertise ten subnets to PE-1 using BGP Once CE-1 starts advertising the subnets to PE-1, the BGP peering state changes to Active.

Referring to the CLI output shown in the exhibit, which statement is correct?

A. CE-1 is advertising its entire routing table.
B. CE-1 is configured with an incorrect peer AS
C. The prefix limit has been reached on PE-1
D. CE-1 is unreachable

Answer: B

Explanation:

The problem in this scenario is that CE-1 is configured with an incorrect peer AS number for its BGP session with PE-1. The CLI output shows that CE-1 is using AS 65531 as its local AS number and AS 65530 as its peer AS number. However, PE-1 is using AS 65530 as its local AS number and AS 65531 as its peer AS number. This causes a mismatch in the BGP OPEN messages and prevents the BGP session from being established. To solve this problem, CE-1 should configure its peer AS number as 65530 under [edit protocols bgp group external] hierarchy level.

Question: 124

By default, which statement is correct about OSPF summary LSAs?

A. All Type 2 and Type 7 LSAs will be summanzed into a single Type 5 LSA
B. The area-range command must be installed on all routers.
C. Type 3 LSAs are advertised for routes in Type 1 LSAs.
D. The metric associated with a summary route will be equal to the lowest metric associated with an individual

contributing route

Answer: C

Explanation:

OSPF uses different types of LSAs to describe different aspects of the network topology. Type 1 LSAs are also known as router LSAs, and they describe the links and interfaces of a router within an area. Type 3 LSAs are also known as summary LSAs, and they describe routes to networks outside an area but within the same autonomous system (AS). By default, OSPF will summarize routes from Type 1 LSAs into Type 3 LSAs when advertising them across area boundaries.

Question: 125

When building an interprovider VPN, you notice on the PE router that you have hidden routes which are received from your BGP peer with family inet labeled-unica3t configured.

Which parameter must you configure to solve this problem?

A. Under the family inet labeled-unicast hierarchy, add the explicit null parameter.
B. Under the protocols ospf hierarchy, add the traffic-engineering parameter.
C. Under the family inet labeled-unicast hierarchy, add the resolve-vpn parameter.
D. Under the protocols mpls hierarchy, add the traffic-engineering parameter

Answer: C

Explanation:

The resolve-vpn parameter is a BGP option that allows a router to resolve labeled VPN-IPv4 routes using unlabeled IPv4 routes received from another BGP peer with family inet labeled-unicast configured. This option enables interprovider VPNs without requiring MPLS labels between ASBRs or using VRF tables on ASBRs. In this scenario, you need to configure the resolve-vpn parameter under [edit protocols bgp group external family inet labeled-unicast] hierarchy level on both ASBRs.

Question: 126

Which two EVPN route types are used to advertise a multihomed Ethernet segment? (Choose two)

A. Type 1
B. Type 3
C. Type 4
D. Type 2

Answer: A,C

Explanation:

EVPN is a solution that provides Ethernet multipoint services over MPLS networks. EVPN uses BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. EVPN uses different route types to

convey different information in the control plane.

The following are the main EVPN route types:

â Type 1 - Ethernet Auto-Discovery Route: This route type is used for network-wide messaging and discovery of other PE devices that are part of the same EVPN instance. It also carries information about the redundancy mode and load balancing algorithm of the PE devices.

â Type 2 - MAC/IP Advertisement Route: This route type is used for MAC and IP address learning and advertisement between PE devices. It also carries information about the Ethernet segment identifier (ESI) and the label for forwarding traffic to the MAC or IP address.

â Type 3 - Inclusive Multicast Ethernet Tag Route: This route type is used for broadcast, unknown unicast, and multicast (BUM) traffic forwarding. It also carries information about the multicast group and the label for forwarding BUM traffic.

â Type 4 - Ethernet Segment Route: This route type is used for multihoming scenarios, where a CE device is connected to more than one PE device. It also carries information about the ESI and the designated forwarder (DF) election process.

Question: 127

Exhibit
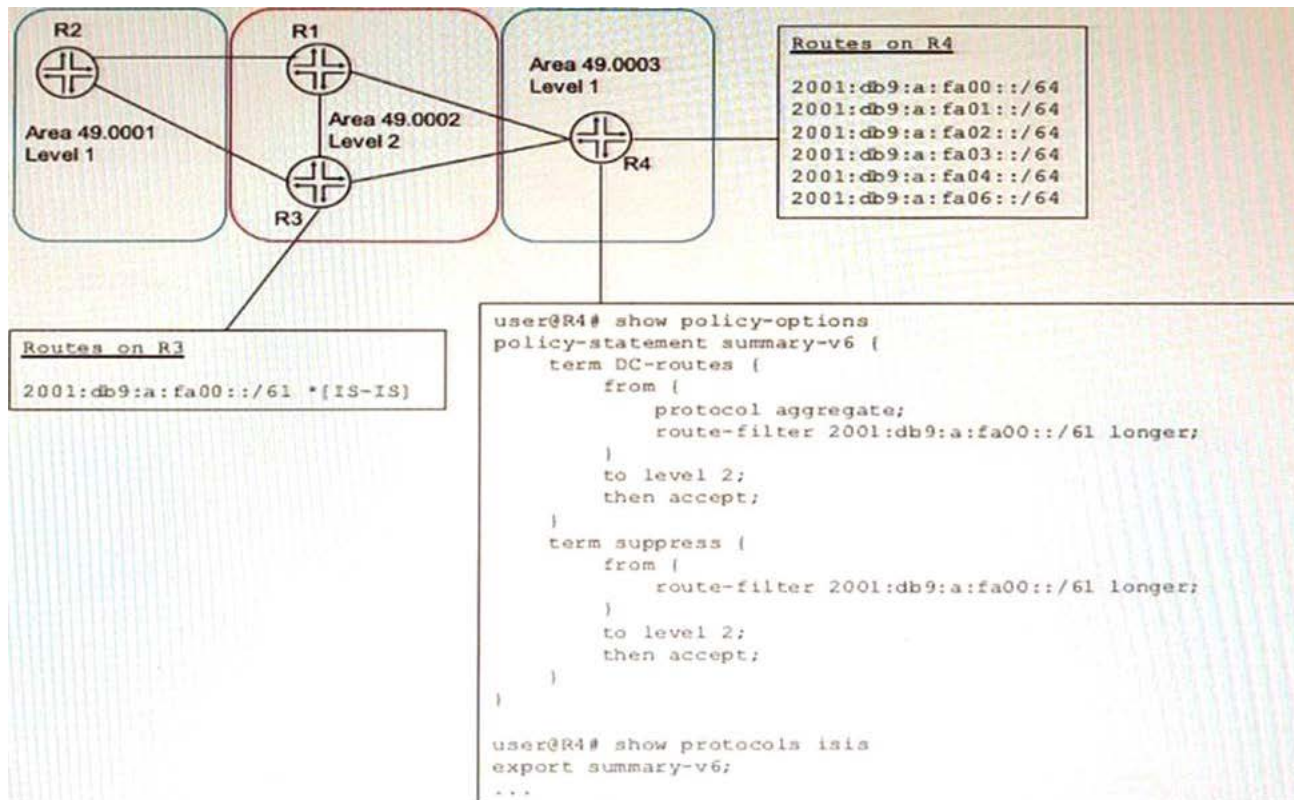
A network designer would like to create a summary route as shown in the exhibit, but the configuration is not working.

Which three configuration changes will create a summary route? (Choose three.)

A. set policy-options policy-statement leak-v6 term DC-routes then reject
B. delete policy-options policy-statement leak-v6 term DC-routes from route-filter 2001: db9:a: fa00 : :/6l longer
C. set policyâoptions policy-statement leak-v term DCâroutes from route-filter 2001:db9:a:faOO::/61 exact
D. delete protocols isis export summary-v6
E. set protocols isis import summary-v6

Answer: A,B,C,D

Explanation:

To create a summary route for IS-IS, you need to configure a policy statement that matches the prefixes to be summarized and sets the next-hop to discard. You also need to configure a summary-address statement under the IS-IS protocol hierarchy that references the policy statement. In this case, the policy statement leak-v6 is trying to match the prefix 2001:db9:a:fa00::/61 exactly, but this prefix is not advertised by any router in the network. Therefore, no summary route is created. To fix this, you need to delete the longer keyword from the route-filter term and change the prefix length to /61 exact. This will match any prefix that falls within the /61 range. You also need to delete the export statement under protocols isis, because this will export all routes that match the policy statement to other IS-IS routers, which is not desired for a summary route.

Question: 128

An interface is configured with a behavior aggregate classifier and a multifield classifier How will the packet be processed when received on this interface?

A. The packet will be discarded.
B. The packet will be processed by the BA classifier first, then the MF classifier.
C. The packet will be forwarded with no classification changes.
D. The packet will be processed by the MF classifier first, then the BA classifier.
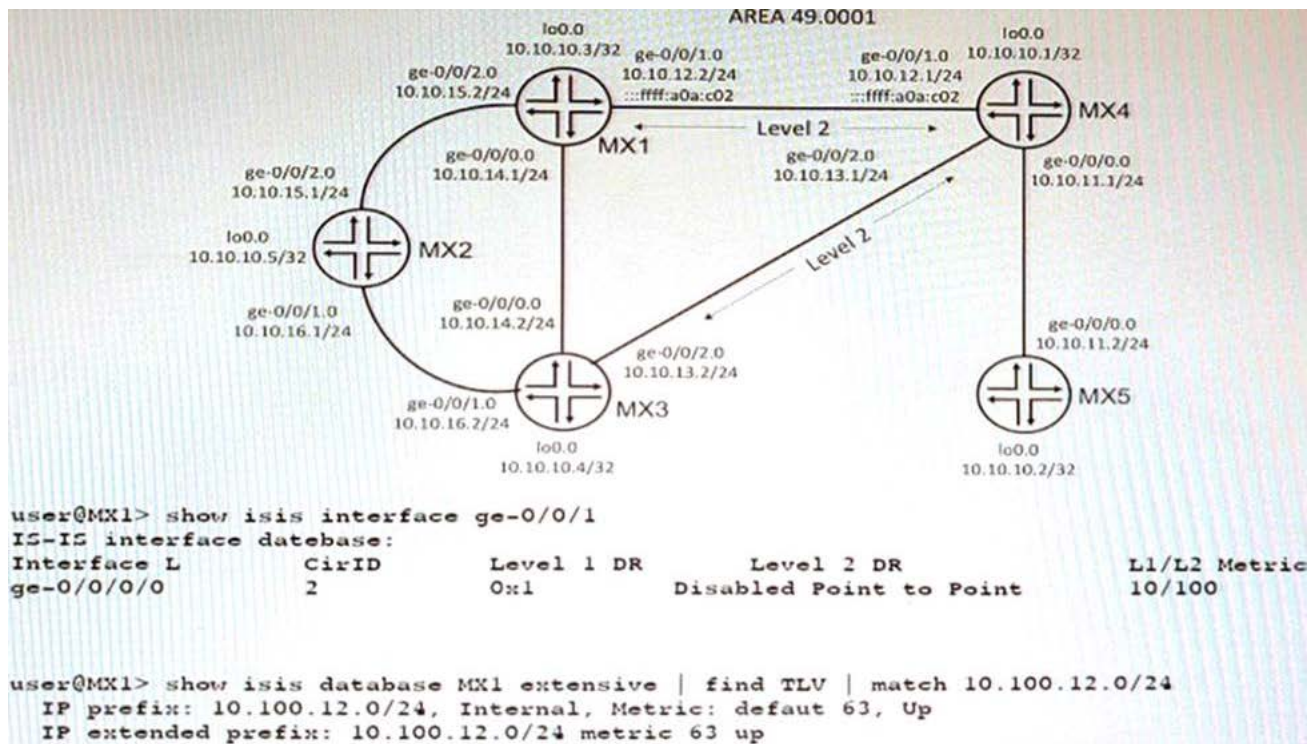
Answer: C

Explanation:

behavior aggregate (BA) classifiers and multifield (MF) classifiers are two types of classifiers that are used to assign packets to a forwarding class and a loss priority based on different criteria. The forwarding class determines the output queue for a packet. The loss priority is used by a scheduler to control packet discard during periods of congestion.

A BA classifier maps packets to a forwarding class and a loss priority based on a fixed-length field in the packet header, such as DSCP, IP precedence, MPLS EXP, or IEEE

Question: 129

Exhibit

AREA 49.0001

lo0.0
10.10.10.3/32 ge-0/0/1.0
ge-0/0/2.0   10.10.12.2/24   ge-0/0/1.0   10.10.10.1/32
10.10.15.2/24   ::ffff:a0a:c02   10.10.12.1/24
                                     ::ffff:a0a:c02   MX4
                          Level 2
ge-0/0/2.0   ge-0/0/0.0   ge-0/0/2.0   ge-0/0/0.0
10.10.15.1/24 10.10.14.1/24   10.10.13.1/24   10.10.11.1/24
             MX1

lo0.0   MX2
10.10.10.5/32

                                              Level 2
ge-0/0/1.0   ge-0/0/0.0
10.10.16.1/24 10.10.14.2/24   ge-0/0/0.0
                        ge-0/0/2.0   10.10.11.2/24
                        10.10.13.2/24
ge-0/0/1.0   MX3   MX5
10.10.16.2/24

lo0.0   lo0.0
10.10.10.4/32   10.10.10.2/32

```
user@MX1> show isis interface ge-0/0/1
IS-IS interface datebase:
Interface L        CirID       Level 1 DR        Level 2 DR            L1/L2 Metric
ge-0/0/0/0         2           0x1               Disabled Point to Point  10/100


user@MX1> show isis database MX1 extensive | find TLV | match 10.100.12.0/24
  IP prefix: 10.100.12.0/24, Internal, Metric: defaut 63, Up
  IP extended prefix: 10.100.12.0/24 metric 63 up
```

A network is using IS-IS for routing.

In this scenario, why are there two TLVs shown in the exhibit?

A. There are both narrow and wide metric devices in the topology
B. The interface specified a metric of 100 for L2.
C. Wide metrics have specifically been requested
D. Both IPv4 and IPv6 are being used in the topology

Answer: A

Explanation:

TLVs are tuples of (Type, Length, Value) that can be advertised in IS-IS packets. TLVs can carry different kinds of information in the Link State Packets (LSPs). IS-IS supports both narrow and wide metrics for link costs. Narrow metrics use a single octet to encode the link cost, while wide metrics use three octets. Narrow metrics have a maximum value of 63, while wide metrics have a maximum value of 16777215. If there are both narrow and wide metric devices in the topology, IS-IS will advertise two TLVs for each link: one with the narrow metric and one with the wide metric. This allows backward compatibility with older devices that only support narrow metrics12.

Question: 130

In IS-IS, which two statements are correct about the designated intermediate system (DIS) on a multi-access network segment? (Choose two)

A. A router with a priority of 10 wins the DIS election over a router with a priority of 1.
B. A router with a priority of 1 wins the DIS election over a router with a priority of 10.
C. On the multi-access network, each router forms an adjacency to every other router on the segment
D. On the multi-access network, each router only forms an adjacency to the DI

Answer: A,D

Explanation:

In IS-IS, a designated intermediate system (DIS) is a router that is elected on a multi-access network segment (such as Ethernet) to perform some functions on behalf of other routers on the same segment. A DIS is responsible for sending network link-state advertisements (LSPs), which describe all the routers attached to the network. These LSPs are flooded throughout a single area. A DIS also generates pseudonode LSPs, which represent the multi-access network as a single node in the link-state database. A DIS election is based on the priority value configured on each routerâs interface connected to the multi-access network. The priority value ranges from 0 to 127, with higher values indicating higher priority. The router with the highest priority becomes the DIS for the area (Level 1, Level 2, or both). If routers have the same priority, then the router with the highest MAC address is elected as the DIS. By default, routers have a priority value of 64. On a multi-access network, each router only forms an adjacency to the DIS, not to every other router on the segment. This reduces the amount of hello packets and LSP

Question: 131

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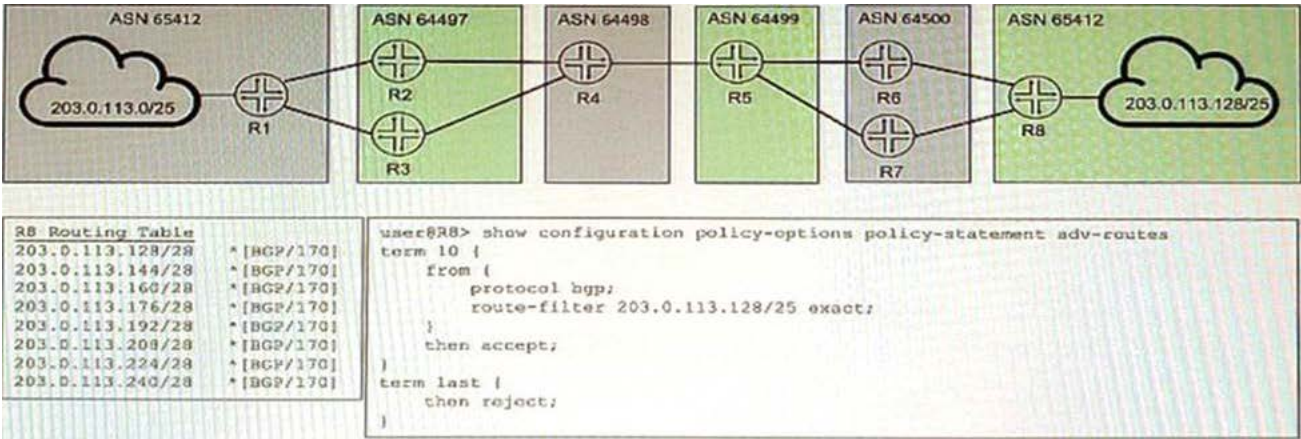
You are attempting to summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500. You implement the export policy shown in the exhibit and all routes from the routing table stop being advertised.

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 tables.
C. Add the set routing-options static route 203.0.113.123/25 discard command.
D. Replace exact in the export policy with orlonger.

Answer: A,C,D

Explanation:

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

â 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.

â 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.

Question: 132

Exhibit
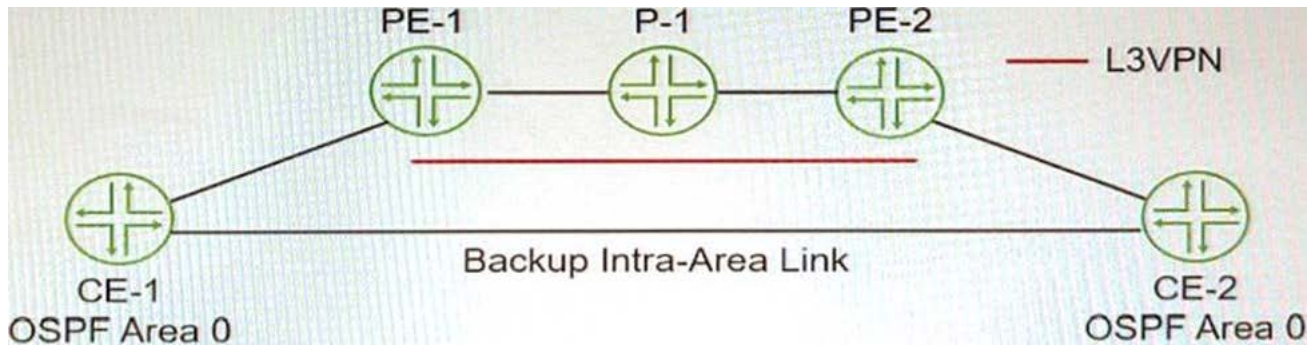
You must ensure that the VPN backbone is preferred over the back door intra-area link as long as the VPN is available. Referring to the exhibit, which action will accomplish this task?

A. Configure an import routing policy on the CE routers that rejects OSPF routes learned on the backup intra-area link.
B. Enable OSPF traffic-engineering.
C. Configure the OSPF metric on the backup intra-area link that is higher than the L3VPN
link.
D. Create an OSPF sham link between the PE routers.

Answer: D

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. To create a sham link, you need to configure the local and remote addresses of the PE routers under the [edit protocols ospf area area-id] hierarchy level1.

Question: 133

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
EM -- encapsulation mismatch         WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down        NP -- interface hardware not present
CM -- control-word mismatch          -> -- only outbound connection is up
CN -- circuit not provisioned        <- -- only inbound connection is up
OR -- out of range                   Up -- operational
OL -- no outgoing label              Dn -- down
LD -- local site signaled down       CF -- call admission control failure
RD -- remote site signaled down      SC -- local and remote site ID collision
LN -- local site not designated      LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status      IL -- no incoming label
MM -- MTU mismatch                   MI -- Mesh-Group ID not available
BK -- Backup connection              ST -- Standby connection
PF -- Profile parse failure          PB -- Profile busy
RS -- remote site standby            SN -- Static Neighbor
LB -- Local site not best-site       RB -- Remote site not best-site
VM -- VLAN ID mismatch               HS -- Hot-standby Connection
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: A,D

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:

â 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.

â 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.

â 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.

â 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.

SAMPLE QUESTIONS

*These questions are for demo purpose only. **Full version** is up to date and contains actual questions and answers.*

Killexams.com is an online platform that offers a wide range of services related to certification exam preparation. The platform provides actual questions, exam dumps, and practice tests to help individuals prepare for various certification exams with confidence. Here are some key features and services offered by Killexams.com:

Actual Exam Questions: Killexams.com provides actual exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these actual questions, candidates can familiarize themselves with the content and format of the real exam.

Exam Dumps: Killexams.com offers exam dumps in PDF format. These dumps contain a comprehensive collection of questions and answers that cover the exam topics. By using these dumps, candidates can enhance their knowledge and improve their chances of success in the certification exam.

Practice Tests: Killexams.com provides practice tests through their desktop VCE exam simulator and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice tests cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

Guaranteed Success: Killexams.com offers a success guarantee with their exam dumps. They claim that by using their materials, candidates will pass their exams on the first attempt or they will refund the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exams.

Updated Content: Killexams.com regularly updates its question bank and exam dumps to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.

Technical Support: Killexams.com provides free 24x7 technical support to assist candidates with any queries or issues they may encounter while using their services. Their certified experts are available to provide guidance and help candidates throughout their exam preparation journey.

For More exams visit https://killexams.com/vendors-exam-list
Kill your exam at First Attempt....Guaranteed!