PassQueen
a.k.a dumpsqueen.com

HCIP-Datacom-Advanced Routing & Switching Technology V1.0

Huawei H12-831_V1-0

Version Demo

Total Demo Questions: 42

Total Premium Questions: 422

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, OSPF	112
Topic 2, IS-IS	49
Topic 3, BGP	46
Topic 4, Routing Policy	12
Topic 5, IPv6	3
Topic 6, Multicast	3
Topic 7, MPLS	45
Topic 8, MPLS VPN	44
Topic 9, EVPN	2
Topic 10, Network Security	19
Topic 11, Network Services and Applications	43
Topic 12, Mix Questions	44
Total	422

QUESTION NO: 1

On the ISIS network shown in the figure, R1 imports a default route using the default-route-advertise always level-1 command. In this case, R3 can learn this default route through IS-IS.

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct because `default-route-advertise always level-1` generates and advertises a default route only in Level-1 IS-IS information. The scope of that advertisement is the Level-1 area in which R1 participates; it is carried in Level-1 LSPs and is available to Level-1 routers in that area. IS-IS maintains separate Level-1 and Level-2 databases and flooding domains. Therefore, a default route originated specifically at Level-1 is not automatically propagated into the Level-2 domain for a router such as R3 outside that Level-1 area. To make a default route available at Level-2, the route must be advertised at Level-2, for example through the appropriate Level-2 default-route advertisement configuration on a suitable router. The `always` keyword ensures that R1 originates the configured default route regardless of whether it has another default route in its routing table; it does not expand the route's Level-1 flooding scope. This behavior follows the IS-IS distinction between Level-1 intra-area routing and Level-2 inter-area routing. See Huawei's VRP configuration documentation and the IS-IS protocol specification in RFC 1195.

QUESTION NO: 2

On the OSPFv3 network shown in the figure, the IPv6 address of Loopback0 on R3 is 2000-3/128, and OSPFv3 is enabled. Which of the following statements are true?

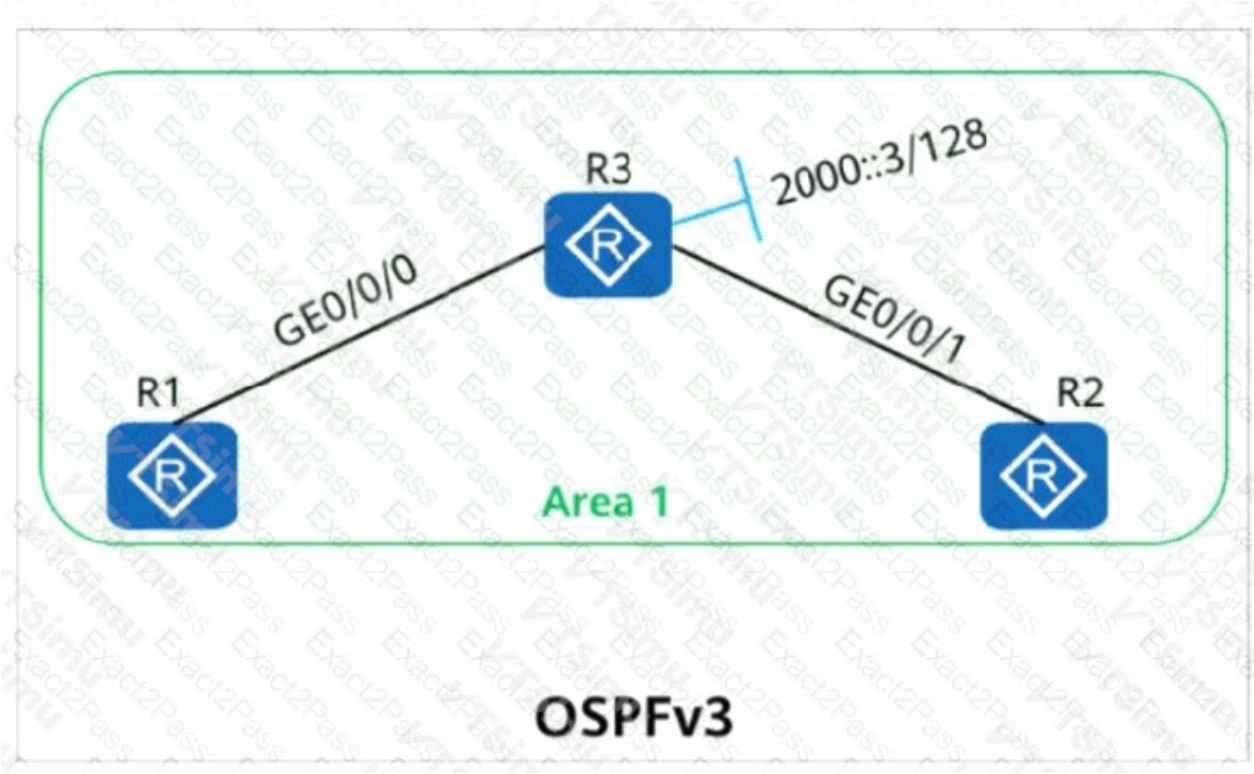

A. R2 definitely generates a Link-LSA.

B. The Link-LSA generated by R2 does not exist in the LSDB of RI.

C. RI definitely generates a Network-LSA.

D. R3 definitely generates an Intra-Area-Prefix-LSA.

ANSWER: A B D

Explanation:

R2 definitely generates a Link-LSA because an OSPFv3 router originates a separate Link-LSA for each OSPFv3-enabled link. This LSA describes link-specific information, including the router's link-local address and the IPv6 prefixes associated with that link. Link-LSAs use link-local flooding scope, so they are retained and used only on the originating router's local link rather than being flooded throughout the area. Consequently, the Link-LSA generated by R2 does not exist in the LSDB of R1; R1 does not receive or store R2's link-scoped LSA.

R3 definitely generates an Intra-Area-Prefix-LSA for its OSPFv3-advertised Loopback0 IPv6 prefix. OSPFv3 separates topology information from IPv6 prefix information: Router-LSAs describe topology, while Intra-Area-Prefix-LSAs carry IPv6 prefixes associated with Router-LSAs or Network-LSAs. A loopback interface is represented as a stub link in the originating router's topology information, and its configured OSPFv3 prefix is advertised through an Intra-Area-Prefix-LSA. These mechanisms enable routers in the same area to calculate routes to R3's loopback address. See RFC 5340, OSPF for IPv6, especially the Link-LSA and Intra-Area-Prefix-LSA definitions.

QUESTION NO: 3 - (SIMULATION)

On the OSPFv3 network shown in the figure, the **LSDB of R2 contains ____ Router-LSAs** . (**Enter only digits.**)

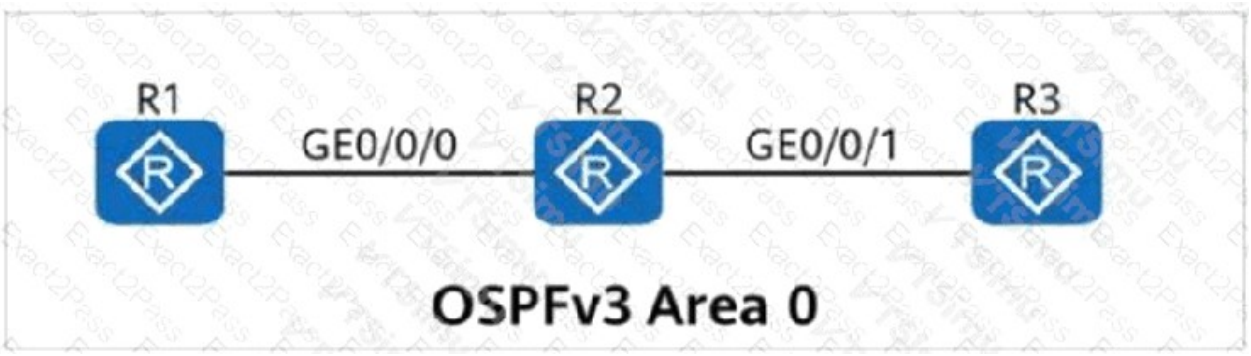

ANSWER: See the explanation for the answer

Explanation:

3

All three routers (R1, R2, and R3) belong to OSPFv3 Area 0. Each router originates one Type 1 Router-LSA for that area, so R2's Area 0 LSDB contains the Router-LSAs originated by R1, R2, and R3.

QUESTION NO: 4

In an operator's MPLSVPN network, there are two devices PE1 and PE2 for MPLSVPN data forwarding. PE1 receives a private network route of 172.16.1.0/24 from the client, and converts it to a VPNv4 route on PE1 and assigns a label of 1027 Released to PE2. The outgoing label of the MPLS LSR-ID of PE2 reaching PE1 is 1025o When a client on PE2 accesses 172.16.1.0/24, the frame sent by PE2, the inner and outer labels should be a combination of the following 6 options?

A. outer layer label:1025:inner label:1027

B. outer layer label:1027:inner label:1027

C. Layer Label:1027:inner label:1025

D. Layer Label,1025:inner label:1025

ANSWER: A

Explanation:

The correct label stack is **outer layer label:1025:inner label:1027**. In an MPLS Layer 3 VPN, the ingress provider edge device uses a two-label stack when forwarding traffic across the provider backbone toward a remote provider edge device. The inner VPN label identifies the destination VPN forwarding context and the egress-side service forwarding action. PE1 allocated VPN label 1027 for the route to 172.16.1.0/24 and advertised that labeled VPNv4 route to PE2; therefore, PE2 pushes label 1027 as the inner label for traffic destined for that private-network prefix.

The outer transport label is used solely to deliver the packet through the MPLS core to PE1. PE2 has learned that the LSP toward PE1 uses outgoing label 1025, so it pushes 1025 above the VPN label. Core provider routers use and swap only this outer transport label while forwarding the packet toward PE1. At PE1, the transport label is removed, and the remaining label 1027 identifies the relevant VPN route or forwarding entry for delivery to the attached customer network. This separation of transport forwarding and VPN/service forwarding is the fundamental MPLS VPN label-stack operation described in Huawei MPLS VPN documentation: Huawei MPLS VPN Configuration Guide and Huawei MPLS Technology Documentation.

QUESTION NO: 5

From the user's point of view, any phenomenon that affects the business can be defined as a fault.

A. True

B. False

ANSWER: A

Explanation:

True is correct because fault management is ultimately business- and service-oriented. From an end user’s perspective, a fault is any condition that degrades, interrupts, or otherwise adversely affects the service they are trying to use. The underlying cause may be a physical device failure, a link problem, a routing convergence event, a configuration error, resource exhaustion, excessive delay, packet loss, or an application-related issue. What makes it a fault from the user viewpoint is the observable impact on normal business operation, rather than whether a network element has generated a specific hardware alarm.

This perspective is important in enterprise network operations because the objective is to maintain the availability and quality of delivered services. Effective monitoring therefore combines device and link indicators with service experience indicators, such as reachability, response time, loss, and transaction success. Huawei’s enterprise networking resources emphasize network operations and maintenance around service assurance, while standard fault-management terminology also treats faults in terms of conditions requiring corrective action to restore expected operation. See Huawei’s Enterprise Support documentation portal and the IETF’s Network Management Datastore Architecture terminology for related network-management concepts.

QUESTION NO: 6

When deploying BGPIMPLS IPVPN, when two VPNs have a common site, the common site must not use overlapping address space with other sites of the two VPNs.

A. True

B. False

ANSWER: A

Explanation:

True is correct. In a BGP/MPLS IP VPN, overlapping private address space can be supported for sites that are isolated in separate VPN instances, because route distinguishers make otherwise identical IPv4 prefixes unique in the provider’s VPN-IPv4 address family. A common site, however, participates in both VPNs and must exchange routes with the other sites belonging to each VPN. Its routes are therefore installed and resolved within the relevant VPN routing and forwarding contexts alongside routes from those other sites. To ensure unambiguous destination selection and correct bidirectional

communication, the common site's prefix space must be unique relative to the other sites that it can reach through both VPNs. This is an important design consideration for an extranet or shared-site topology: route targets control which routes are imported and exported, but they do not make two identical reachable IP prefixes a meaningful, unambiguous addressing design for the common site. Huawei's IP VPN documentation describes how VPN route targets govern route advertisement and import between VPN instances, while RFC 4364 defines the BGP/MPLS VPN model and the role of route distinguishers in making VPN routes unique. See RFC 4364: BGP/MPLS IP Virtual Private Networks and Huawei: Understanding MPLS L3VPN.

QUESTION NO: 7

The **router ID of R4 is 10.0.4.4**. The **LSA shown in the figure is displayed in the LSDB of R4**.

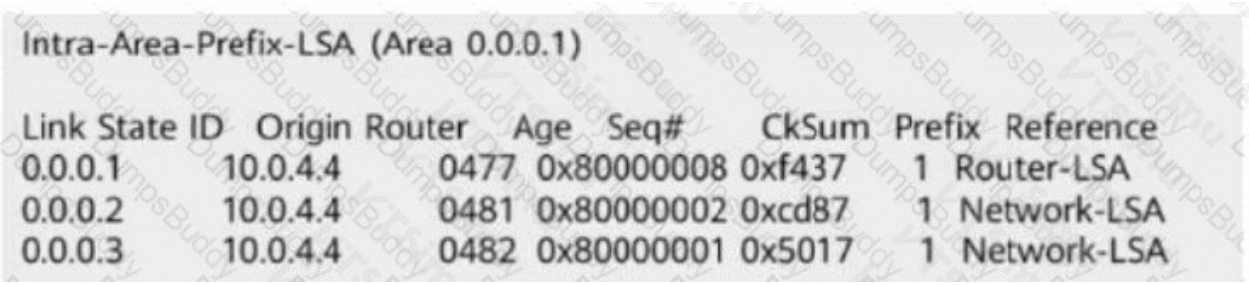

Intra-Area-Prefix-LSA (Area 0.0.0.1)

Link State ID	Origin Router	Age	Seq#	CkSum	Prefix	Reference
0.0.0.1	10.0.4.4	0477	0x80000008	0xf437	1	Router-LSA
0.0.0.2	10.0.4.4	0481	0x80000002	0xcd87	1	Network-LSA
0.0.0.3	10.0.4.4	0482	0x80000001	0x5017	1	Network-LSA

Which of the following statements are **true** about the LSA?

A. The three LSAs belong to area 1.

B. Area 1 is a stub area.

C. No OSPFv3 neighbor relationship is established on at least one link of R4.

D. R4 is the DR on a link.

ANSWER: A C D

Explanation:

The three displayed entries are Intra-Area-Prefix-LSAs in area 0.0.0.1, which is area 1 in dotted-decimal notation. Therefore, the entries belong to area 1. In OSPFv3, an Intra-Area-Prefix-LSA carries IPv6 prefix information and identifies the Router-LSA or Network-LSA to which those prefixes are attached. This relationship permits the topology and prefix information to be evaluated separately.

The references in the displayed LSAs show both router-associated and network-associated prefixes. The router-associated information identifies at least one R4 link that is represented as a stub link rather than as a transit network with an established OSPFv3 neighbor adjacency. Consequently, it is valid to conclude that at least one R4 link does not have an established OSPFv3 neighbor relationship.

In addition, the network-associated prefix information references a Network-LSA advertised by R4, whose router ID is 10.0.4.4. OSPFv3 Network-LSAs are originated by the Designated Router for a broadcast or NBMA transit network. Thus, R4 must be the DR on at least one link. These OSPFv3 LSA functions are specified in RFC 5340, section 4.4.3.9 and section 4.4.3.2.

QUESTION NO: 8

Which of the following statements about 1S-IS (IPv6) is true?

A. to supportIPv6Processing and counting of routesBJS-ISexist129 TLVadded inNLPID

B. IS-ISWorking at the data link layer only needs to add newTLVready to supportPv6

C. to supportIPv6Processing and calculation of routesIS-ISnewly addedTLV232236

D. by defaultJS-ISTurn on the multi-topology feature

ANSWER: A B C

Explanation:

IS-IS was designed to operate independently of a particular network-layer protocol, with its protocol packets carried directly over the data link layer. This architecture allows IPv6 routing support to be added through IS-IS extensions rather than requiring a different routing protocol. IPv6 capability is advertised by including the IPv6 Network Layer Protocol Identifier (NLPID), value 0x8E, in the existing Protocols Supported TLV (TLV 129). This identifies IPv6 as a network-layer protocol supported by the IS-IS speaker.

IPv6 route calculation also relies on the IPv6-specific TLVs defined for IS-IS. TLV 232 carries IPv6 interface addresses, enabling a router to advertise addresses associated with its IS-IS interfaces. TLV 236 carries IPv6 reachability information, including IPv6 prefixes and their associated metrics; these prefixes are used by the SPF calculation to build IPv6 routing entries. Therefore, adding the IPv6 NLPID to TLV 129 and using TLVs 232 and 236 are the essential standardized mechanisms for IPv6 support in Integrated IS-IS. The relevant extension is specified in RFC 5308, while the base Integrated IS-IS framework and Protocols Supported TLV are described in RFC 1195.

QUESTION NO: 9

Which of the following are static information coke analysis?

A. Interface Type

B. Packet loss rate

C. License

D. Equipment type

ANSWER: A C D

Explanation:

Static information analysis uses relatively stable inventory and capability attributes collected from network devices rather than measurements that vary with current traffic conditions. **Interface Type** is static device and port inventory data: it identifies the physical or logical interface form and capability, such as Ethernet interface types, without depending on a particular moment of network operation. **License** is also static information because it records the feature entitlements and capacity capabilities authorized on a device. This information is essential when assessing whether services and functions can be deployed across the network. **Equipment type** is a fundamental static attribute that identifies the device model or category, providing the basis for inventory management, capability matching, and topology-oriented analysis.

These attributes are normally obtained through device discovery, inventory synchronization, and configuration/capability collection. They remain available as descriptive metadata until hardware, interface configuration, or licensing changes, which makes them appropriate inputs to static information analysis. Huawei's enterprise networking portfolio and network-management solutions use this device inventory and capability context to support unified operation and planning. See the Huawei enterprise networking overview and the Huawei iMaster NCE-Campus product page for product context on managed device and network capability information.

QUESTION NO: 10

On the network shown in the figure, **single-hop BFD** is configured on **R1 and R2**.

A network engineer finds that the **BFD session goes down**. To locate the fault, the network engineer queries the **BFD configurations on R1 and R2**.

According to the configuration information marked in the figure, **it can be determined that mismatched time parameter settings on R1 and R2 cause the BFD session to go down**.

Is this statement **TRUE or FALSE**?

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct because single-hop BFD peers do not require their configured timing parameters to be identical in order to establish or maintain a session. BFD communicates each peer's Desired Min TX Interval and Required Min RX Interval in BFD control packets. Each direction then derives its effective control-packet transmission interval by using the applicable advertised constraints from both peers. This negotiation allows routers with different locally configured minimum transmit and receive intervals to interoperate normally.

Therefore, observing different min-tx-interval or min-rx-interval values on R1 and R2 is not, by itself, sufficient evidence that the session went down because of a timer mismatch. The timing values affect the negotiated transmission and detection behavior after the session is operating; they are not a requirement for exact peer-to-peer equality. The BFD base specification defines this exchange and the calculation of the actual transmission interval in its state-machine and timer rules. Huawei VRP BFD configuration follows this standards-based behavior, with local interval settings being advertised to the peer and used in negotiation.

For the protocol timer definitions, see RFC 5880: Bidirectional Forwarding Detection. Huawei product documentation is available through the Huawei Enterprise Documentation portal.

QUESTION NO: 11

Part of the configuration of R1 is as follows. Which statement about IS-IS on R1 is correct?

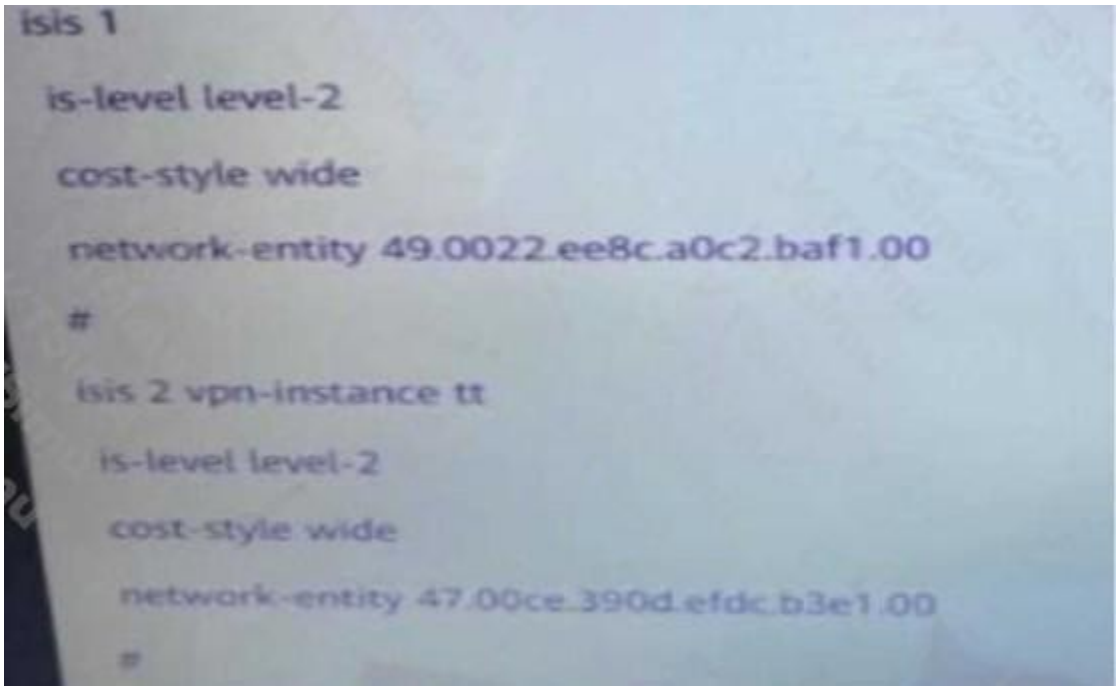

3vpn-instarwce rr

level- 2

cost-style wide

-erntity 47.ccOa 3eft)btxlci-dael

A. Routing information in process 2 and process 3 is isolated.

B. Process 1 contains only public network routing information.

C. Process 1 also contains public network routing information and VRF routing information.

D. Routing information in process 2 and process 3 is shared.

ANSWER: A B

Explanation:

Routing information in process 2 and process 3 is isolated because IS-IS processes bound to distinct VPN instances operate in separate routing and forwarding contexts. Each VPN instance has its own route table and IS-IS link-state database context, so routes learned by one VPN-bound IS-IS process are not automatically visible to another. This separation preserves VPN route isolation even when the processes run on the same physical router.

Process 1 contains only public network routing information because it is the IS-IS process that is not associated with a VPN instance. In Huawei VRP, an IS-IS process configured in the global routing context advertises and calculates routes for the public network routing table. VPN-instance IS-IS processes instead maintain routing information for their respective VPNs. The level-2 setting, wide metric style, and NET identify the IS-IS operational characteristics of a process, but they do not merge the public-network and VPN routing domains. Huawei's IS-IS overview and VPN routing documentation describe the independent routing contexts used for public-network and VPN services: Huawei IS-IS overview and Huawei MPLS L3VPN overview.

QUESTION NO: 12

An administrator configures the following IP prefix list:

```
ip ip-prefix P1 index 10 permit 10.0.0.0 8 greater-equal 16 less-equal 24
```

Which route can be matched by this prefix list entry?

A. 10.0.0.0/8

B. 10.1.0.0/16

C. 10.1.1.0/25

D. 11.1.0.0/16

ANSWER: B

Explanation:

The route **10.1.0.0/16** is matched. The base prefix is 10.0.0.0/8, so a candidate route must be a subnet of 10.0.0.0/8. The `greater-equal 16` and `less-equal 24` parameters further require that the route mask length be from 16 through 24, inclusive.

The route 10.1.0.0/16 satisfies both conditions. The route 10.0.0.0/8 does not satisfy the minimum mask-length condition, 10.1.1.0/25 exceeds the maximum mask length, and 11.1.0.0/16 is not a subnet of 10.0.0.0/8. Huawei IP prefix lists are commonly referenced by routing policies and filter policies to control route advertisement and route acceptance.

QUESTION NO: 13

CORRECT TEXT

As shown in the figure, there is a packet containing a three-layer label header. Please select the values (decimal) corresponding to the x, and Z fields in the figure.

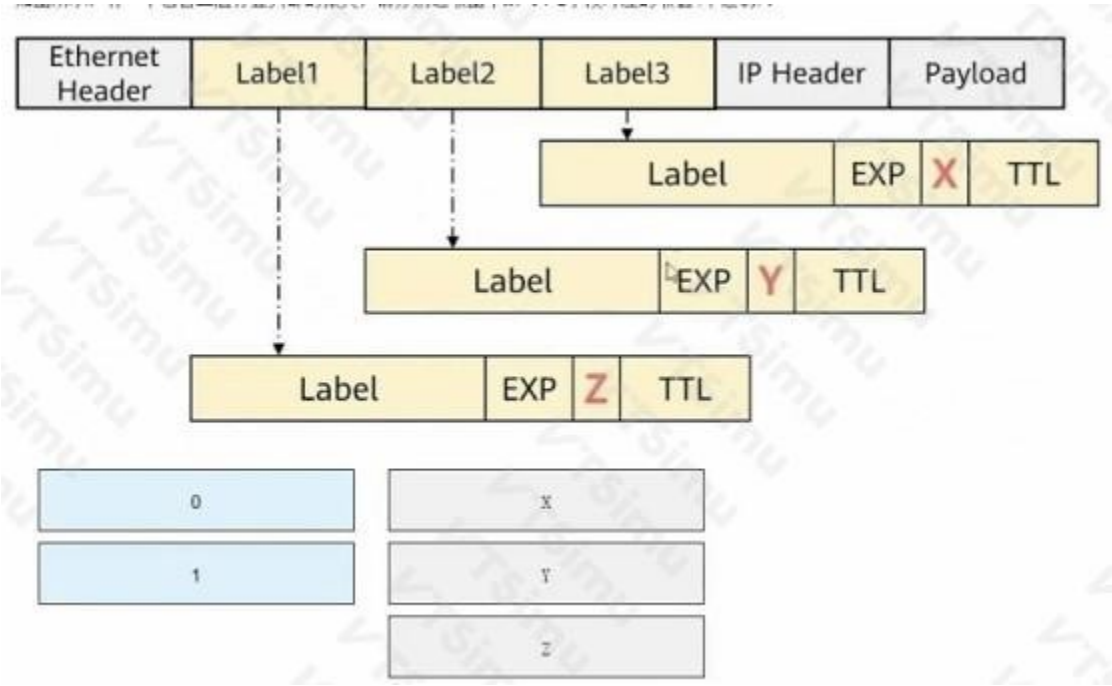

A. ()()()0-Y0-Z1-X

ANSWER: A

Explanation:

The correct values are represented by "()()()0-Y0-Z1-X": X = 1, Y = 0, and Z = 0. An MPLS label stack consists of one or more 32-bit label stack entries. Each entry contains a 20-bit Label Value, 3-bit Traffic Class field, 1-bit Bottom of Stack field, and 8-bit TTL field. The Bottom of Stack field identifies whether the current label is the final label entry in the MPLS stack. Its value is 0 when another label stack entry exists below the current entry, and it is 1 only in the lowest, final label stack entry.

For a three-layer label stack, the two entries that have another label beneath them must carry a Bottom of Stack value of 0. The final entry must carry a Bottom of Stack value of 1. In the illustrated arrangement, Y and Z correspond to entries that are not at the bottom of the stack, while X corresponds to the bottom label entry. Therefore, the required decimal bit values are Y = 0, Z = 0, and X = 1. This setting lets an MPLS forwarding device determine where label processing ends and where the encapsulated payload begins. The MPLS label-stack encoding and the Bottom of Stack semantics are specified in RFC 3032, MPLS Label Stack Encoding.

QUESTION NO: 14

When 0SPFV2 evolves to 0SPFV3, the LSA format and function are exactly the same, but the network layer in the LSA address from IPv4 to IPv6o

A. True

B. False

ANSWER: B

Explanation:

False is correct because OSPFv3 is not simply OSPFv2 with IPv4 addresses replaced by IPv6 addresses. Although OSPFv3 retains the fundamental link-state routing approach and many familiar concepts, its protocol packets and LSAs were redesigned to support IPv6 operation and clearer separation between routing information and address information. For example, OSPFv3 changes the LSA header semantics, uses a 24-bit LSA type field that includes flooding-scope information, and moves IPv6 prefix information out of several topology LSAs into dedicated Link-LSAs and Intra-Area-Prefix-LSAs. Router-LSAs and Network-LSAs therefore describe topology differently from their OSPFv2 counterparts, rather than carrying the same information in an otherwise unchanged format.

OSPFv3 also uses IPv6 link-local addresses for neighbor communication and next-hop handling, while protocol behavior such as authentication is no longer defined in the OSPF packet header as it was in OSPFv2. These architectural changes mean that both the format and functions of relevant LSAs are not exactly identical across the two versions. RFC 5340 defines the OSPFv3 packet and LSA changes in detail: RFC 5340, OSPF for IPv6. Huawei's routing documentation also describes OSPFv3 as a distinct IPv6-oriented protocol implementation: Huawei: Understanding OSPFv3.

QUESTION NO: 15

On the OSPFv3 network shown in the figure:

Area 1 is a stub area

Area 2 is a common area

Area 3 is an NSSA

The IPv6 address of **Loopback0 on R6 is 2000::6/128** .

The **router ID of each router is 10.0.X.X, where X is the router number** .

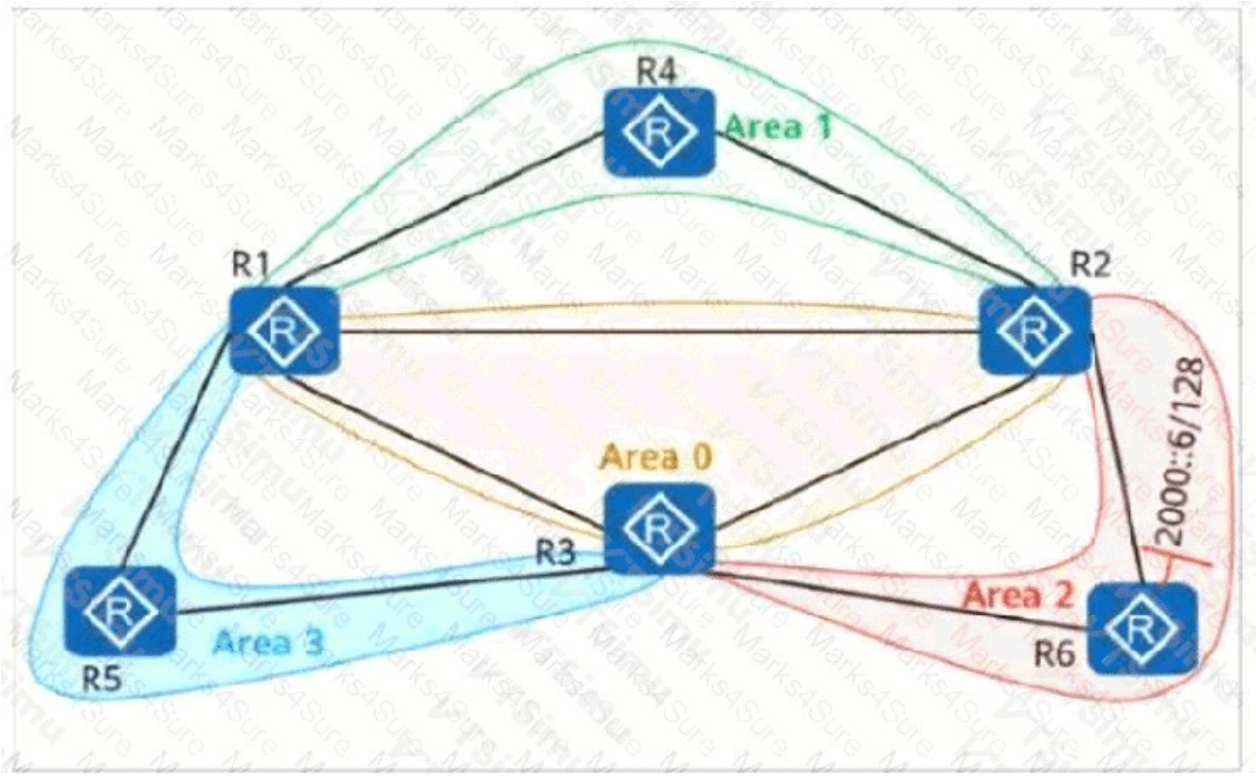

Which of the following statements are true?

A. The **Inter-Area-Prefix-LSA that is generated by R1 and describes 2000::6/128 exists in Area 0** .

B. The **Inter-Area-Prefix-LSA that is generated by R1 and describes 2000::6/128 exists in Area 3** .

C. The **Inter-Area-Prefix-LSA that is generated by R2 and describes 2000::6/128 exists in Area 1** .

D. The **Inter-Area-Prefix-LSA that is generated by R3 and describes 2000::6/128 exists in Area 2** .

ANSWER: A C

Explanation:

The **Inter-Area-Prefix-LSA that is generated by R1 and describes 2000::6/128 exists in Area 0** is correct because R1 is the ABR connecting Area 2, where R6's Loopback0 prefix is learned as an intra-area route, to the backbone. An OSPFv3

ABR originates an Inter-Area-Prefix-LSA (LSA type 3) into another attached area to advertise reachable IPv6 prefixes. Therefore, R1 originates a type 3 LSA for 2000::6/128 in Area 0.

The **Inter-Area-Prefix-LSA that is generated by R2 and describes 2000::6/128 exists in Area 1** is also correct. R2 learns the prefix through the backbone and, as the ABR connected to Area 1, advertises that inter-area destination into Area 1 using an Inter-Area-Prefix-LSA. A stub area does not prohibit inter-area-prefix advertisements. It excludes AS-external information and relies on an ABR-provided default route for external destinations, while internal and inter-area reachability information can still be advertised using type 3 LSAs. OSPFv3 defines Inter-Area-Prefix-LSAs and their ABR origination behavior in RFC 5340, section 3.4.3; the stub-area handling of summary information follows the OSPF area model described in RFC 2328, section 3.6.

QUESTION NO: 16

Which of the following statements about IPv6 Neighbor Discovery are correct?

A. Neighbor Solicitation and Neighbor Advertisement messages can be used for address resolution.

B. Duplicate Address Detection uses Neighbor Discovery messages.

C. IPv6 Neighbor Discovery provides router discovery functionality.

D. IPv6 Neighbor Discovery uses broadcast packets in the same way as IPv4 ARP.

ANSWER: A B C

Explanation:

IPv6 Neighbor Discovery uses ICMPv6 messages to perform functions that include router discovery, neighbor reachability detection, address resolution, duplicate address detection, and redirection. Unlike IPv4 ARP, IPv6 address resolution uses Neighbor Solicitation and Neighbor Advertisement messages rather than a separate ARP protocol.

Duplicate Address Detection is performed by sending a Neighbor Solicitation for the tentative address. If no conflicting Neighbor Advertisement is received, the address can be assigned to the interface. Neighbor Discovery does not use broadcast packets; IPv6 uses multicast addresses, including solicited-node multicast addresses, for these operations. Neighbor Discovery is specified in RFC 4861 and IPv6 Stateless Address Autoconfiguration is specified in RFC 4862.

QUESTION NO: 17

What check items can be set by IPSG?

A. MA, Caddress

B. Outgoing interface

C. VLANID

D. IPaddress

ANSWER: A C D

Explanation:

Huawei IP Source Guard (IPSG) protects an access network by validating the source information in received user packets against a trusted binding entry, typically learned through DHCP snooping or configured statically. The binding identifies the authorized endpoint using a combination of its IP address, MAC address, VLAN, and the interface on which that endpoint is permitted to appear. Therefore, the check items represented by “MA, Caddress”, “VLANID”, and “IPaddress” are valid IPSG validation attributes. Using these attributes together prevents a host from impersonating another endpoint simply by changing one field, such as its IPv4 address or hardware address. In practical deployment, IPSG is normally enabled on user-facing Layer 2 access ports after the network has reliable DHCP-snooping bindings or equivalent static bindings. This creates a consistent relationship between the user identity attributes and the local access segment, allowing the switch to discard traffic whose claimed source does not match the authorized binding. Huawei’s enterprise documentation portal

provides product configuration guidance and feature references for access-security functions, including source-validation capabilities: Huawei Enterprise Documentation. Huawei also describes campus access and endpoint-security solutions through its official enterprise networking resources: Huawei Campus Network Solutions.

QUESTION NO: 18

BGP4+ carries the next hop address of the IPv6 route through the Next Hop attribute in the

Update packet. (Moved into Attribute 14)

A. True

B. False

ANSWER: B

Explanation:

False is correct. For IPv6 unicast routing in BGP-4+, the next-hop information is encoded in the Multiprotocol Reachable NLRI (MP_REACH_NLRI) path attribute, which is attribute type code 14. This multiprotocol extension was designed so that BGP can advertise reachability information for address families other than IPv4 while carrying an address-family-specific next hop. MP_REACH_NLRI includes the Address Family Identifier, Subsequent Address Family Identifier, next-hop length, next-hop address, and the reachable NLRI. Thus, an IPv6 route's next hop is not conveyed by the original BGP NEXT_HOP path attribute used by conventional IPv4 BGP UPDATE processing; it is carried in the multiprotocol reachability attribute. The phrase "Moved into Attribute 14" identifies the relevant mechanism and confirms that the statement's claim about the NEXT_HOP attribute is incorrect. RFC 4760 defines MP_REACH_NLRI and explicitly specifies that it carries the next-hop information for the advertised multiprotocol routes. Huawei VRP documentation also describes IPv6 BGP route advertisement as using the BGP multiprotocol extensions for IPv6 address-family routing. See RFC 4760: Multiprotocol Extensions for BGP-4 and Huawei VRP Product Documentation.

QUESTION NO: 19

Engineer Guan used two routers to test IPv6 services. By running the BGP4+ module to communicate with the branch, as shown in Juxin, an engineer captured packets to view the Updte packets sent by R1. Which of the following statements is true about the description of the message information?

A. This message describes only oneBGPAttribute: Multi-protocol unreachableNLRIAttributes

B. This message describes the current cancellationIPv6routing

C. The maximum and the preceding length of the routing address described in this packet is:2001:db8:2345:l::l/128

D. The next hop address of the route described in this packet is:2001:db8:2345:l::l\

ANSWER: A B C

Explanation:

The captured BGP UPDATE is an IPv6 route-withdrawal advertisement. Its path-attribute section contains the Multiprotocol Unreachable NLRI attribute, formally named MP_UNREACH_NLRI. This attribute carries the AFI/SAFI and the withdrawn IPv6 NLRI, allowing BGP4+ to withdraw IPv6 unicast reachability independently of the legacy IPv4 Withdrawn Routes field. In the displayed information, the MP_UNREACH_NLRI data identifies the IPv6 unicast address family and contains the withdrawn prefix `2001:db8:2345:1::1/128`; a prefix length of 128 denotes one specific IPv6 host route. Therefore, the message has one listed BGP attribute, describes withdrawal of the current IPv6 route, and identifies that /128 NLRI. A next-hop field is not part of MP_UNREACH_NLRI: next-hop information is associated with reachable multiprotocol NLRI advertisements, whereas a withdrawal only needs to identify the route being removed. This behavior follows the multiprotocol BGP extension defined in RFC 4760, particularly the MP_UNREACH_NLRI attribute format. Huawei's BGP IPv6 documentation also describes BGP4+ as using multiprotocol extensions to exchange IPv6 routing information: Understanding BGP4+.

QUESTION NO: 20

Which of the following sequences can be matched by the regular expression 100.S

A. 100

B. 1000

C. 10000

D. 1001

ANSWER: B D

Explanation:

In Huawei VRP regular-expression matching, the period (`.`) represents any single character. The expression is intended to match the prefix `100` followed by exactly one arbitrary character, with the end-of-string marker ensuring that no further characters occur after that final character. Therefore, `1000` matches because its fourth and final character is `0`, and `1001` matches because its fourth and final character is `1`. Both sequences have the required three-character prefix and exactly one additional character. This syntax is especially relevant in VRP features such as AS-path filtering, where regular expressions are used to select routes according to a defined character pattern. Huawei documentation describes regular-expression-based route filtering and the use of regular expressions in routing-policy configuration. See the Huawei CloudEngine documentation and Huawei's VRP configuration documentation for regular-expression and routing-policy reference material.

QUESTION NO: 21

From this picture, we can judge

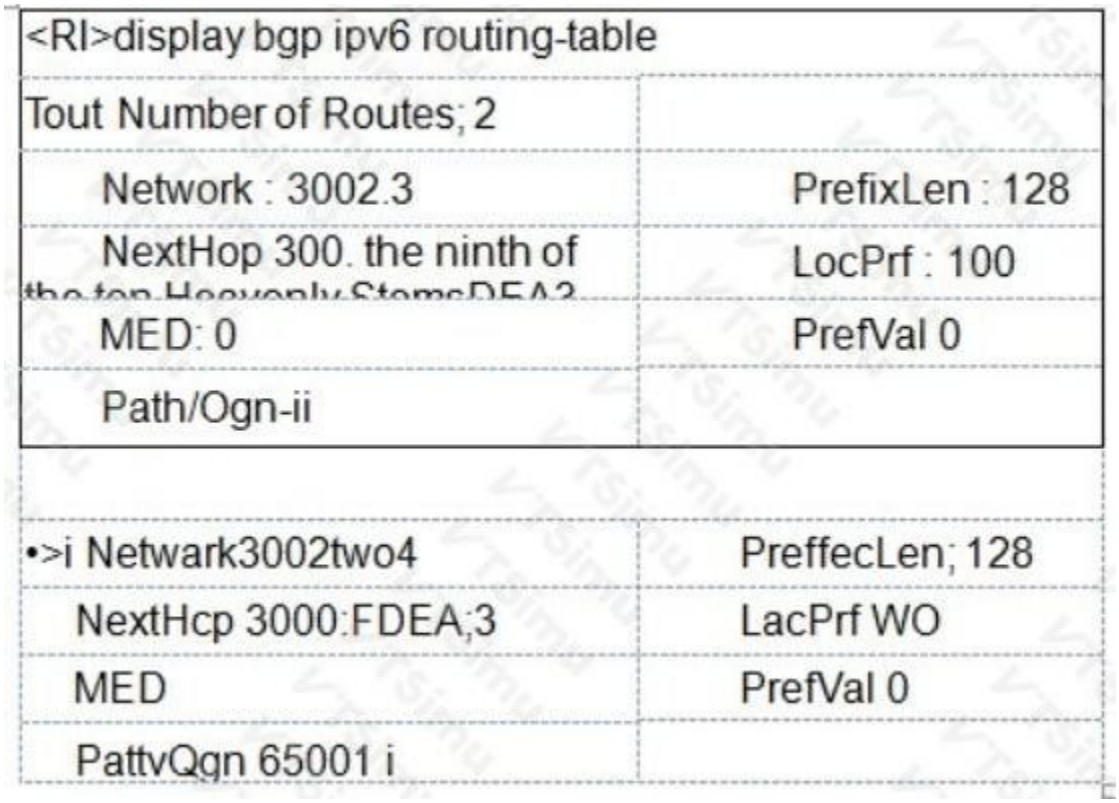

A. If not configured aboutAS_Pathrouting strategy, then3002:-4/128 —must not originate fromAS 65001

B. R1does not have3002:3/128and3002.-4/128routing

C. R1have3002.-3/128and3002:.4/128routing

D. If not configured aboutAS_Pashrouting strategy, then3002:4/28must have originated fromAS 65001

ANSWER: C D

Explanation:

The displayed BGP information shows that R1 has reachability entries for both 3002:3/128 and 3002:4/128. Therefore, these prefixes are present in R1's BGP routing information as indicated by the output. The AS-path information associated with 3002:4/128 also supports the conclusion that, when no AS-Path route-policy is configured to modify or replace AS-path attributes, the route originated in AS 65001. In BGP, the originating autonomous system is represented by the rightmost AS in the AS_PATH attribute; subsequent eBGP advertisements prepend AS numbers to the left of that path. This behavior allows the AS-path shown for a received route to be used to identify its original source AS, assuming no policy has altered the attribute. The relevant BGP AS_PATH processing rules are defined in RFC 4271, Section 5.1.2. Huawei's BGP configuration and route-policy documentation is available through the Huawei Enterprise Documentation portal, including material on AS-path matching and modification.

QUESTION NO: 22

In the MPLS network, there are different operation types for labels, among which the meaning of the "pop" action is:

A. Replace the top label with another value

B. existMPLSAdd the top label to the label stack

C. Replace the top label with another set of labels

D. existMPLSRemove the top label from the label stack

ANSWER: D

Explanation:

In MPLS forwarding, the **pop** operation removes the label currently at the top of the MPLS label stack. MPLS packets can carry one label or multiple stacked labels, and the top label is the one inspected by the current label-switching router to determine its forwarding operation. When that operation is pop, the router strips this outermost label before forwarding the packet. If another MPLS label remains underneath, it becomes the new top label and is processed by the next MPLS hop. If the removed label was the final label, the packet is forwarded as its native payload, such as an IP packet.

This operation is particularly important at an MPLS label-switched path egress, where the forwarding label is no longer needed. It is also used by penultimate-hop popping, in which the router immediately before the egress removes the transport label, reducing processing required at the egress router. Huawei MPLS forwarding behavior follows these standard label-stack operations: push adds label(s), swap replaces the top label, and pop removes the top label. Therefore, **Remove the top label from the label stack** precisely defines the pop action. See the MPLS label-stack operation definitions in RFC 3031 and Huawei's MPLS Configuration Guide.

QUESTION NO: 23

During an attempt to log in to a NetEngine AR6120 through the console port, the network engineer finds that the password is incorrect and the login fails. If a Telnet account can be used to log in to the device and the account has the permission to change the console port password, the network engineer can log in to the device through Telnet to change the console port password.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. On a NetEngine AR6120, the console port is managed through the console user interface and its configured authentication method. When the device remains reachable through Telnet, an authorized administrator can use that remote management session to access the command line and modify the credentials applied to the console login. This recovery approach depends on the Telnet user having sufficient authorization, typically the required privilege level and AAA permissions, to enter the relevant user-interface or local-user configuration view and change the password used by console authentication. Once the revised configuration is committed or saved as appropriate, the engineer can authenticate at the physical console using the new password. This is a standard administrative recovery path: a working management access method can be used to repair another unavailable access method, provided the administrator is explicitly authorized to alter the applicable authentication settings. Huawei documentation describes user-interface authentication configuration and AAA-based local user management in its VRP product documentation. See the Huawei NetEngine AR product documentation and the Huawei VRP configuration documentation.

QUESTION NO: 24

A network engineer obtains a **multicast data frame** whose **destination MAC address** is **01-00-5e-00-01-01**. Given this, which of the following **cannot be the destination IPv4 address** of this frame?

Options:

A. 225.0.1.1

B. 224.0.1.1

C. 239.128.1.1

D. 224.0.0.1

ANSWER: D

Explanation:

224.0.0.1 cannot be the destination IPv4 address because its IPv4 multicast-to-Ethernet mapping produces the destination MAC address **01-00-5e-00-00-01**, not **01-00-5e-00-01-01**. For IPv4 multicast traffic, Ethernet uses the fixed prefix 01-00-5e and copies only the low-order 23 bits of the multicast IPv4 address into the remaining 23 MAC-address bits. The supplied MAC address therefore carries the mapped bit pattern 00-01-01 in its final three octets. In contrast, the address 224.0.0.1 has low-order 23 bits of 00-00-01, so its resulting multicast MAC differs in the fifth octet. The mapping intentionally omits the upper 9 bits of the IPv4 multicast address, meaning multiple multicast group addresses can map to a single Ethernet multicast MAC; however, the retained low-order 23-bit value must still match exactly. This mapping behavior is specified for IP multicast over Ethernet in RFC 1112. See RFC 1112, Host Extensions for IP Multicasting and Huawei's IP Multicast Configuration Guide.

QUESTION NO: 25

Ethernet is a network that supports broadcasting, and once there is a loop in the network, this simple broadcasting mechanism can cause catastrophic consequences. Which of the following phenomena may be caused by loops

A. The device cannot log in remotely

B. pass throughPINGSerious packet loss during network test with command

C. CPUOccupancy exceeds70%

D. use on the deviceDISPLAY INTERFA, CEcommand to view interface statistics, it is found that the interface receives a large number of
broadcast message

ANSWER: A B C D

Explanation:

All listed phenomena may occur when an Ethernet Layer 2 loop is present. Because Ethernet frames do not contain a hop-limit field, broadcast frames and unknown-unicast frames can circulate repeatedly around the loop. Each circulation causes switches to replicate and forward the traffic again, rapidly creating a broadcast storm. Consequently, interface statistics can show a very large number of received broadcast messages, which is a key operational symptom of a loop.

The storm consumes link bandwidth and switch forwarding resources. Legitimate unicast traffic, including ICMP echo traffic used by the PING test, can be delayed or discarded, resulting in severe packet loss. The network may also suffer MAC-address flapping as the same source MAC address is learned on different ports repeatedly. Processing these constant frame arrivals, MAC-table updates, and control-plane events can drive device CPU utilization above 70%, especially on devices where exceptional traffic is punted to the CPU. With bandwidth and device resources exhausted, management protocols such as Telnet or SSH may time out, making remote login unavailable. Loop-prevention mechanisms such as STP/RSTP/MSTP are therefore essential in redundant Layer 2 topologies. Huawei describes STP as a mechanism for eliminating Layer 2 loops and preventing broadcast storms in its STP overview; further configuration guidance is available in Huawei's STP configuration documentation.

QUESTION NO: 26

From this figure, it can be known that

<R1^>tracert 172.1 7.1.5

1 10.1.12.2 40 ms	1 O ms	1 O ms
2 10.1.24.2 30 ms	20 ms	20 ms
3 10/134 1 20 ms	20 ms	20 ms
A 10.1.1 3.1 20 ms	20 ms	1 O ms
5 10/1.12.2 20 ms	30 ms	20 ms
6 10.1.24, 2 30 ms	30 ms	30 ms
10.134-1 50 ms	40 ms	40 ms
8 IO 1.13.1 20 ms	30 ms	30 ms

9 10.1.1 2.2 50 ms 40 ms 40 ms

A. R1access172.17.1.5There is a loop

B. R1have access172.17.1 5routing

C. R1no access172.17.1.5routing

D. R1access172.17.1.5There is no problem with the path

ANSWER: A B

Explanation:

The correct conclusions are “R1access172.17.1.5There is a loop” and “R1have access172.17.1 5routing.” The route-display and forwarding-path information in the figure indicates that R1 has a route entry for the destination network containing 172.17.1.5; therefore, R1 can perform a route lookup and forward packets toward that destination. The existence of a route in the routing table does not, however, guarantee successful end-to-end delivery. The illustrated forwarding behavior shows that traffic is returned through previously traversed routing devices, producing a cyclic next-hop path. This is a Layer 3 routing loop: each participating device has forwarding information, but the combined forwarding decisions fail to move the packet closer to its destination. In operational troubleshooting, this condition is commonly identified by examining route entries and using a traceroute/tracert-style path test, which reveals repeated hops until the packet’s TTL expires. Huawei VRP provides route-display and path-detection commands for validating both the selected route and the actual forwarding path. See Huawei’s VRP command reference and tracert command documentation.

QUESTION NO: 27

CORRECT TEXT

As shown in the figure, OSPF is enabled on all interfaces of the router, and the ip address identified in the figure is the device's IP address.

Loopback0 presses the IP address of the port, RI, R2, R3EJLoopbacD is advertised in the area 1, R4

LoopbacfkD advertises in area D and which of the following functional addresses can be pinged between the LoopbackD announcements of R5 and R5 in area 2?

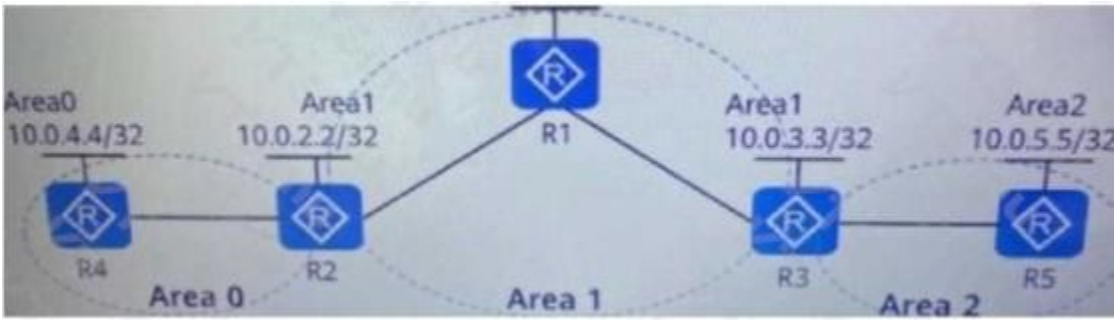

Areal

10.0.1.1/32

A. 10.0.3.3 10.0.5.5

B. 10.0.4.4and10.0.2.2

C. 10.0.2.2and10.0.3.3

D. 10.0.2.2 10.0.5.5

ANSWER: C

Explanation:

10.0.2.2and10.0.3.3 is correct because these are the Loopback0 addresses of R2 and R3, and both loopback routes are advertised within OSPF area 1. Routers in the same OSPF area exchange Type 1 Router LSAs, construct an identical area link-state database after convergence, and run SPF to calculate intra-area routes. Therefore, R2 learns a route to R3's 10.0.3.3/32 loopback and R3 learns a route to R2's 10.0.2.2/32 loopback. With OSPF enabled on the relevant interfaces, the underlying paths in area 1 provide reachability to each loopback address, so ICMP echo packets can be routed in both directions.

This result relies on normal OSPF intra-area operation and does not require inter-area route propagation through an area border router. OSPF treats a loopback prefix advertised by its attached router as a reachable host route, provided the area has converged and the forwarding path to the advertising router is available. The OSPF architecture and intra-area SPF behavior are defined in RFC 2328. Huawei's OSPF configuration documentation is available through the Huawei Enterprise documentation portal.

QUESTION NO: 28

Which of the following statements about OSPF DR and BDR election on a broadcast network are correct?

A. An interface with OSPF priority 0 cannot become a DR or BDR.

B. The DR originates the Network-LSA for a broadcast transit network.

C. All routers on a broadcast network must form Full adjacencies with each other.

D. The BDR is elected only after the DR fails.

ANSWER: A B

Explanation:

On an OSPF broadcast network, the Designated Router and Backup Designated Router reduce the number of full adjacencies required between routers. Routers other than the DR and BDR normally form full adjacencies with the DR and BDR, while the DR originates the Network-LSA for the transit network.

OSPF interface priority participates in DR and BDR election. An interface with priority 0 is not eligible to become either the DR or the BDR. If priorities are equal, the router ID is used as the tie-breaker. OSPF DR and BDR election rules are specified in RFC 2328.

QUESTION NO: 29

On the OSPF network shown in the figure, areas 1, 2, and 3 are common areas. The IP address of **Loopback0 on R5 is 10.0.5.5/32** , and OSPF is enabled on the interface using the **network** command.

ACL 2000 (with the matching rule shown in the figure) is configured on all routers. If the command **filter 2000 import** is run in **area 0** (where R1, R2, and R3 reside), which of the following routers do **not** have the route **10.0.5.5/32** in their routing tables?

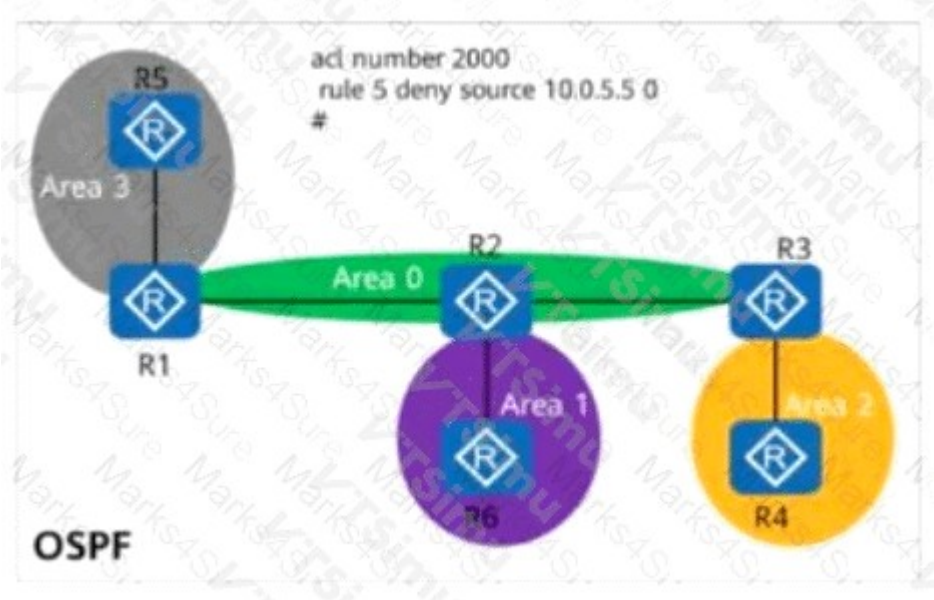

A. R1

B. R3

C. R4

D. R2

ANSWER: A B C D

Explanation:

`filter 2000 import` under the backbone-area view applies an inter-area route filter at the area boundary. The access-control list rule for 10.0.5.5/32 prevents that prefix, which originates from R5's Loopback0 in a non-backbone area, from being imported into area 0 as an inter-area route. Consequently, the three routers located only in the backbone do not install this prefix through backbone flooding and inter-area route calculation.

The route also cannot be advertised onward from the backbone into the separate non-backbone area containing R4. Inter-area OSPF reachability between distinct ordinary areas must traverse the backbone. Since the prefix is stopped when it is being imported into the backbone, there is no eligible inter-area advertisement for R4 to use. Therefore, R1, R3, R4, and R2 do not have 10.0.5.5/32 in their routing tables. R5 retains its directly connected Loopback0 route locally; that local route is not dependent on an inter-area import.

This behavior follows OSPF's backbone-based inter-area routing model, specified in RFC 2328. Huawei's documentation portal for OSPF configuration and route-filter behavior is available at Huawei Enterprise Support Documentation.

QUESTION NO: 30

Service routes of an enterprise are transmitted on the network. Some configurations of PE1 and PE2 are shown in the figure. Which of the following LSAs may CE2 receive?

A. Type 5 LSA

B. Type 1 LSA

C. Type 3 LSA

D. Type 7 LSA

ANSWER: B C

Explanation:

CE2 may receive a **Type 1 LSA** and a **Type 3 LSA**. In an OSPF VPN deployment, the PE participates in the OSPF process facing its attached CE. On the CE2-facing OSPF link, PE2 originates a router-LSA to describe its own router interface and the links belonging to that area. This router-LSA is a Type 1 LSA, so CE2 can learn it as part of the area's normal link-state database synchronization. OSPF routes learned from the remote CE site are carried through the VPN and then advertised by the destination PE into the CE-facing OSPF domain. When the route is presented as inter-area reachability, PE2 originates a summary-LSA; a Type 3 LSA advertises the destination prefix and its cost into CE2's area. Thus, CE2 can receive both the local-area topology information represented by Type 1 LSAs and remote-area route reachability represented by Type 3 LSAs. The OSPF LSA type definitions and the function of router-LSAs and summary-LSAs are specified in RFC 2328, section 12.4.1 and section 12.4.3.

QUESTION NO: 31

In the **hub-spoke networking** shown in the figure, the **Hub-PE can be configured to allow repetitive local AS numbers for correct route transmission** .

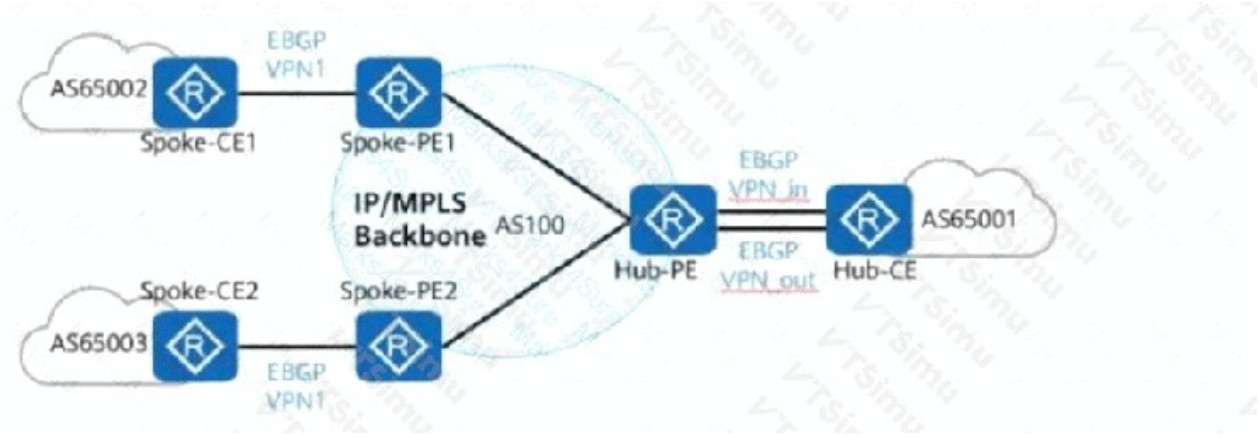

Is this statement TRUE or FALSE?

Options:

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. BGP normally performs AS-path loop detection: when a router receives an update whose AS_PATH already contains its own local AS number, it rejects that update. In a hub-spoke design, however, valid traffic paths can require a route to traverse the hub and be advertised in a manner that includes a repeated customer/local AS. The Hub-PE can use Huawei BGP's `allow-as-loop` capability for the applicable peer or address family to permit reception of routes

containing the local AS a specified number of times. This removes the otherwise intentional AS-path rejection for the defined hub-spoke routing case and allows the required VPN route propagation to continue.

The feature is controlled rather than unrestricted: the permitted repetition count can be configured, allowing the design to accommodate the expected path while retaining AS-path-based loop protection beyond that limit. This is useful where the topology and AS-numbering plan legitimately cause an AS to reappear during route exchange through the hub. The underlying BGP AS-path loop-detection behavior is specified in RFC 4271. Huawei's Enterprise support documentation, including BGP configuration references and command documentation, is available through the Huawei Enterprise Documentation portal.

QUESTION NO: 32

Which of the following statements about LDP Label Retention - Free Way is true:

A. Requires more memory and tab space

B. Keep all labels distributed by neighbors

C. Only keep labels from next-hop neighbors and discard all labels from next-hop neighbors

D. Save memory and label space

ANSWER: A B

Explanation:

"Free Way" label retention in this context refers to liberal label retention. With liberal retention, an LSR keeps label mappings received from all LDP peers for a given FEC, including mappings advertised by peers that are not currently the routing next hop. Therefore, "Requires more memory and tab space" is correct: retaining these additional mappings consumes more control-plane memory and LFIB/label-related table resources than retaining only currently useful mappings. "Keep all labels distributed by neighbors" is also correct because the defining behavior is to preserve label advertisements from all neighbors rather than immediately discard mappings that are not associated with the selected next hop. This approach can speed up forwarding convergence after an IGP route change, since a usable label mapping may already be available when a different neighbor becomes the next hop. The LDP specification describes liberal label retention as retaining label mappings received from peers regardless of whether they are currently used for forwarding. See the IETF's LDP Specification, label retention modes and Huawei's enterprise product documentation portal.

QUESTION NO: 33

If the display ospfv3 peer verbose command is run to check OSPFv3 neighbor information, the command output contains information such as the peer router ID, global unicast address of the peer interface, and neighbor status.

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct because OSPFv3 neighbor communication on a link uses IPv6 link-local addresses, not the peer interface's global unicast address. The detailed neighbor information shown by `display ospfv3 peer verbose` includes peer identification and adjacency details, such as the neighbor router ID, neighbor state, interface-related information, and the neighbor's IPv6 address used for OSPFv3 communication. That address is the link-local address associated with the interface on the shared link. This behavior is fundamental to OSPFv3 operation: routers send OSPFv3 Hello packets using link-local IPv6 source addresses, and neighbors use those addresses for communication across the local link. A globally routable IPv6 unicast address may be configured on the interface for ordinary IPv6 forwarding, but it is not the address type relied on for establishing or maintaining an OSPFv3 adjacency and therefore should not be described as the peer address reported for this purpose. RFC 5340 specifies OSPFv3's use of link-local addresses for neighbor interactions and packet

exchange on an interface; see RFC 5340: OSPF for IPv6. Huawei's OSPFv3 configuration and operation documentation is available through the Huawei Enterprise documentation portal.

QUESTION NO: 34

In what ways can log information be viewed:

A. Network management system interface view

B. Log in to the device and run `display logbuffer`.

C. there is no way

D. Transfer log files in `flash:/logfile/` to a PC through FTP/TFTP.

ANSWER: A B D

Explanation:

Huawei VRP provides several practical paths for accessing log information, supporting both immediate troubleshooting and centralized operations. A network management system can receive and present device log messages through its management interface, allowing administrators to monitor events across multiple devices from a central location. On the device itself, the `display logbuffer` command displays messages currently retained in the log buffer, which is useful when directly connected through the console, SSH, or Telnet. For offline review, retention, or escalation, log files stored in the device's flash file system can be transferred to a PC using FTP or TFTP and then opened or archived there. These methods complement one another: the management system offers centralized visibility, the log buffer supports direct real-time diagnosis, and file transfer enables persistent analysis beyond the device. Huawei's command and configuration documentation describes log buffering, log-file storage, and log output mechanisms; consult the Huawei Enterprise Documentation portal and Huawei Enterprise Support for documentation matching the specific VRP software release.

QUESTION NO: 35

On the **OSPFv3 network** shown in the figure, **area 1 is a stub area**.

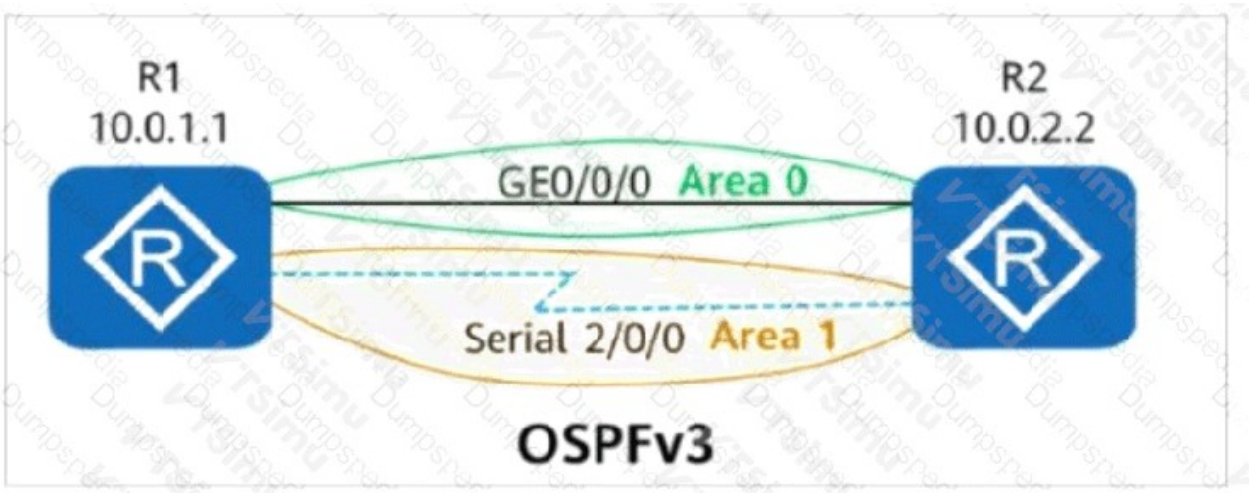

Which of the following **LSAs is generated by only one router**?

A. Inter-Area-Prefix-LSA

B. Link-LSA

C. Network-LSA

D. Intra-Area-Prefix-LSA

ANSWER: C

Explanation:

Network-LSA is correct. In OSPFv3, a Network-LSA is originated for a transit multi-access network, such as an Ethernet broadcast segment, by the designated router (DR) for that segment. The DR represents the shared network in the link-state database and lists the routers that are fully adjacent on it. Because the DR is the sole originator for that particular transit network, exactly one router generates its Network-LSA at a given time. If DR election changes, the newly elected DR originates the relevant LSA, preserving the one-originator rule for the segment.

This behavior reduces duplicate topology advertisements for multi-access networks and allows all routers in the area to derive a consistent view of the network topology. The designation of area 1 as a stub area affects the inter-area and external routing information permitted into that area, but it does not alter the fundamental OSPFv3 rule governing Network-LSA origination on a multi-access segment. RFC 5340 specifies that a router originates a Network-LSA when it is the designated router for the link; see RFC 5340, Section 4.4.2. The associated DR function is described in RFC 5340, Section 7.1.

QUESTION NO: 36

Network administrator A wants to use the AS-Path Fiter to match the routing entries of the BGP routing attribute AS_PATH [100 200 300]. Network administrator A sorts out the four configurations in the figure. Which of the following configurations can satisfy network administrator A's requirements? Require?

	ip as-path-filter TEST permit 100 200 300	
	lip as-path-filter TEST permit 300$	
\C	ip as-path«filter TEST permit A[A100]	
D one	ip as-path-filter TEST permit *.	1

A. EC

B. WayA C. *D

C. WayB

ANSWER: B C

Explanation:

“WayA C. *D” and “WayB” are correct because an AS-path filter uses a regular expression to evaluate the textual AS_PATH representation. For the path `100 200 300`, a regular expression can match the complete path by anchoring it with `^` and `$`, or it can match the required AS sequence by using a pattern that accepts the AS-number separators in the path. Huawei’s BGP AS-path filtering supports regular-expression matching and uses boundary-aware expressions when matching complete AS numbers, so the relevant configurations shown in the figure match the sequence of AS numbers 100, 200, and 300 in the specified order. The AS-path filter is then referenced by routing policy or BGP route filtering to select the matching routes. This behavior allows an administrator to filter paths according to either the entire AS path or the required ordered AS sequence, depending on the expression used. See Huawei’s AS-path filter documentation search and the Huawei BGP configuration documentation for AS_PATH filtering concepts and regular-expression usage.

QUESTION NO: 37 - (SIMULATION)

On the **OSPF network** shown in the figure, drag the values on the left to the correct positions so that:

Traffic from PC1 to PC2 follows the path:PC1 -> R1 -> R3 -> R2 -> PC2

Traffic from PC2 to PC1 follows the path:PC2 -> R2 -> R1 -> PC1

Which of the following **OSPF cost values** should be assigned to **Cost-A, Cost-B, and Cost-C** to achieve this routing behavior?

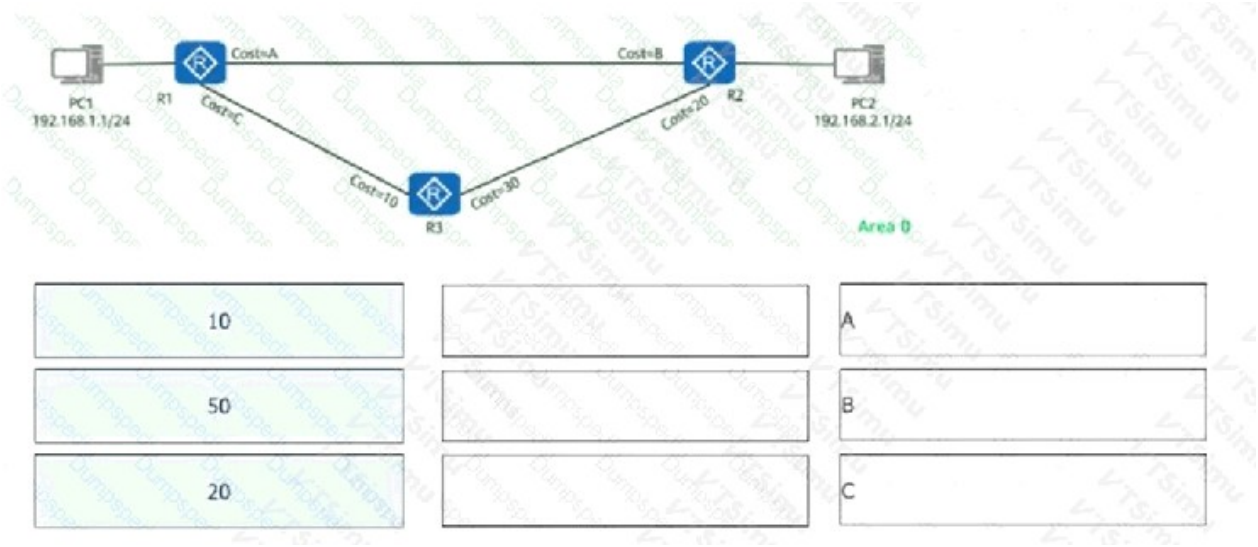

ANSWER: See the explanation for the answer

Explanation:

Assign the values as follows:

Verification:

This produces the required asymmetric routing: `PC1 → R1 → R3 → R2 → PC2` and `PC2 → R2 → R1 → PC1`.

QUESTION NO: 38

OSPFv3 packets are encapsulated in IPv6 packets . Which of the following is the value of the **Next Header** field in the **IPv6 packet header** ?

Options:

A. 6

B. 189

C. 17

D. 89

ANSWER: D

Explanation:

89 is the correct IPv6 Next Header value for an OSPFv3 packet. OSPFv3 is the version of the Open Shortest Path First protocol designed to operate over IPv6. It is carried directly as the payload of an IPv6 packet; it does not use TCP or UDP as a transport-layer wrapper. In the IPv6 base header, the Next Header field identifies the type of header immediately following the IPv6 header. When no IPv6 extension header precedes OSPFv3, this field identifies OSPF directly. The Internet Assigned Numbers Authority assigns protocol number 89 to the Open Shortest Path First Interior Gateway Protocol, and that same assigned value is used in IPv6's Next Header field. RFC 5340, which specifies OSPF for IPv6, explicitly states that OSPF packets are sent as IPv6 packets with protocol number 89. This protocol identifier enables the receiving IPv6 stack to deliver the packet to the OSPFv3 process, where Hello messages, link-state advertisements, database description packets, and other OSPF control packets are processed. See RFC 5340 and the IANA Assigned Internet Protocol Numbers registry.

QUESTION NO: 39

When troubleshooting a network, if you need to confirm the business traffic path, which of the following may need to be done ()

A. Confirm data link layer service traffic path

B. Confirm the network layer service traffic path

C. Investigate service traffic path planning in the network design phase

D. Confirm the frequency of failure

ANSWER: A B C

Explanation:

Confirming the end-to-end business traffic path requires validating forwarding at every relevant layer and comparing the observed path with the intended design. **Confirm data link layer service traffic path** is required where traffic depends on Layer 2 forwarding, such as VLAN membership, MAC-address learning, trunk links, link aggregation, or an Ethernet switching domain. This establishes how frames reach the Layer 3 gateway or the destination within the same broadcast domain.

Confirm the network layer service traffic path is also required because IP traffic may traverse multiple routers or Layer 3 switches. The effective route, next hop, routing-table selection, and any equal-cost paths determine the actual packet path. Tools such as route-table inspection, ping, and tracert/traceroute help validate this forwarding behavior.

Investigate service traffic path planning in the network design phase provides the expected baseline against which the current forwarding path is checked. The design documents identify intended gateways, transit devices, primary and backup paths, and traffic-engineering expectations. Comparing live Layer 2 and Layer 3 forwarding with that baseline is essential for identifying path deviations. Huawei's enterprise documentation portal is available at Huawei Enterprise Support Documentation; IP forwarding and routing fundamentals are also described in RFC 1812.

QUESTION NO: 40

To prevent hackers from attacking user devices or networks through MA, C addresses, you can configure the MA, C addresses of untrusted users as black hole MA, C addresses to filter out illegal MA, C addresses.

A. True

B. False

ANSWER: A

Explanation:

True is correct. A blackhole MAC-address entry is a Layer 2 security mechanism that instructs a switch to discard frames associated with a specified MAC address. Administrators can add the MAC addresses of untrusted or malicious hosts to the blackhole table, preventing traffic sourced from those devices from being forwarded through the switch. This can mitigate attacks in which an attacker uses a known, spoofed, or otherwise unauthorized hardware address to access network resources, send unwanted traffic, or target connected users and services.

The protection is enforced locally on the device holding the blackhole entry: when traffic matching the configured address arrives, the switch drops it rather than learning or forwarding it normally. This makes blackhole MAC-address configuration useful as a rapid containment control when a specific endpoint or address has been identified as unsafe. It should be maintained alongside normal access-control and port-security practices so that the blocked-address list remains accurate and operationally useful. Huawei documentation and support resources for this feature can be found through the Huawei Enterprise blackhole MAC-address documentation search and the Huawei Enterprise Support portal.

QUESTION NO: 41

What is the LSA whose flooding range is the entire autonomous system in SPFV3?

A. Link-LSA

B. As-external-LSA

C. Inter-Area-Prefix-LSA

D. Intra-Area-Prefix-LSA

ANSWER: B

Explanation:

As-external-LSA is correct because, in OSPFv3, an AS-external-LSA carries routes to IPv6 destinations that are external to the OSPF autonomous system, such as routes redistributed from another routing protocol or learned from another source. This LSA has autonomous-system flooding scope, meaning that it is flooded throughout the OSPF routing domain rather than being restricted to the originating link or area. Routers use the external routing information advertised in these LSAs to calculate routes to destinations outside the OSPF domain, with the metric type and forwarding-address information contributing to external route selection where applicable.

OSPFv3 defines the AS-external-LSA as LSA function code 0x4005. The U-bit in this function code specifies that the LSA is flooded with AS scope, while the function-code definition identifies it as an external-route advertisement. This scope enables a normal OSPFv3 ASBR to distribute external reachability consistently across all areas in the autonomous system. The authoritative OSPFv3 LSA-type and flooding-scope definitions are specified in RFC 5340, section 3.4, and the AS-external-LSA format and purpose are described in RFC 5340, section 3.5.5.

QUESTION NO: 42

Run the display ip interface GEO/D/D command to view statistics about ping packets received on the interface.

A. True

B. False

ANSWER: B

Explanation:

False is correct. The `display ip interface GE0/0/0` command displays IPv4 operational and traffic statistics for the specified interface, such as the interface IP address and aggregate IP packet and byte counters. These interface-level counters are not a dedicated measurement of ping traffic. A ping uses ICMP Echo Request and Echo Reply messages, so identifying received ping packets requires statistics that specifically classify ICMP traffic, typically through protocol-specific monitoring, traffic statistics, or packet capture rather than the general IP-interface display. ICMP defines Echo and Echo Reply as distinct message types used by ping, as specified in RFC 792. Huawei's command-reference documentation for VRP can be accessed through the Huawei Enterprise documentation portal; its IP interface display commands provide interface state and aggregate IP traffic information, not a separate received-ping counter. Therefore, running the stated command cannot be used to view statistics specifically for ping packets received on that interface.