PassQueen
a.k.a dumpsqueen.com

HCIP-Datacom-Core Technology V1.0

Huawei H12-821_V1-0

Version Demo

Total Demo Questions: 33

Total Premium Questions: 333

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Datacom Network Fundamentals	53
Topic 2, IP Unicast Routing	176
Topic 3, IP Multicast	27
Topic 4, MPLS	5
Topic 5, Network Security	37
Topic 6, Network Services and Applications	35
Total	333

QUESTION NO: 1

In an IPv4 address space, Class D addresses are used for multicast. Among Class D addresses, which of the following is the permanent group address range reserved for routing protocols?

A. 232.0.0.0 to 232.255.255.255

B. 239.0.0.0 to 239.255.255.255

C. 224.0.0.0 to 224.0.0.255

D. 224.0.1.0 to 231.255.255.255

ANSWER: C

Explanation:

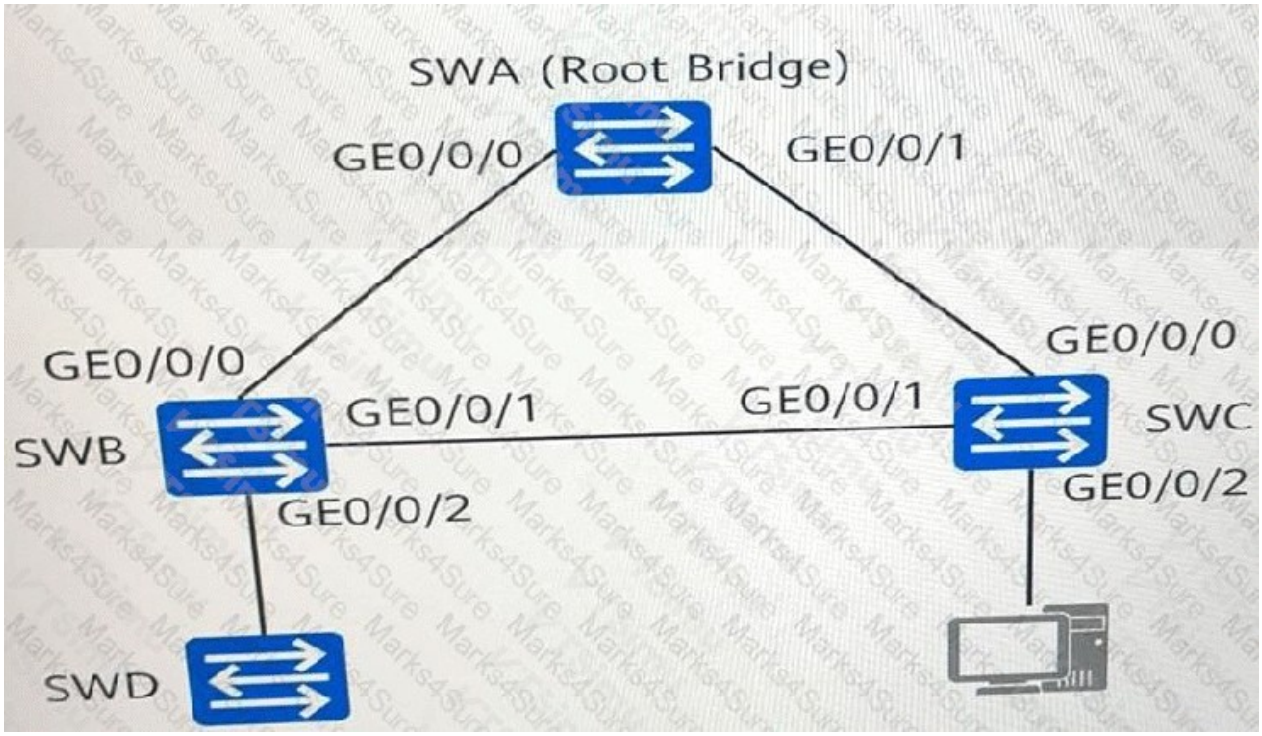
224.0.0.0 to 224.0.0.255 is the IPv4 Local Network Control Block, commonly called the link-local multicast range. These permanent multicast group addresses are reserved for control and routing functions whose traffic is intended to remain on the local network segment. Routers do not forward packets sent to this range, which ensures that protocol-control traffic is constrained to the relevant link.

Several well-known routing protocol multicast groups are allocated from this block. For example, OSPF routers use 224.0.0.5 (AllSPFRouters) and 224.0.0.6 (AllDRouters), while RIPv2 uses 224.0.0.9 (RIP Routers). Using these fixed group addresses allows routers participating in the same protocol to receive routing messages without requiring each device to maintain a list of every peer's unicast address.

The IANA multicast-address registry identifies 224.0.0.0/24 as the Local Network Control Block and lists the protocol-specific assignments within it. RFC 5771 also defines this block as reserved for local network control and specifies that packets in it are not forwarded beyond the local link. See the IANA IPv4 Multicast Address Space Registry and RFC 5771.

QUESTION NO: 2

As shown in the figure, SWA, SWB, SWC, and SWD run the Rapid Spanning Tree Protocol (RSTP). Which of the following statements are true?



A. Once receiving BPDUs, the edge port re-participates in the calculation of the spanning tree.

B. After a port is Configured as an edge port, the port can quickly enter the Forwarding state.

C. You can enable the edge port on SWD ' s GEO/0/2 connected to the terminal so that this port can quickly enter the Forwarding state.

D. You can enable the edge port on SWC ' s GE0/0/2 connected to the terminal so that this port can quickly enter the forwarding state.

ANSWER: A B C D

Explanation:

RSTP edge ports are intended for interfaces that connect directly to terminal devices rather than to other switches. When an interface is configured as an edge port, it bypasses the usual proposal/agreement convergence process and can move to the Forwarding state immediately after the link becomes active. Therefore, configuring the terminal-facing GE0/0/2 interfaces on SWD and SWC as edge ports provides fast connectivity for the attached terminals. This is the standard RSTP access-port deployment practice and reduces the delay that would otherwise occur while spanning-tree states converge.

An interface's edge status is conditional on it behaving as an end-device-facing port. If an edge port receives a BPDU, RSTP recognizes that the port may be connected to another bridge. The port consequently loses its edge-port role and participates in spanning-tree calculation, protecting the Layer 2 topology from a potential loop. Huawei documents the edge-port configuration behavior in its STP configuration documentation; the RSTP protocol behavior is standardized by IEEE 802.1w and incorporated into IEEE 802.1D, summarized in IEEE 802.1D.

QUESTION NO: 3

Which of the following BGP attributes are well-known mandatory attributes?

A. ORIGIN

B. AS_PATH

C. NEXT_HOP

D. LOCAL_PREF

E. MED

ANSWER: A B C

Explanation:

ORIGIN, AS_PATH, and NEXT_HOP are well-known mandatory BGP attributes. A BGP UPDATE message that advertises reachable NLRI must carry these attributes. ORIGIN identifies how the route was introduced into BGP, AS_PATH records the autonomous systems through which the route has passed, and NEXT_HOP identifies the next-hop address used to reach the destination.

LOCAL_PREF is a well-known discretionary attribute rather than a mandatory one. MED is an optional non-transitive attribute. BGP defines these attribute categories and their processing requirements in RFC 4271.

QUESTION NO: 4

Which of the following statements regarding an IP prefix are true?

A. An IP prefix filter is used to filter IP address prefixes and cannot match an IP prefix number and a prefix length at the same time.

B. An IP prefix filter cannot be used to filter data packets.

C. An IP prefix filter is used to filter IP address prefixes and can match an IP prefix number and a prefix length at the same time.

D. An IP prefix filter can be used to filter data packets.

ANSWER: C

Explanation:

“An IP prefix filter is used to filter IP address prefixes and can match an IP prefix number and a prefix length at the same time.” is correct. On Huawei VRP devices, an IP prefix list defines matching rules for network prefixes used in route-policy processing, such as route filtering and route attribute manipulation. Each index entry contains an IPv4 or IPv6 prefix together with a mask length, and it can also define a permitted range of prefix lengths through the `greater-equal` and `less-equal` parameters. This means matching evaluates both the prefix bits and the route’s prefix length. For example, a rule based on `10.0.0.0 8 greater-equal 16 less-equal 24` matches routes within 10.0.0.0/8 whose masks range from /16 through /24. This precise control is essential when a routing policy needs to accept, deny, advertise, import, or redistribute only selected route prefixes. Huawei documents IP prefix lists as route-filtering objects that can be referenced by route policies and related routing configuration, with the configured network address and mask length forming the basis of the match. See Huawei’s IP Prefix List documentation and the Route-Policy documentation.

QUESTION NO: 5

On an OSPF network, there are multiple types of packets, each with a different function. Which of the following are the main functions of Hello packets?

A. Neighbor relationship maintenance

B. Neighbor discovery

C. Neighbor deletion

D. Neighbor relationship establishment

ANSWER: A B D

Explanation:

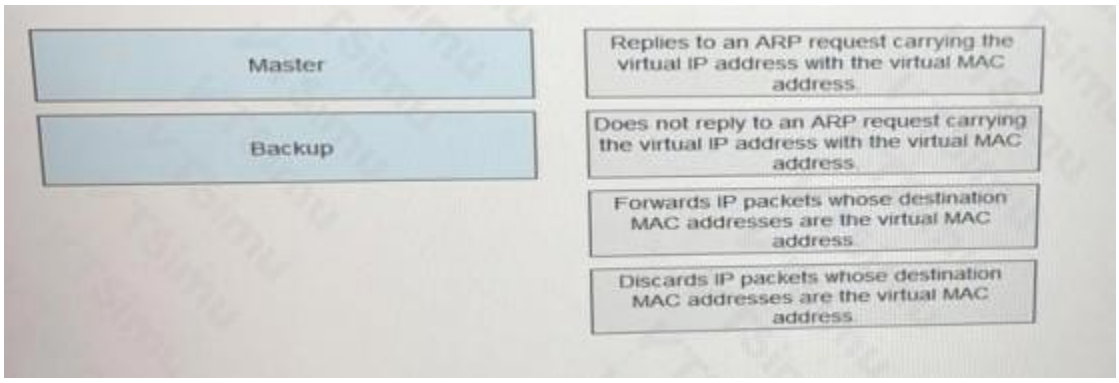
“Neighbor discovery,” “Neighbor relationship establishment,” and “Neighbor relationship maintenance” are the principal functions of OSPF Hello packets. An OSPF router sends Hellos periodically on an enabled interface so that other routers on the attached network can identify it as an OSPF-speaking neighbor. This discovery process uses Hello packet fields such as the area ID, network mask where applicable, Hello and Dead intervals, authentication information, and options; compatible values allow the routers to progress their neighbor relationship.

Hello packets also carry a list of neighbors from which a router has recently received Hellos. Seeing its own router ID in a neighbor’s Hello lets a router recognize bidirectional communication and reach the 2-Way state. On network types where an adjacency is required, this provides the basis for the subsequent adjacency-establishment exchange. Finally, the ongoing periodic Hello transmission and receipt maintains liveness. Each received Hello resets the inactivity timer; consequently, the relationship remains valid while timely Hellos continue to arrive. These functions are defined in the OSPF specification’s Hello-protocol description and neighbor-state processing sections: RFC 2328, Section 10.5 and RFC 2328, Section 10.3.

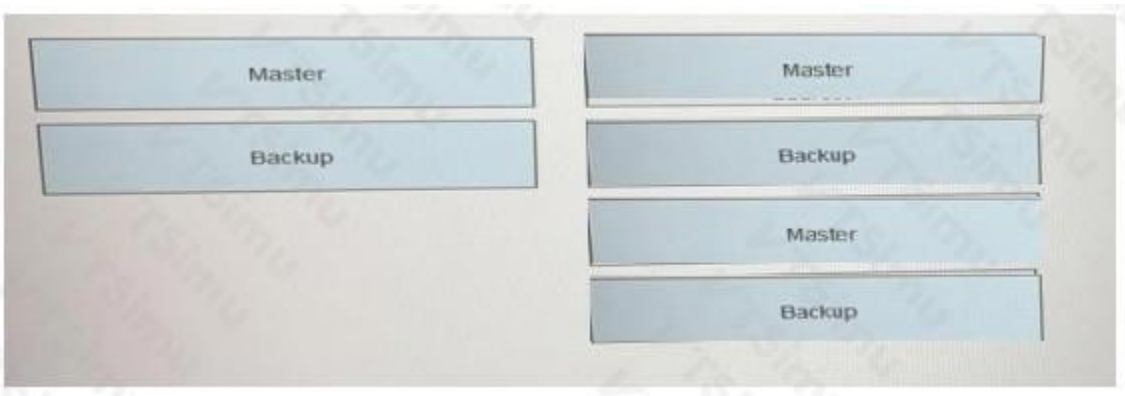
QUESTION NO: 6 - (DRAG DROP)

DRAG DROP

Drag the following VRRP states to the corresponding working mechanisms.



ANSWER:



Explanation:

In VRRP, routers in a virtual router group use a shared virtual IP address and virtual MAC address to provide a resilient default gateway. The router in the Master state is the active forwarding device for that virtual router. Because it currently owns the virtual gateway function, it answers ARP requests for the virtual IP address using the virtual MAC address. Hosts therefore learn the virtual MAC as the Layer 2 destination associated with their default gateway, while the active router remains able to take responsibility even if the physical device acting as Master changes.

The Master router also accepts and forwards packets that arrive with the virtual MAC address as their destination MAC address. This behavior is fundamental to transparent gateway redundancy: end hosts send traffic to one stable virtual gateway identity, and the Master handles that traffic. Huawei's VRRP documentation describes the virtual router concept and the Master/Backup roles in its VRRP overview.

A router in the Backup state monitors the Master and is prepared to take over when the Master is unavailable, but it does not actively provide the virtual gateway service during normal operation. Accordingly, it does not answer ARP requests for the virtual IP with the virtual MAC, and it discards frames addressed to that virtual MAC rather than forwarding them. This prevents two routers from simultaneously forwarding traffic for the same virtual router and preserves a single active forwarding role. Therefore, the displayed sequence of Master, Backup, Master, Backup correctly matches the four VRRP mechanisms shown.

QUESTION NO: 7

What parameters can a DHCP6 server assign to a DHCPv6 client?

A. Gateway address

B. DNS server address

C. IPV6 address/prefix

D. SNTP server address

ANSWER: B C D

Explanation:

DHCPv6 can provide several categories of IPv6 host configuration. **DNS server address** is a standard DHCPv6 configuration item delivered through the Domain Name System Recursive Name Server option, allowing a client to learn recursive DNS resolver addresses. **IPV6 address/prefix** is also correct: in stateful DHCPv6, a server can allocate an IPv6 address using an Identity Association for Non-temporary Addresses, and it can delegate an IPv6 prefix using Prefix Delegation. This enables both ordinary host addressing and downstream-network addressing for devices such as routers. **SNTP server address** can likewise be assigned through the DHCPv6 Simple Network Time Protocol Servers option, which supplies time-server information to a client. These capabilities illustrate that DHCPv6 supports both address or prefix allocation and delivery of supplementary network-service parameters. The applicable DHCPv6 base protocol is specified in RFC 8415; DNS recursive-server configuration is defined by RFC 3646. In normal IPv6 host configuration, default-router information is learned from IPv6 Router Advertisements rather than being a DHCPv6-assigned parameter.

QUESTION NO: 8

Which of the following statements regarding different LSA types is false?

A. LS Request packets contain only LS Type, LS ID, and Advertising Router.

B. LS Ack packets contain complete LSA information.

C. DD packets contain only LSA summary information, including LS Type, LS ID, Advertising Router, and LS Sequence Number.

D. LS Update packets contain complete LSA information.

ANSWER: B

Explanation:

LS Ack packets contain complete LSA information. is the false statement. In OSPF, Link State Acknowledgment packets are used to reliably acknowledge receipt of LSAs flooded by another router. An LS Ack packet does not repeat or carry the complete body of each acknowledged LSA. Instead, it carries the LSA header for each acknowledged advertisement. The header provides the identifying and version-control fields needed to identify precisely which LSA was received, such as the LS type, link-state ID, advertising router, sequence number, checksum, and length. This lets the sender match the acknowledgment with the flooded LSA while avoiding unnecessary retransmission of the entire advertisement.

The complete LSA is conveyed in a Link State Update packet, which is the OSPF packet type used for flooding new or changed link-state information. Separating the full-update function from the acknowledgment function reduces protocol overhead and supports OSPF's reliable flooding process. RFC 2328 specifies that an LS Acknowledgment packet contains a list of LSA headers, whereas Link State Update packets carry one or more complete LSAs. See RFC 2328, LS Acknowledgment Packet and RFC 2328, Link State Update Packet.

QUESTION NO: 9

Which of the following statements regarding OSPF is true?

A. OSPF does not have an acknowledgement mechanism. Therefore, OSPF relies on the upper-layer protocol, TCP, for acknowledgement.

B. OSPF performs LSDB update every 30 minutes.

C. OSPF uses the Bellman-Ford algorithm, and each router independently runs this algorithm.

D. OSPF floods a LSU packet at an interval of 5s.

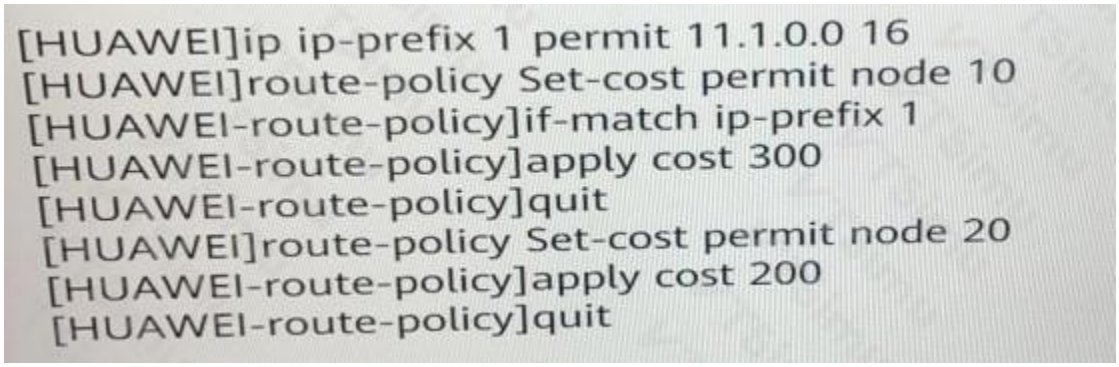
ANSWER: B

Explanation:

“OSPF performs LSDB update every 30 minutes.” is correct in the sense that OSPF refreshes its self-originated link-state advertisements at the default LSRefreshTime of 1,800 seconds (30 minutes). This periodic refresh prevents link-state information from becoming stale and ensures that the information maintained in the link-state database continues to be advertised throughout the OSPF area, even when no topology change has occurred. A refreshed LSA is reoriginated with the appropriate sequence-number handling and is flooded through the area, allowing routers to retain a synchronized, current view of the topology before LSAs reach their maximum lifetime. The relevant timer behavior is defined by the OSPF specification: LSRefreshTime is 1,800 seconds, while MaxAge is 3,600 seconds. In practical Huawei OSPF deployments, this standard LSA aging and refresh behavior supports reliable SPF topology calculations and database synchronization among adjacent routers. See the OSPF version 2 specification in RFC 2328, particularly its definition of LSRefreshTime and LSA aging, and Huawei’s OSPF configuration documentation at Huawei Support Documentation.

QUESTION NO: 10

Regarding the route-policy set-cost configuration below, which of the following statements is true?



A. The route 11.1.0.0/16 is permitted by node 10, and its cost is set to 300.

B. The cost of all routes is set to 200.

C. All the routes that are not permitted by node 10 will be denied.

D. The route 11.1.0.0/16 will continue to match node 20 after permitted by node 10, and (he final cost is set to 200.

ANSWER: A

Explanation:

The route 11.1.0.0/16 is permitted by node 10, and its cost is set to 300. In Huawei VRP, route-policy nodes are processed in ascending node-number order. When a route matches the `if-match` condition in a permit node, the actions configured in that node are performed. Here, the matching prefix 11.1.0.0/16 reaches node 10, is permitted, and the `apply cost 300` action modifies its route cost to 300.

Route-policy processing ends once the route has matched a node and that node has permitted it; the route is not subsequently evaluated against later nodes for additional cost changes. Therefore, the cost assigned at node 10 is the resulting cost for this route. This behavior is important when building policies with multiple nodes: node sequence determines which matching rule takes effect, and a successful match does not continue through later permit nodes.

Huawei documents route-policy matching, node sequencing, and apply actions in its VRP configuration documentation; see Huawei VRP documentation and the route-policy configuration reference.

QUESTION NO: 11

An enterprise administrator views the following details about a BGP route during routine O & M:

yaml

CopyEdit

< HUAWEI > display bgp routing-table 192.168.1.1

BGP local router ID: 10.1.1.1

Local AS number: 100

Paths: 2 available, 0 best, 0 select

BGP routing table entry information of 192.168.1.1/32:

From: 10.1.1.2 (10.1.1.2)

Route Duration: 00h01m31s

Relay IP Nexthop: 0.0.0.0

Relay IP Out-Interface: --

Original nexthop: 172.16.1.2

AS-path: 200, origin incomplete, MED 0, localpref 100, pref-val 0, internal, pre 255, invalid for IP unreachable

Not advertised to any peer yet

Which of the following statements are true about the BGP route?

A. The original next hop of the route is 172.16.1.2.

B. The local BGP router ID is 10.1.1.1.

C. The local AS number is 100.

D. The route is preferentially selected because its Local_Pref has a higher priority.

ANSWER: A B C

Explanation:

The original next hop of the route is 172.16.1.2 is correct because the route-detail output explicitly presents `Original nexthop: 172.16.1.2`. On Huawei VRP, the original next hop records the next-hop attribute carried with the received BGP path, independently of any relay next-hop information that might be used for recursive forwarding resolution.

The local BGP router ID is 10.1.1.1 is also correct. This is a process-level BGP value displayed in the command header as `BGP local router ID: 10.1.1.1`; it identifies the local BGP speaker, rather than the peer from which this specific route was learned.

The local AS number is 100 is correct because the same header directly states `Local AS number: 100`. This identifies the autonomous system configured for the local BGP process. The displayed route was received from a path containing AS 200, but that path information does not alter the locally configured AS number shown by the device. These fields are direct operational-state values reported by the Huawei BGP routing-table command. BGP route attributes and the role of the NEXT_HOP attribute are standardized in RFC 4271; Huawei product documentation and command references are available through the Huawei Enterprise Support portal.

QUESTION NO: 12

Which of the following statements are false about VRRP timers? (Select all that apply)

A. Skew_Time = (255 - Priority)/255

B. By default, the VRRP preemption delay is 1s.

C. The default interval for sending VRRP Advertisement packets is 2s.

D. Master_Down = (3 × Adver_Interval) + Skew_Time

ANSWER: A B C

Explanation:

The statements “Skew_Time = (255 - Priority)/255”, “By default, the VRRP preemption delay is 1s.”, and “The default interval for sending VRRP Advertisement packets is 2s.” are false. VRRP calculates Skew_Time using the router priority relative to 256 and includes the Advertisement Interval: Skew_Time = (256 − Priority) / 256 × Advertisement_Interval. This calculation produces a priority-dependent offset so that routers do not all attempt takeover at precisely the same instant after a master failure. The VRRP default Advertisement Interval is 1 second, not 2 seconds. Consequently, timer calculations based on default behavior use a one-second advertisement interval unless the interval has been explicitly configured otherwise. On Huawei VRP devices, a higher-priority backup router can preempt without an added delay by default; a preemption delay is an optional configured behavior rather than a default one-second timer. These timer rules are consistent with the VRRP specification’s definitions of Advertisement Interval, Skew_Time, and master-down detection. See RFC 5798, VRRP packet Advertisement Interval and RFC 5798, Skew_Time and Master_Down_Interval calculations.

QUESTION NO: 13

In PIM-DM, which of the following processes are involved in SPT establishment?

A. Prune

B. Graft

C. Flooding

D. State-refresh

ANSWER: A B C

Explanation:

In PIM Dense Mode, multicast forwarding begins with a source-rooted shortest-path tree: routers use their unicast-routing RPF information to identify the upstream interface toward the source and initially distribute packets outward on eligible downstream interfaces. Therefore, **Flooding** is part of establishing the initial forwarding tree and making traffic available throughout the dense-mode domain. As receivers and routers determine that particular branches have no interested downstream members, **Prune** processing removes those unnecessary branches from the distribution tree, leaving the active source tree paths that are needed for delivery. If a previously pruned branch subsequently requires the multicast stream, **Graft** processing is used to rapidly reattach that branch to the source tree and resume forwarding without waiting for a periodic flood-and-prune cycle. Together, flooding, pruning, and grafting form the core PIM-DM tree construction and adaptation behavior. The PIM Dense Mode protocol specification describes the initial broadcast/flood behavior, prune state, and graft operation in its forwarding state machine and message definitions. See RFC 3973: PIM Dense Mode and RFC 3208: PIM-DM State Refresh.

QUESTION NO: 14

A non-client is an IBGP peer that functions as neither an RR nor a client. A non-client must establish fully meshed connections with the RR and all the other non-clients.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. In a route-reflector topology, a non-client is an internal BGP peer of the route reflector that has not been configured as a route-reflector client. Route reflection has specific advertisement behavior: routes learned by a route reflector from a non-client peer are reflected to its clients, while the non-client side retains the normal internal BGP split-horizon behavior for routes learned from other internal peers. Consequently, non-client peers need ordinary internal BGP

connectivity among themselves to ensure that routes can be exchanged throughout that part of the autonomous system. They also establish internal BGP peerings with the route reflector so the reflector can distribute eligible routes to its client groups and participate in the overall route-reflection design.

This distinction is fundamental when designing a scalable internal BGP topology. Client peers can avoid a full mesh because the route reflector reflects appropriate routes between clients and toward non-clients. Non-client peers do not receive the same full-mesh reduction benefit from a given route reflector, so they are normally fully meshed with the route reflector infrastructure and with one another. This behavior follows the route-reflection rules defined in RFC 4456: BGP Route Reflection, particularly the treatment of routes received from client and non-client peers. Huawei's BGP configuration guidance also describes route reflectors as the mechanism used to reduce the full-mesh requirement specifically for client peers: Huawei BGP Configuration Documentation.

QUESTION NO: 15

On an OSPF network, an algorithm is used to prevent loops within an area, but loops may occur between areas. Therefore, OSPF defines a loop prevention mechanism for inter-area routes. Which of the following statements are true about the loop prevention mechanism?

A. Inter-area routes cannot be directly transmitted between non-backbone areas.

B. All non-backbone areas must be logically connected to area 0.

C. Inter-area routes need to be forwarded through area 0.

D. An ABR cannot inject Type 3 LSAs that describe routes to a network segment in an area back to the same area.

ANSWER: A B C D

Explanation:

OSPF prevents inter-area routing loops through its hierarchical area design and strict route-advertisement rules. Inter-area routing information is exchanged through the backbone, Area 0; non-backbone areas do not directly exchange inter-area routes with one another. Therefore, traffic and route information between separate non-backbone areas are carried through Area 0. Each non-backbone area must be logically connected to the backbone, either through a physical Area 0 connection or, where necessary, through a virtual link that provides logical backbone connectivity.

Area Border Routers preserve this hierarchy when generating summary LSAs. An ABR must not inject a Type 3 summary LSA describing destinations from an area back into that same area. This prevents a router in the area from learning an intra-area destination as an inter-area path via the ABR, which could create inconsistent route selection and enable looping behavior. These constraints ensure that inter-area paths follow a controlled, backbone-centered topology while OSPF's SPF calculation provides loop-free paths within each individual area. The OSPF inter-area routing and summary-LSA behavior are defined in RFC 2328, particularly its descriptions of area border routers and inter-area routing.

QUESTION NO: 16

The IP prefix list configuration on a Huawei router is shown below. Which of the following routes can match the IP prefix list and be permitted?

ip ip-prefix huawei index 10 deny 10.1.1.0 24

ip ip-prefix huawei index 20 permit 10.1.1.1 32

ip ip-prefix huawei index 30 permit 10.1.2.0 24 less-equal 32

A. 10.1.1.2/32

B. 10.1.2.5/32

C. 10.1.2.0/25

D. 10.1.1.0/24

E. Both 10.1.2.5/32 and 10.1.2.0/25

ANSWER: E

Explanation:

Both **10.1.2.5/32 and 10.1.2.0/25** can be permitted by this IP prefix list. Huawei evaluates IP-prefix entries in ascending index order. The entry `10.1.2.0 24 less-equal 32` permits prefixes whose address range falls within `10.1.2.0/24` and whose prefix length is from /24 through /32, inclusive. A host route for 10.1.2.5 has a /32 mask and lies inside the 10.1.2.0/24 range. Likewise, 10.1.2.0/25 is a more-specific subnet of 10.1.2.0/24 and has a mask length within the permitted /24-to-/32 range. Consequently, each route satisfies the same permit entry. The original choices therefore contain two individually valid routes, which conflicts with the stated Single Choice format; the combined answer accurately represents all matching permitted routes. Huawei IP-prefix matching uses the configured base prefix together with the optional mask-length range to determine whether a route matches. See Huawei's ip ip-prefix command reference and the Huawei IP routing documentation.

QUESTION NO: 17

An OSPF area ID identifies an OSPF area. Which of the following OSPF area IDs is equivalent to 0.0.1.0?

A. 1

B. 255

C. 256

D. 10

ANSWER: C

Explanation:

256 is the equivalent decimal OSPF area identifier. An OSPF area ID is a 32-bit value and may be displayed either as a decimal number or in dotted-decimal notation that resembles an IPv4 address. The dotted form represents four octets of the same 32-bit value. For 0.0.1.0, the first two octets contribute zero, the third octet has a value of one in the 256s position, and the final octet contributes zero. Its numeric value is therefore calculated as $(0 \times 256^3) + (0 \times 256^2) + (1 \times 256) + 0$, which equals 256.

This is only a representation conversion; it does not create a different OSPF area. A router configured with area 0.0.1.0 refers to the same area as a router configured with area 256, provided that all other OSPF adjacency requirements are met. The OSPF specification defines the Area ID as a 32-bit field and permits it to be expressed in either decimal or dotted-decimal form. See RFC 2328, OSPF Version 2 and Huawei's OSPF configuration documentation for OSPF area configuration concepts.

QUESTION NO: 18

An enterprise administrator wants to configure single-hop BFD to implement fast detection of direct links. Which of the following configurations are mandatory?

A. Configure the remote discriminator of a BFD session.

B. Configure the local discriminator of a BFD session.

C. Configure a multicast IP address for BFD.

D. Enable BFD globally.

ANSWER: A B D

Explanation:

For a manually configured single-hop BFD session on Huawei VRP devices, BFD must first be enabled globally so that the device can create and run BFD sessions. A static session also requires a local discriminator, which is the locally assigned nonzero identifier carried in BFD control packets. The remote discriminator must be configured as the identifier used by the directly connected peer. Together, these discriminator values uniquely associate received BFD control packets with the intended session and ensure that each end of the link has a consistent view of the peer relationship. In practice, the local discriminator configured on one device corresponds to the remote discriminator configured on its neighbor, and the reverse relationship is configured at the other end. Once the BFD session is established, its negotiated detection timers allow rapid recognition of a direct-link or neighbor forwarding failure, enabling upper-layer features such as routing protocols to react quickly. Huawei's BFD configuration guidance documents global BFD enablement and the local/remote discriminator settings as the core elements of static BFD session configuration. See the Huawei BFD configuration documentation and the BFD protocol specification (RFC 5880) for discriminator behavior and session establishment details.

QUESTION NO: 19

BGP can select routes based on the AS_Path attribute. Therefore, in some cases, a route-policy needs to be used to modify the AS_Path attribute for route selection. Which of the following parameters can be specified in the apply as-path command for a route-policy to modify this attribute?

A. Delete

B. Additive

C. Copy

D. Overwrite

ANSWER: A B D

Explanation:

Huawei route-policies can alter the BGP AS_Path attribute with `apply as-path`, allowing an administrator to influence BGP route selection and advertisement behavior. **Delete** is supported for removing AS numbers that match a configured AS-path filter from the route's existing AS_Path. This permits controlled path manipulation according to the policy's matching criteria. **Additive** is supported when specified AS numbers must be added while retaining the original AS_Path. In practice, this is used for AS-path prepending, making an advertised route appear to have traversed additional autonomous systems and therefore less preferable to external BGP peers that compare AS-path lengths. **Overwrite** is also supported to replace the current AS_Path with the AS numbers specified by the route-policy. This provides a direct way to set a replacement path where that behavior is required by the routing policy. These operations are performed when the relevant route-policy node is applied to BGP route processing. Huawei's BGP documentation describes AS_Path as a path-vector attribute used in route selection, and Huawei's route-policy command documentation covers applying AS-path changes through routing policy actions. See Huawei VRP Route Policy Configuration Guide and Huawei BGP documentation.

QUESTION NO: 20

On a WLAN, the HSB service sets up an HSB channel between two devices that back up each other, maintains the channel status, and backs up data. Which of the following can HSB back up in real time?

A. CAPWAP tunnel information

B. DHCP address information

C. AP entries

D. User data information

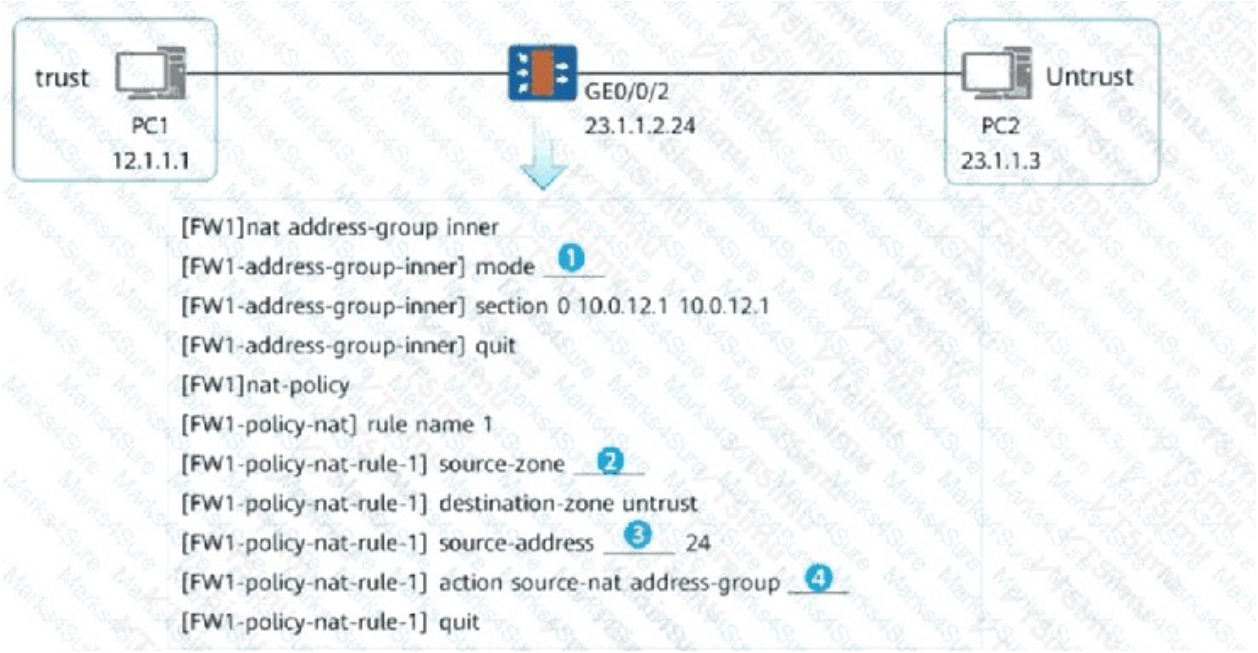
ANSWER: A B C D

Explanation:

Huawei Hot Standby Backup for WLAN controllers synchronizes dynamic service state between the active and standby devices so that the standby controller can take over with minimal service interruption. This real-time synchronization includes CAPWAP tunnel information, allowing the standby device to retain the tunnel-related state needed for managed APs. It also includes AP entries, which preserve the AP registration and operational information required to continue AP management after a switchover.

User data information is synchronized because active wireless client service state, such as associated user/session information, must be available on the standby controller to support a smooth service takeover. DHCP address information is also included in the protected dynamic state where the controller provides DHCP services: synchronizing address-allocation or lease-related information helps prevent inconsistent client address assignments after an active/standby switchover. Together, these data types allow both AP connectivity and wireless-user services to be restored or maintained quickly when the active device fails. Huawei's enterprise support portal provides the product documentation and configuration guides for WLAN controller high-availability features: Huawei Enterprise Support. See also Huawei's WLAN product resources at Huawei WLAN.

QUESTION NO: 21 - (SIMULATION)



ANSWER: See the explanation for the answer

Explanation:

The four blanks are:

`mode pat` enables NAPT (PAT) for the address group. The NAT policy matches traffic entering from the `trust` zone with source subnet `12.1.1.0/24`, and translates the source address using NAT address group `inner`.

QUESTION NO: 22

Which of the following is the default dead time of an interface on an OSPF P2P or broadcast network?

A. 40s

B. 120s

C. 60s

D. 10s

ANSWER: A

Explanation:

The default dead time for an interface operating in an OSPF point-to-point or broadcast network is **40s**. OSPF uses Hello packets to discover neighbors and maintain established neighbor relationships. On these network types, Huawei VRP uses a default Hello interval of 10 seconds. The dead interval is normally four times the Hello interval, so a router declares a neighbor down if it does not receive a Hello packet from that neighbor within 40 seconds. This timer directly affects failure-detection speed: after the 40-second period expires without valid Hellos, OSPF removes the neighbor relationship and initiates the appropriate link-state recalculation and flooding behavior. For routers to form and maintain an OSPF adjacency, parameters including the Hello interval and dead interval must be compatible between the two neighboring interfaces. The 40-second default applies specifically to broadcast and point-to-point OSPF interface network types; timer values can be manually adjusted when operational requirements call for faster or slower detection. Huawei's OSPF configuration documentation identifies the default Hello timer as 10 seconds and the default dead timer as 40 seconds for these interface types. See Huawei: Understanding OSPF and Huawei: Configuring OSPF Interface Parameters.

QUESTION NO: 23

Which of the following statements are true about the **Local_Pref** attribute in BGP? (Choose all that apply)

A. The default Local_Pref value is 0.

B. The Local_Pref attribute can be transmitted to all BGP peers.

C. The Local_Pref value remains unchanged during transmission unless it is affected by a routing policy.

D. The greater the value of the Local_Pref attribute, the higher the priority.

ANSWER: C D

Explanation:

The statements "The Local_Pref value remains unchanged during transmission unless it is affected by a routing policy." and "The greater the value of the Local_Pref attribute, the higher the priority." correctly describe BGP local preference behavior. LOCAL_PREF is an attribute used within an autonomous system to express the preferred exit path for destinations. When several valid BGP paths are available, BGP prefers the route with the highest LOCAL_PREF value before considering many later best-path attributes. This lets an administrator implement consistent outbound-path selection for the entire AS.

LOCAL_PREF is carried between internal BGP peers, allowing the routing decision made or configured at one router to be distributed to other routers in the same AS. A received or locally assigned value is retained as the route is propagated through iBGP unless a routing policy deliberately changes it. On Huawei devices, route-policies can set Local_Pref values for selected routes, making it a key tool for traffic-engineering policy and for preferring one external connection over another. The standard default local-preference value is 100 when no policy supplies a different value. See RFC 4271, Section 5.1.5 and Huawei's BGP configuration documentation for the attribute's internal-AS use and best-path role.

QUESTION NO: 24 - (SIMULATION)

In BGP, Keepalive messages are used to maintain BGP peer relationships. When a BGP router receives a Keepalive message from a peer, the BGP router sets the state of the peer to Established and periodically sends Keepalive messages to maintain the connection. By default, the device sends Keepalive messages every seconds.

ANSWER: See the explanation for the answer

Explanation:

Answer: 60 seconds.

By default, Huawei BGP sends a Keepalive message every `60` seconds to maintain an Established BGP neighbor relationship. The default Hold timer is typically 180 seconds, which is three times the Keepalive interval.

QUESTION NO: 25

iMaster NCE-Fabric networks. It helps customers improve service efficiency from multiple dimensions. Which of the following functions is not supported by the controller?

A. Management

B. Control

C. Forwarding

D. Analysis

ANSWER: C

Explanation:

Forwarding is the function not supported by the iMaster NCE-Fabric controller itself. iMaster NCE-Fabric is Huawei's network management and control platform for data center networks. It provides centralized management capabilities, automated service provisioning, policy-based control, telemetry collection, and network analysis/optimization functions. These capabilities allow the controller to maintain a global view of the fabric and translate service intent into device configurations and control policies.

Actual packet forwarding is performed by the physical network devices in the fabric, such as switches and routers. Their data planes receive packets, look up forwarding information, apply encapsulation or policy actions where required, and transmit traffic through the appropriate interfaces. The controller can calculate, distribute, and adjust the information that guides traffic handling, but it is not in the packet-by-packet forwarding path. This separation of centralized management/control from distributed device forwarding is fundamental to controller-based network architecture and enables scalable fabric operations. Huawei describes iMaster NCE-Fabric as an automation and management platform for data center networks in its iMaster NCE-Fabric product information and related technical documentation.

QUESTION NO: 26

On an IS-IS network, each device needs to be configured with a **network entity title (NET)** , which consists of three parts, each containing different fields. Which of the following fields are contained in the **Area Address part** ?

A. High Order DSP

B. AFI

C. IDI

D. System ID

ANSWER: A B C

Explanation:

An IS-IS Network Entity Title uses the OSI Network Service Access Point addressing structure. Its area-address portion is made up of the Authority and Format Identifier, the Initial Domain Identifier, and the High Order Domain Specific Part. Therefore, **High Order DSP**, **AFI**, and **IDI** are all fields contained in the Area Address part.

The AFI identifies the authority and format used for the address. The IDI identifies the initial domain under that authority. Together with the High Order DSP, these fields form the variable-length area address that identifies the IS-IS routing area. In practical IS-IS NET notation, the area portion precedes the fixed-length system ID and the final N-selector. For example, in a NET such as `49.0001.1921.6800.1001.00`, `49.0001` is the area address, `1921.6800.1001` is the system ID, and `00` is the N-selector used by IS-IS. The High Order DSP is specifically part of the area-address hierarchy, which is why it must be selected along with AFI and IDI. See the ISO addressing description in RFC 1195 and Huawei's IS-IS configuration documentation.

QUESTION NO: 27

Which of the following is used as the destination port for single-hop BFD?

A. UDP port 3784

B. UDP port 4784

C. TCP port 3784

D. TCP port 4784

ANSWER: A

Explanation:

UDP port 3784 is the destination port used by the BFD control protocol for single-hop sessions. Single-hop BFD monitors reachability between directly connected systems, and its control packets are carried in UDP. The standard assigns UDP destination port 3784 to single-hop BFD control packets, allowing devices to identify and process these packets as BFD traffic. The sending device uses a dynamically allocated UDP source port, while the receiver listens on the well-known destination port 3784. This port assignment is part of the BFD single-hop encapsulation defined by the IETF, ensuring interoperable operation between devices from different vendors. On Huawei routers, when BFD is configured for a directly connected peer or link, the protocol behavior follows this standardized single-hop UDP encapsulation. The IANA Service Name and Transport Protocol Port Number Registry lists *bfd-control* on UDP port 3784, and RFC 5881 specifies the single-hop BFD control-packet UDP destination port. References: RFC 5881: BFD for IPv4 and IPv6 (Single Hop) and IANA port registry.

QUESTION NO: 28

On a **BGP network** , a **route-policy** can be used to modify route attributes. Which of the following statements is **false** about **route-policies** ?

A. When OSPF routes are imported using the import command, a **route-policy** can be used to change their default **MED value** .

B. When direct routes are imported using the import command, a **route-policy** can be used to filter out the routes that you do not want to advertise.

C. When routes are imported using the network command, a **route-policy** cannot be used to modify the attributes of the imported routes.

D. In the address family view, the peer command can be run to reference a **route-policy** to modify the **local preference** of the routes to be advertised.

ANSWER: D

Explanation:

“In the address family view, the peer command can be run to reference a **route-policy** to modify the **local preference** of the routes to be advertised.” is false. On Huawei VRP, local preference is an intra-AS BGP path-selection attribute: it is used by routers within the local AS to choose an exit path, and it is propagated only to IBGP peers rather than to EBGP peers. A route-policy action that sets local preference is therefore intended for BGP routes processed on import, so that the device can use the resulting value in its own BGP decision process and, where applicable, advertise the value within its AS. It is not an export-policy action for altering the local-preference attribute of routes being advertised to a peer. In contrast, import-route policies can filter redistributed routes and apply permitted BGP attribute actions, while a policy attached to a BGP **network** statement is used to control route origination rather than modify the attributes of the originated route. Huawei’s BGP documentation describes local preference as an attribute used within an AS and documents route-policy use in BGP route processing. See Huawei BGP Configuration Guide and Huawei VRP Routing Policy Configuration Guide.

QUESTION NO: 29

Which of the following statements regarding LACP are correct?

A. LACP member interfaces exchange LACPDUs.

B. An active interface can establish an LACP aggregation with a passive interface.

C. Two active interfaces can establish an LACP aggregation.

D. Two passive interfaces can establish an LACP aggregation without receiving LACPDUs from an active peer.

E. LACP does not use system priority during negotiation.

ANSWER: A B C

Explanation:

LACP is defined by IEEE 802.1AX and is used to dynamically negotiate and maintain an Eth-Trunk between connected devices. LACP-enabled interfaces exchange LACPDUs to determine whether member links can join the aggregation group. An interface in active mode initiates LACP negotiation, whereas an interface in passive mode responds to received LACPDUs. Therefore, an active/passive pair and an active/active pair can establish an LACP aggregation.

Two passive interfaces do not initiate LACP negotiation and therefore cannot establish an LACP aggregation by themselves. LACP also uses system priority and system ID information when selecting the Actor and Partner system roles. See IEEE 802.1AX.

QUESTION NO: 30

On a network, each router has a local core routing table and protocol routing tables. A routing entry in the local core routing table has multiple key fields. Which of the following are included?

A. Destination address of a route

B. Routing protocol preference of a route

C. Inbound interface that learns a route

D. Routing protocol that learns a route

ANSWER: A B D

Explanation:

A local core routing table, called the IP routing table in Huawei VRP documentation, is the unified table used for route selection and packet forwarding. Each installed route identifies the network that can be reached, the source protocol that supplied the route, and the route preference used to compare routes learned through different routing mechanisms. Therefore, *Destination address of a route* is a fundamental route-entry field: it identifies the destination prefix, normally together with its mask or prefix length. *Routing protocol preference of a route* is also included as the Preference (Pre) value. Huawei devices use this value to select between otherwise comparable routes to the same destination, with a smaller preference being more preferred. *Routing protocol that learns a route* is represented by the Proto field, which records the route source, such as Direct, Static, OSPF, IS-IS, or BGP. These fields are shown by the Huawei `display ip routing-table` command along with operational forwarding information such as cost, next hop, and outbound interface. Huawei's routing documentation describes the routing table as the basis for selecting optimal routes and forwarding IP packets. See the Huawei IP Routing Configuration Guide and the Huawei technical documentation portal.

QUESTION NO: 31

On a WLAN, MAC address authentication controls network access rights of a user based on the user ' s access interface and terminal MAC address.

Which of the following statements is false about MAC address authentication?

A. User passwords can be processed using PAP or CHAP.

B. By default, the MAC address of a terminal is used as the user name and the last 6 hexadecimal digits of the terminal MAC address are used as the password for MAC address authentication.

C. No client software needs to be installed on user terminals.

D. A MAC address authentication system consists of three types of entities: terminal, access device, and authentication server.

ANSWER: B

Explanation:

The statement “By default, the MAC address of a terminal is used as the user name and the last 6 hexadecimal digits of the terminal MAC address are used as the password for MAC address authentication.” is false. In Huawei MAC address authentication, the access device derives the authentication identity from the connecting terminal’s MAC address. Under the default MAC-address-based credential behavior, the terminal MAC address is used for both the user name and password; Huawei does not define the default password as merely the final six hexadecimal digits of that address. Administrators can configure the MAC address format used for authentication credentials—for example, choices concerning delimiters, letter case, and whether portions of an address are used—when required by the RADIUS server’s account format. Such configuration flexibility does not alter the documented default behavior. This mechanism enables transparent admission control because the device obtains the MAC address directly during access processing rather than requiring a user to enter credentials. Huawei’s enterprise documentation portal provides the relevant WLAN and authentication configuration references at Huawei Enterprise Documentation and Huawei Support Documentation Center.

QUESTION NO: 32

Which of the following statements is false about BFD?

A. The asynchronous mode is the primary BFD operating mode.

B. In asynchronous mode, two systems periodically exchange BFD Control packets at the negotiated interval. If one system does not receive any BFD Control packets from the other within the detection interval, the BFD session is declared down.

C. The asynchronous mode does not support the echo function.

D. In demand mode, after a BFD session is set up, the system does not periodically send BFD Control packets.

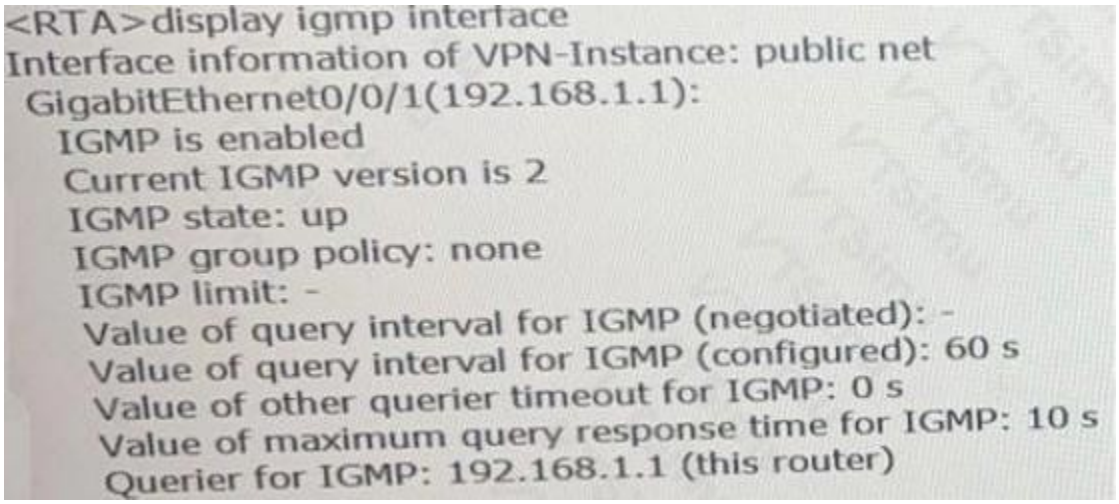
ANSWER: C

Explanation:

The statement “The asynchronous mode does not support the echo function.” is false. BFD asynchronous mode supports the Echo function as an optional mechanism for rapidly verifying forwarding-path continuity. In normal asynchronous operation, peers exchange BFD Control packets to maintain the session and detect failures. When Echo is enabled and supported by the relevant interface and feature configuration, a device also sends BFD Echo packets that are looped back through the forwarding path. The sender can then use the returned Echo packets to confirm that the data-plane path remains operational, while BFD Control packets continue to provide session state and parameter negotiation. This capability can improve failure-detection efficiency because Echo packets are processed through the forwarding path rather than requiring the remote BFD process to generate every return packet. Echo operation is optional rather than mandatory, and platform, interface, and protocol support must be considered when configuring it; however, its optional nature does not mean it is unsupported in asynchronous mode. The BFD base specification describes asynchronous mode and the Echo function, and the single-hop BFD specification defines Echo packet handling in more detail. See RFC 5880 and RFC 5881.

QUESTION NO: 33

Which of the following statements is false, based on the following IGMP information on an interface of RTA?



A. The interval for sending group-specific Query messages is 60s.

B. The maximum time for response to Query messages is 10s.

C. The IP address of the interface is 192.168.1.1.

D. The IGMP version is IGMPV2.

ANSWER: B

Explanation:

The statement “The maximum time for response to Query messages is 10s.” is false because the displayed 10-second value applies to the maximum response time carried in a general IGMP Query, rather than to every type of IGMP Query. In IGMPv2, a group-specific Query is sent when a router is checking whether members remain for a particular multicast group. Its response behavior is controlled by the last-member-query settings, especially the last-member-query interval, which is typically much shorter than the general-query maximum response time. Therefore, describing 10 seconds as the maximum response time for Query messages without distinguishing general Queries from group-specific Queries is inaccurate. This distinction is fundamental to IGMPv2 leave processing: group-specific Queries are designed to confirm remaining membership quickly before multicast forwarding state is removed. The IGMPv2 protocol specification defines the separate general-query and group-specific-query behavior, including use of the Last Member Query Interval for the latter. Huawei VRP implements these IGMP query and membership-maintenance mechanisms on multicast-enabled interfaces. See RFC 2236: Internet Group Management Protocol, Version 2 and the Huawei VRP multicast configuration documentation.