ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

DRAG DROP

Huawei modular devices have multiple hardware modules that provide different functions. Match the following hardware modules with their functions.

Answer:

Explanation:

The hardware modules of Huawei modular devices and their functions are:

Main Processing Unit (MPU): Provides control and management planes for the entire system, responsible for protocol processing, system security, and software upgrades.

Switch Fabric Unit (SFU): Provides the data plane, enabling high-speed data switching between service modules.

Line Processing Unit (LPU): Manages data forwarding, offering various interfaces (optical and electrical) for data access.

Question: 2

A router performs a lookup in its FIB table for a packet. If the tunnel ID in the matching entry is 0, the packet needs to be forwarded through a tunnel, such as an MPLS tunnel.

A. TRUE
B. FALSE

Answer: B

Explanation:

The question indicates that a router performs a lookup in its FIB table for a packet and determines that the tunnel ID in the

matching entry is 0, suggesting that the packet needs to be forwarded through a tunnel such as an MPLS tunnel. However, this is a misunderstanding of the FIB functionality.

FIB Table Overview

The Forwarding Information Base (FIB) is used to make packet-forwarding decisions. Entries in the FIB include next-hop information, which can be directly linked to an interface or a tunnel.

If the Tunnel ID is 0, it indicates that the packet is forwarded via a normal routing path and not through a tunnel.

For MPLS or other tunnels, the Tunnel ID would have a non-zero value pointing to the associated tunnel.

MPLS Tunnel Operation

When a router forwards packets through an MPLS tunnel, a label-switched path (LSP) is set up. The FIB would reflect specific tunnel identifiers for packets that need such encapsulation.

HCIP-Datacom-Core Reference

Routing Principles and MPLSexplain the forwarding mechanisms clearly, stating that if a packet is routed normally, the tunnel ID remains 0.

The section on MPLSclarifies the encapsulation process and the role of tunnel identifiers.

Hence, the claim in the question is incorrect. A Tunnel ID of 0 implies no tunneling is required, and normal IP forwarding occurs

Question: 3

On an OSPF network, one router with P2P as the network type is directly connected to another router with P2MP as the network type. If the Hello intervals on the two routers are changed to be the same, neighbor relationship establishment and LSDB synchronization are not affected.

A. TRUE
B. FALSE

Answer: B

Explanation:

The scenario describes a mismatch in OSPF network types between two connected routers: one set to Point-to-Point (P2P) and the other set to Point-to-Multipoint (P2MP). While aligning Hello intervals may seem sufficient for establishing an OSPF neighbor relationship, the fundamental mismatch in network types introduces issues.

OSPF Network Types

P2P: Assumes a direct connection with a single neighbor, uses faster convergence and simpler LSDB synchronization.

P2MP: Supports multiple neighbors on a single interface, requiring different handling for DR/BDR roles and LSDB

updates.

Impact of Network Type Mismatch

If Hello intervals are aligned, adjacency establishment might occur. However, mismatched network types affect neighbor role assignment and LSDB synchronization.

P2P expects a direct link and would handle updates differently than P2MP, which assumes multiple neighbors. This leads to inconsistencies in route calculation and forwarding.

HCIP-Datacom-Core Reference

OSPF Basics and Configurationclearly outlines the criticality of consistent network type configuration for stable OSPF operation.

Lab examples in the HCIP Datacom Lab Guidefurther demonstrate the consequences of such mismatches, including unstable neighbor states and incomplete LSDB synchronization.

Hence, the statement that neighbor relationships and LSDB synchronization remain unaffected is incorrect. Proper OSPF operation requires matching network types in addition to aligned Hello intervals.

Question: 4

On an enterprise network, the directly connected interfaces of two OSPF routers are on different network segments and have different masks. To establish an OSPF neighbor relationship between the two interfaces, you can change their network types to which of the following?

A. Point-to-point
B. NBMA
C. P2MP
D. Broadcast

Answer: A

Explanation:

When OSPF routers have interfaces on different network segments with different subnet masks, the network type can be adjusted to establish adjacency. A point-to-point (P2P) network type eliminates the requirement for matching subnet masks by treating the link as directly connected without intermediate devices.

P2P Network Characteristics

OSPF treats the link as a direct connection between two routers.

No DR/BDR election occurs, simplifying adjacency establishment.

Subnet mask differences do not hinder neighbor relationships as the link is viewed as a logical tunnel.

HCIP-Datacom-Core Reference

The OSPF configuration section explicitly mentions P2P as a suitable network type for resolving adjacency issues caused by mismatched subnet masks.

Question: 5

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 directly 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 Area Design and Loop Prevention:

OSPF uses a hierarchical structure with areas to improve scalability and efficiency. Area 0, the backbone area, plays a crucial role in ensuring loop-free route distribution between areas. The following mechanisms are key to preventing routing loops:

Strict adherence to hierarchical area design.

Prohibition of direct inter-area route exchanges between non-backbone areas.

Reference: HCIP-Datacom-Core Technology Training Material (OSPF Basics and Advanced Concepts).

Analysis of Each Statement:

A . Inter-area routes cannot be directly transmitted between non-backbone areas.
This statement is TRUE. OSPF mandates that all inter-area routing must pass through Area 0. Direct inter-area route exchanges between two non-backbone areas are not allowed to prevent loops.

Reference: HCIP-Datacom-Core Technology Training Material (Inter-Area Routing Mechanisms).

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

This statement is TRUE. OSPF requires every non-backbone area to connect directly to Area 0 to facilitate loop-free inter-area routing. Virtual links may be configured in exceptional cases where direct connection is not possible.

Reference: HCIP-Datacom-Core Technology Training Material (OSPF Backbone and Area Connectivity).

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

This statement is TRUE. All inter-area traffic must traverse Area 0 to ensure hierarchical routing and loop prevention. This rule is a core design principle of OSPF.

Reference: HCIP-Datacom-Core Technology Training Material (Routing Control and Loop Prevention in OSPF).

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

This statement is TRUE. OSPF explicitly prohibits an ABR from injecting Type 3 LSAs describing a route into the same area where the route originates. This mechanism prevents routing loops within an area.

Reference: HCIP-Datacom Advanced Routing & Switching Technology (OSPF LSA Types and ABR Behavior).

Conclusion:

All options (A, B, C, D) are correct. OSPF enforces a robust loop prevention mechanism through

hierarchical routing, mandatory traversal via Area 0, and strict rules on LSA propagation by ABRs. This ensures reliable and loop-free inter-area routing in OSPF networks.

Question: 6

OSPF has five types of packets, which have the same header format. If the Auth Type field in the packet header is 1, which of the following authentication modes is used?

A. Non-authentication
B. MD5 authentication
C. Plaintext authentication
D. Hash authentication

Answer: C

Explanation:

OSPF Authentication Overview

The Auth Type field in the OSPF packet header determines the authentication mode. If the Auth Type is 1, plaintext authentication is used.

Plaintext authentication involves transmitting the password in an easily readable format, which is less secure compared to MD5.

HCIP-Datacom-Core Reference

Authentication mechanisms, including plaintext authentication, are detailed in the OSPF security configuration chapter, confirming that Auth Type = 1 corresponds to plaintext.

Question: 7

On an OSPF network, if a router receives an update of an LSA that exists in the local LSDB, the router updates the LSDB and floods the LSA.

A. TRUE
B. FALSE

Answer: B

Explanation:

OSPF LSA Flooding Mechanism

If a router receives an LSA identical to one already in its LSDB, it does not flood the LSA again unless the LSA has changed (i.e., the sequence number or content has been updated).

OSPF ensures efficient use of bandwidth by avoiding redundant flooding of unchanged LSAs.

HCIP-Datacom-Core Reference

The OSPF LSDB synchronization process explains that unchanged LSAs are not reflooded, ensuring stability and resource optimization.

Question: 8

An enterprise uses OSPF to implement network communication. To ensure data validity and security, all authentication modes supported by OSPF are enabled on routers. In this case, interface authentication is preferentially used by the routers.

A. TRUE
B. FALSE

Answer: A

Explanation:

OSPF Authentication Overview

OSPF supports three authentication modes:

Null Authentication: No authentication (default).

Plaintext Authentication: Uses clear-text passwords.

MD5 Authentication: Secure cryptographic authentication.

Interface-Level Priority

When both interface-level and area-level authentication are configured, OSPF prioritizes interfacelevel authentication. This ensures that interface-specific security overrides area-wide configurations for greater granularity and security.

HCIP-Datacom-Core Reference

OSPF authentication hierarchy and configurations are detailed in the OSPF security configuration

chapter.

Question: 9

The following figure shows the OSPF network of an enterprise and the OSPF configurations of R1.

Which of the following statements is false about the network?

```
[R1]ospf
[R1 ospf-1 Jarea 0
[R1 -ospf-1 -area-0.0.0.0]network 192.168.4.0 0.0.0255
[R1 -ospf-1 -area-0.0.0.0]quit
[R1-ospf-1]silent-interfa« GigabitEthernet 0/0/1
```

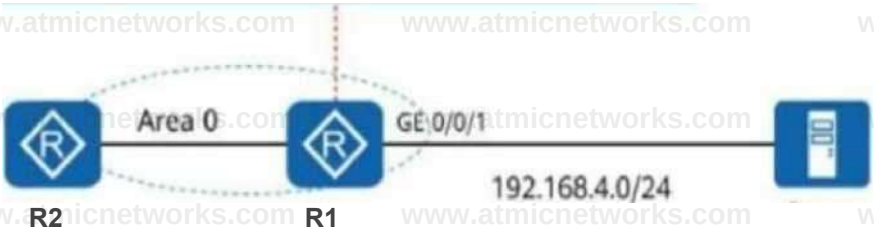

A. R2 can access the server.

B. GE 0/0/1 of R1 cannot send OSPF packets.

C. The network segment to which GE 0/0/1 of R1 belongs cannot be advertised.

D. GE 0/0/1 of R1 cannot accept OSPF packets.

Answer: C

Explanation:

Silent Interface Explanation

The silent-interface command is used to prevent OSPF from sending or receiving OSPF packets on the specified interface (GE 0/0/1). This disables OSPF adjacency establishment and stops route advertisement for that interface.

Network Observations

Statement A: R2 can access the server.

This is correct, as the silent interface does not impact data traffic, only OSPF-related communication.

Statement B: GE 0/0/1 of R1 cannot send OSPF packets.

Correct due to the silent-interface configuration.

Statement C: The network segment to which GE 0/0/1 of R1 belongs cannot be advertised.

This is correct, as the silent interface prevents route advertisement.

Statement D: GE 0/0/1 of R1 cannot accept OSPF packets.

Correct, as the silent interface configuration blocks packet reception.

HCIP-Datacom-Core Reference

OSPF interface command behavior is outlined in the configuration and lab examples sections.

Question: 10

On an OSPF network, if two routers with the same router ID run in different areas and one of the routers is an ASBR, LSA flapping occurs.

A. TRUE

B. FALSE

Answer: A

Explanation:

Understanding Router ID and Its Role in OSPF:

In OSPF, the Router ID uniquely identifies a router within the OSPF domain. If two routers are configured with the same Router ID, it can lead to issues such as LSA conflicts and flapping. This is because the Router ID is used as a key in OSPF operations, including LSA generation and database synchronization.

Reference: HCIP-Datacom-Core Technology Training Material (OSPF Basics - Router ID and LSA Handling).

Scenario Details:

Different Areas: Even if the two routers belong to different areas, the Router ID remains globally significant in the OSPF domain. This means that any duplication of Router IDs will confuse OSPF mechanisms.

ASBR (Autonomous System Boundary Router): When one of the routers is an ASBR, it generates Type 4 and Type 5 LSAs to describe external routes. These LSAs use the Router ID as an identifier. If another router in the network has the same Router ID, conflicts occur during LSDB synchronization.

Reference: HCIP-Datacom Advanced Routing & Switching Technology (LSA Types and ASBR Operations).

Impact of Router ID Duplication:

LSA Flapping: The OSPF process receives conflicting LSAs from routers with the same Router ID. This results in continuous updates and withdrawals of these LSAs, causing flapping.

Routing Instability: LSA flapping leads to frequent recalculations of the OSPF shortest path tree (SPT), affecting overall network stability.

Reference: HCIE-Datacom V1.0 Training Material (OSPF Troubleshooting and Best Practices).

Conclusion:

The statement is TRUE. LSA flapping occurs when two routers in an OSPF network have the same Router ID, even if they are in different areas and one is an ASBR. This is due to the global significance of Router IDs in OSPF and the role they play in LSA generation and propagation.

Question: 11

On an OSPF network, routers learn routing information on the entire network by exchanging LSAs.
Which of the following values is the LS Age in the LSA header when an LSA is deleted?

A. 1800s
B. 3600s
C. 1200s
D. 600s

Answer: B

Explanation:

LSA Lifetime and Deletion

The LS Age field in the LSA header tracks the age of an LSA. When the LS Age reaches its maximum value (3600 seconds), the LSA is marked for deletion. This ensures old or stale LSAs are removed from the network to maintain accurate routing information.

HCIP-Datacom-Core Reference

Detailed explanation of LS Age behavior and LSA deletion processes can be found in the OSPF LSDB and LSA sections.

Question: 12

DRAG DROP

OSPF networks are classified into four types of networks by link layer protocol. Drag the following link layer protocols to the corresponding network types.(Token is reusable)

Answer:

Explanation:

Network Types and Corresponding Link Layer Protocols

Broadcast: Ethernet

Point-to-Point (P2P): PPP, HDLC

Point-to-Multipoint (P2MP): PPP

Non-Broadcast Multi-Access (NBMA): Frame Relay

OSPF Network Types:

OSPF classifies networks based on link layer protocols into the following types:

Broadcast: This type assumes that all routers on the network can communicate directly with one another using multicast or broadcast frames. Ethernet networks are typical examples.

Point-to-Point (P2P): This type is used for links that connect two routers directly. Common protocols include PPP (Point-to-Point Protocol) and HDLC.

Point-to-Multipoint (P2MP): This type simulates multiple point-to-point connections over a single physical network, often used in WAN scenarios where PPP is employed.

Non-Broadcast Multi-Access (NBMA): These networks connect multiple devices but lack native broadcast capability, such as Frame Relay.

Reference: HCIP-Datacom-Core Technology Training Material (OSPF Network Types).

Explanation of Matches:

Broadcast - Ethernet: Ethernet supports broadcast and multicast communication, making it a suitable example of a broadcast OSPF network.

P2P - PPP, HDLC: Both PPP and HDLC are designed for direct communication between two nodes, fitting the P2P category.

P2MP - PPP: In WANs, PPP often operates in a point-to-multipoint configuration, simulating separate connections for each endpoint.

NBMA - Frame Relay: Frame Relay is a classic NBMA technology where direct communication between devices requires manual configuration, as there is no inherent broadcast capability.

Conclusion:

This classification ensures that OSPF operates efficiently over different network types by adapting neighbor discovery and LSA propagation mechanisms to the underlying link layer technology.

Question: 13

On an IS-IS network, each router can generate LSPs. Which of the following events trigger the generation of a new LSP?

A. Related IS-IS interfaces go up or down.
B. Periodic updates occur.
C. Inter-area IP routes change.
D. The IS-IS interface cost is increased.

Answer: A, B, D

Explanation:

LSP Generation in IS-IS

IS-IS routers generate new Link State Packets (LSPs) under the following conditions:

Interface Status Changes: When IS-IS interfaces go up or down, the link state changes, triggering LSP updates.

Periodic Updates: IS-IS periodically regenerates LSPs to ensure link-state information remains synchronized across the network.

Interface Metric Changes: Any modification to interface costs results in a new LSP to reflect the updated cost in the network.

Incorrect Option

C . Inter-area IP routes change is incorrect because IS-IS does not inherently differentiate between areas for LSP generation.

HCIP-Datacom-Core Reference

IS-IS LSP generation rules are detailed in the IS-IS configuration and implementation chapters.

Question: 14

Similar to the OSPF DR, the IS-IS DIS needs to be elected on a broadcast network. However, the OSPF DR is preemptive by default, whereas the IS-IS DIS is not preemptive by default.

A. TRUE

B. FALSE

Answer: A

Explanation:

DIS and DR Election

The IS-IS Designated Intermediate System (DIS) is responsible for generating and updating pseudonode LSPs on a broadcast network.

Unlike OSPF DR, the IS-IS DIS does not preempt by default. This behavior avoids unnecessary flapping in the network due to frequent DIS re-elections.

HCIP-Datacom-Core Reference

The characteristics of DIS and DR behavior are explained in IS-IS network operation chapters.

Question: 15

On an IS-IS network, routers send LSPs to exchange link state information. LSPs are classified into Level-1 LSPs and Level-2 LSPs and have the same format. Which of the following parts constitute the LSP ID in an LSP?

A. LSP Number

B. Pseudonode ID

C. System ID

D. IS Type

Answer: A, B, C

Explanation:

IS-IS Overview: Intermediate System to Intermediate System (IS-IS) is a link-state routing protocol. Routers exchange Link State Packets (LSPs) to maintain a synchronized link-state database. These LSPs are categorized into Level-1 LSPs (intra-area routing) and Level-2 LSPs (inter-area routing). Both types share the same packet format.

Reference: HCIP-Datacom-Core Technology Training Material (IS-IS LSPs and Packet Structure).

LSP ID Format: The LSP ID uniquely identifies each LSP and ensures accurate routing information. It comprises the following components:

System ID (C): A 6-byte identifier assigned to each router, derived from the router's NET (Network Entity Title). This identifier ensures unique identification of routers within the IS-IS domain.

Pseudonode ID (B): Assigned when a router acts as a Designated Intermediate System (DIS) on a broadcast network. It differentiates LSPs generated by the DIS from other routers.

LSP Number (A): A 1-byte field indicating the sequence number of the LSP. It helps distinguish multiple LSPs generated by the same router for the same level.

Reference: HCIP-Datacom Advanced Routing & Switching Technology (IS-IS LSP Format).

IS Type Exclusion:

IS Type (D) is not part of the LSP ID itself. It is a field within the IS-IS PDU that indicates the type of Intermediate System (Level-1, Level-2, or both) but does not contribute to the composition of the LSP ID.

Reference: HCIE-Datacom V1.0 Training Material (IS-IS Basics and Levels).

Conclusion: The LSP ID in IS-IS consists of System ID, Pseudonode ID, and LSP Number. These components uniquely identify each LSP within the IS-IS domain.

Question: 16

On an OSPF network, interfaces are classified into four types based on link layer protocols. Which of the following types can interfaces on an IS-IS network be classified into based on physical links?

A. P2P
B. Broadcast
C. P2MP
D. NBMA

Answer: A, B

Explanation:

IS-IS Interface Types

IS-IS interfaces are categorized based on physical link characteristics:

Point-to-Point (P2P): Direct connections between two routers.

Broadcast: Shared medium networks where multiple routers communicate.

Incorrect Options

C . P2MP and D. NBMA are not standard interface classifications in IS-IS.

HCIP-Datacom-Core Reference

IS-IS physical link classifications are elaborated in IS-IS link configuration sections.

Question: 17

Which of the following attributes must be carried when BGP sends route update messages?

A. MED
B. Next_Hop
C. AS_Path
D. Local_Preference

Answer: B, C

Explanation:

Mandatory BGP Attributes

Next_Hop: Specifies the next hop to reach the destination.

AS_Path: Lists the autonomous systems traversed, crucial for loop prevention and route selection.

Optional Attributes

MED (Multi-Exit Discriminator) and Local_Preference are optional attributes that aid in route preference but are not mandatory.

HCIP-Datacom-Core Reference

BGP attribute behavior and classifications are detailed in BGP path selection principles.

Question: 18

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:

Non-Client Definition in IBGP

In an IBGP setup with a Route Reflector (RR), a non-client is an IBGP peer that is neither the RR itself nor its client.

Non-clients must establish fully meshed IBGP connections with all other non-clients and the RR because IBGP rules prohibit route propagation between non-clients without a direct connection.

HCIP-Datacom-Core Reference

The behavior of non-clients in an RR topology is clearly outlined in the BGP implementation chapters.

Question: 19

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: 60

Explanation:

BGP Keepalive Message Behavior

Keepalive messages are used to maintain the Established state of a BGP peer relationship.

The Keepalive timer determines the frequency of these messages and defaults to 60 seconds, as per the BGP specification.

The Keepalive timer default value is covered in the BGP configuration and operational principles.

Question: 20

Which of the following statements is true about BGP?

A. If the export routing policy applied to a BGP peer changes, manual intervention is required so that the device resends Update messages to the peer.
B. IGP routes can be converted into BGP routes only through the network command.
C. A router cannot be configured with multiple BGP processes.
D. Open messages carry only the BGP header.

Answer: A

Explanation:

Export Routing Policy Changes

When an export routing policy is modified, BGP does not automatically resend affected routes. Manual intervention, such as a clear ip bgp command, is required to resend Update messages reflecting the new policy.

Incorrect Options

B . IGP routes can also be advertised into BGP using redistribution, not just the network command.

C . A router can be configured with multiple BGP processes using different AS numbers (multiinstance BGP).

D . Open messages carry additional parameters such as AS number, Hold Time, and Router ID, not just the header.

HCIP-Datacom-Core Reference

BGP policy and update behavior are detailed in the route control and redistribution chapters.

Question: 21

In BGP, Notification messages are used to request peers to resend routing information after routing policies are changed.

A. TRUE
B. FALSE

Answer: B

Explanation:

Notification Message Purpose

BGP Notification messages are used to report errors or terminate a connection. They do not request peers to resend routing information after routing policies are changed.

Routing updates following policy changes require manual resynchronization, not Notification messages.

HCIP-Datacom-Core Reference

The purpose and usage of Notification messages are discussed in the BGP operation chapters.

Question: 22

When a BGP device sends an Open message to establish a peer connection, which of the following information is carried?

A. Local AS number
B. Router ID
C. NLRI
D. Hold time

Answer: A, B, D

Explanation:

BGP Open Message Components

The Open message contains the following critical parameters:

Local AS Number: The autonomous system of the router.

Router ID: A unique identifier for the router.

Hold Time: The maximum time the router will wait for Keepalive or other messages from its peer.

Incorrect Option

C . NLRI: Network Layer Reachability Information is not included in Open messages; it is carried in Update messages.

HCIP-Datacom-Core Reference

The structure and contents of Open messages are explained in BGP protocol details.

Question: 23

During BGP route summarization configuration, the keyword can be used to suppress all specific routes so that only the summary route is advertised. The summary route carries the Atomic- aggregate attribute rather than the community attributes of specific routes.

Answer: suppress spec-map

Explanation:

Understanding BGP Route Summarization:

In Border Gateway Protocol (BGP), route summarization is a technique used to aggregate multiple specific prefixes into a broader summary prefix. This reduces the size of routing tables and improves routing efficiency.

Summarization helps to hide unnecessary details from other parts of the network while still maintaining connectivity.

Reference: HCIP-Datacom-Core Technology Training Material (BGP Route Summarization).

Suppressing Specific Routes:

When summarizing routes, the suppress-spec-map keyword is used to suppress specific prefixes so that only the summarized route is advertised.

The suppressed routes are not advertised to BGP peers; instead, only the summary route is propagated.

Reference: HCIP-Datacom Advanced Routing & Switching Technology (BGP Attributes and Summarization Techniques).

Atomic-Aggregate Attribute:

The summary route generated during BGP route summarization carries the Atomic-aggregate attribute. This attribute indicates that the summary route might not provide the exact path information available in the original specific routes.

Additionally, when using the suppress-spec-map option, the specific routes' attributes, such as community attributes, are not included in the summary route.

Reference: HCIP-Datacom-Core Technology Training Material (BGP Path Attributes).

Conclusion:

The suppress-spec-map keyword is used to suppress specific routes when performing BGP summarization. The summarized route includes the Atomic-aggregate attribute but does not carry community attributes from the suppressed routes.

Question: 24

BGP is generally applied to complex networks where routes change frequently. Frequent route flapping consumes a large number of bandwidth and CPU resources, and even affects the normal operation of the network. This is an unavoidable problem that cannot be solved in BGP.

A. TRUE
B. FALSE

Answer: B

Explanation:

BGP Route Flapping and Mitigation

While route flapping is a concern in large-scale networks, BGP employs mechanisms such as Route Dampening to mitigate its impact. Route dampening suppresses frequently flapping routes for a period of time to reduce resource consumption and network instability.

Therefore, it is incorrect to state that the issue cannot be resolved in BGP.

HCIP-Datacom-Core Reference

The mechanism of route dampening is detailed in the BGP optimization chapters.

Question: 25

The Next_Hop attribute in BGP records the next hop of a route. Similar to the next hop in an IGP, the Next_Hop attribute in BGP must be the IP address of a peer interface.

A. TRUE
B. FALSE

Answer: B

Explanation:

BGP Next_Hop Attribute

Unlike IGP, the Next_Hop attribute in BGP does not necessarily have to be the IP address of a peer interface. For example, in multi-hop BGP configurations, the Next_Hop can point to a different router or interface within the network.

HCIP-Datacom-Core Reference

Details of the Next_Hop attribute and its behavior are outlined in BGP path selection principles.

Question: 26

A BGP device receives a route carrying an unknown attribute from a peer but does not know whether other devices need the attribute. In this case, the BGP device retains this attribute when advertising the route to other peers. Which of the following attributes is of this type?

A. Community
B. AS.Path
C. MED
D. OriginatorID

Answer: A

Explanation:

Transitive vs. Non-Transitive Attributes

The Community attribute is a transitive optional attribute, meaning that if a router receives a route with this attribute

and does not understand its purpose, the router retains and propagates it to other peers.

Other options, such as AS_Path and MED, are well-defined mandatory or optional attributes with specific purposes.

HCIP-Datacom-Core Reference

The behavior of optional transitive attributes is detailed in the BGP protocol and attribute chapters.

Question: 27

When configuring an ACL on a router, you can specify a unique number or name to identify the ACL. Once a named ACL is created, it cannot be modified. You can only delete the named ACL and reconfigure it.

A. TRUE

B. FALSE

Answer: B

Explanation:

ACL Modification

Named ACLs can be modified after their creation. Unlike numbered ACLs, named ACLs provide greater flexibility for editing individual rules without deleting the entire ACL.

Therefore, the statement is incorrect.

HCIP-Datacom-Core Reference

The flexibility and editability of named ACLs are discussed in the ACL configuration sections.

Question: 28

ACLs are a common tool for matching routes. ACLs are classified into multiple types based on ACL rule functions. An ACL can be identified by a number, and the number range for each type of ACLs is different. Which of the following type of ACLs are numbered from 4000 to 4999?

A. User-defined ACL

B. Basic ACL

C. Layer 2 ACL

D. Advance ACL

Answer: A

Explanation:

ACL Number Ranges

Basic ACLs: Numbered from 2000 to 2999.

Advanced ACLs: Numbered from 3000 to 3999.

Layer 2 ACLs: Numbered from 5000 to 5999.

User-defined ACLs: Numbered from 4000 to 4999.

HCIP-Datacom-Core Reference

ACL classifications and their number ranges are covered in the ACL principles and configuration chapters.

Question: 29

ACLs are a common matching tool in routing policies. An ACL can be configured on a router to match routes.

A. TRUE

B. FALSE

Answer: A

Explanation:

ACL in Routing Policies

ACLs are frequently used in routing policies to match specific routes based on criteria such as source

IP, destination IP, and more. This allows ACLs to influence route redistribution, filtering, and forwarding decisions.

HCIP-Datacom-Core Reference

ACL applications in routing policies are discussed in the routing policy chapters.

Question: 30

Both MQC and PBR can be applied on device interfaces to filter received and sent packets or control packet forwarding paths.

A. TRUE
B. FALSE

Answer: A

Explanation:

MQC and PBR

MQC (Modular QoS Command-Line Interface): Applied to interfaces for classifying and controlling traffic.

PBR (Policy-Based Routing): Used to influence packet forwarding based on policies rather than traditional routing tables.

Both MQC and PBR can be configured on device interfaces to filter incoming/outgoing packets or control their forwarding paths.

HCIP-Datacom-Core Reference

The use of MQC and PBR on device interfaces is elaborated in the QoS and routing control chapters.

Question: 31

An IP prefix list is a common matching tool used in routing policies. Which of the following cannot be configured as matching conditions in an IP prefix list on a Huawei router?

A. Port number
B. Mask
C. Action
D. Index

Answer: A

Explanation:

IP Prefix List Matching Conditions

An IP prefix list matches based on:

Mask: Specifies the subnet mask length.

Action: Specifies whether to permit or deny.

Index: Orders the rules within the prefix list.

Port numbers are not applicable as matching conditions in an IP prefix list.

HCIP-Datacom-Core Reference

IP prefix list configurations are detailed in the routing policy and route filtering chapters.

Question: 32

When receiving a packet, a Huawei router matches the packet against ACL rules. The default ACL matching order used by the Huawei router is.

Answer: Sequential

Explanation:

ACL Matching in Huawei Routers:

Access Control Lists (ACLs) are used to filter packets based on specific criteria, such as source/destination IP, ports, or protocols.

When a packet arrives, the Huawei router processes it against the configured ACL rules to decide whether to permit or deny the packet.

Reference: HCIP-Datacom-Core Technology Training Material (ACL Principles and Configuration).

Sequential Matching Order:

By default, Huawei routers match packets against ACL rules sequentially. This means:

The router checks the packet against rules in the order they are listed, starting from the top of the ACL.

The first rule that matches the packet's attributes is applied, and no further rules are checked. This is known as the first-match principle.

If no rules match, the packet is denied by default (implicit deny).

Reference: HCIP-Datacom Advanced Routing & Switching Technology (ACL Matching Mechanism).

Example of Sequential Matching:

Consider the following ACL rules:

Rule 10: Permit IP 192.168.1.0/24

Rule 20: Deny IP 192.168.1.1

If a packet with source IP 192.168.1.1 arrives:

The router matches it against Rule 10 (Permit 192.168.1.0/24) and allows the packet.

Rule 20 is not evaluated because the first match (Rule 10) already applies.

Alternative Matching Orders:

Some routers or configurations allow batch matching (evaluating all rules) for specific scenarios, but this is not the default behavior in Huawei routers.

Conclusion:

The default ACL matching order on Huawei routers is sequential, and the first matching rule determines the action applied to the packet.

Question: 33

DRAG DROP

On an STP network, the root bridge, root port, and designated port are elected in sequence. The election rules of these ports are different. List the steps for electing the root port in sequence.

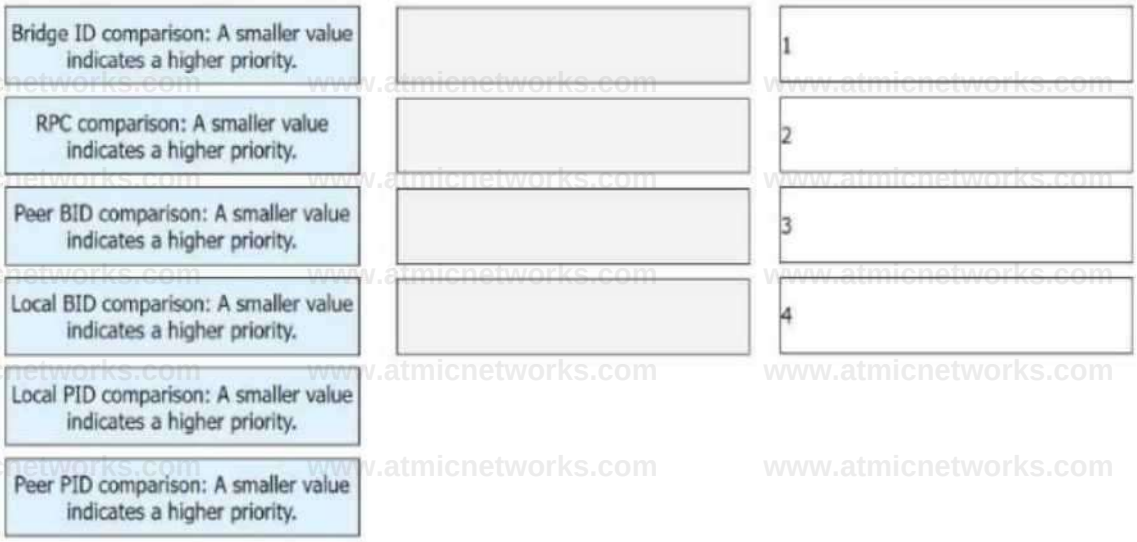

Answer:

Explanation:

The sequence of steps for electing the root port in an STP (Spanning Tree Protocol) network is as follows:

Bridge ID Comparison: The Bridge ID (BID) is compared between the bridges in the network. A smaller value indicates a higher priority, meaning the bridge with the lowest Bridge ID is elected as the root bridge.

RPC (Root Path Cost) Comparison: The path cost to reach the root bridge is calculated. The router with the lowest Root Path Cost (RPC) to the root bridge will have a higher priority for the election of the root port.

Peer BID Comparison: If there is a tie in the Root Path Cost, the Peer BID is compared. A smaller Peer BID indicates a higher priority. This step ensures that if two routers have the same RPC, the one with the lower Peer Bridge ID wins.

Local BID Comparison: If there is still a tie, the Local BID is compared. A smaller Local BID indicates a higher priority. This final step ensures that the router with the lowest local identifier is selected.

Bridge ID Comparison:

The first step in electing the root port is comparing the Bridge IDs. The bridge with the lowest Bridge ID becomes the root bridge. The Bridge ID is made up of the bridge priority and MAC address. The root bridge is the center of the network for STP, and all other ports will calculate their paths based on this root.

Reference: HCIP-Datacom-Core Technology Training Material (STP Concepts and Election Process).

RPC (Root Path Cost) Comparison:

Once the root bridge is selected, the network needs to determine the best path to the root. Each port on a non-root bridge will calculate the Root Path Cost (RPC), which is the cumulative cost of reaching the root bridge from that port. The root port is the one that has the lowest RPC, meaning it provides the best path to the root bridge.

Reference: HCIP-Datacom-Core Technology Training Material (STP Path Selection).

Peer BID Comparison:
If multiple paths have the same Root Path Cost, the next step is to compare the Peer Bridge IDs. The bridge with the lowest Peer BID is chosen as the root port. This ensures a tie-breaking mechanism based on the neighbor's identifier.

Reference: HCIP-Datacom-Core Technology Training Material (STP Election Process).

Local BID Comparison:
If there is still a tie after comparing the Peer Bridge IDs, the Local Bridge ID is compared. A smaller Local BID indicates a higher priority, and the port with the lower Local BID will be selected as the root port.

Reference: HCIP-Datacom-Core Technology Training Material (STP Local Port Selection).

Question: 34

Compared with RSTP, which of the following port roles are added to MSTP?

A. Backup port
B. Master port
C. Edge port
D. Regional edge port

Answer: B, D

Explanation:

Additional Port Roles in MSTP

Master Port: Indicates the port on the shortest path to the root bridge in a region.

Regional Edge Port: Identifies a port at the boundary of the MST region.

Backup and edge ports exist in RSTP and are not newly introduced by MSTP.

HCIP-Datacom-Core Reference

MSTP port roles are elaborated in the MSTP configuration sections.

Question: 35

On an RSTP network, if a port receives an RST BPDU and finds that its buffered RST BPDU is superior to the received RST BPDU, the port discards the received RST BPDU without responding.

A. TRUE

B. FALSE

Answer: A

Explanation:

RST BPDU Handling

On an RSTP network, if a port receives an RST BPDU and determines its own buffered BPDU is superior, it discards the received BPDU without responding. This ensures stability and proper convergence in the network.

HCIP-Datacom-Core Reference

BPDU handling is described in the RSTP operation chapters.

Question: 36

An edge port is a new port role added to RSTP to overcome the disadvantages of STP. Which of the following statements is false about this port role?

A. The port does not participate in RSTP calculation.

B. The port can directly enter the Forwarding state from the Discarding state.

C. After receiving a configuration BPDU, the port is still in the Forwarding state.

D. The Up and Down states of the port do not cause network topology changes.

Answer: C

Explanation:

Edge Port Behavior

An edge port transitions directly to the Forwarding state but reverts to participating in RSTP calculations if it receives a valid configuration BPDU. Thus, statement C is false.

HCIP-Datacom-Core Reference

The behavior of edge ports and their transitions are discussed in the RSTP enhancement sections.

Question: 37

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:

Permanent Group Addresses for Routing Protocols

The range 224.0.0.0 to 224.0.0.255 is reserved for local network protocols, including routing protocols. Examples include OSPF (224.0.0.5/6) and RIP (224.0.0.9).

HCIP-Datacom-Core Reference

Reserved multicast address ranges are detailed in the multicast configuration sections.

Question: 38

Based on IGMP snooping, IGMP snooping proxy enables a switch to act as a substitute for an upstream Layer 3 device to send IGMP Query messages to downstream hosts, and also to act as a substitute for downstream hosts to send IGMP Report/Leave messages to an upstream device. As such, this function conserves bandwidth between the upstream device and the local device.

A. TRUE
B. FALSE

Answer: A

Explanation:

GMP Snooping Proxy

IGMP snooping proxy allows the switch to substitute for upstream Layer 3 devices and downstream hosts, reducing unnecessary bandwidth consumption by consolidating IGMP messages.

HCIP-Datacom-Core Reference

The IGMP snooping proxy functionality is described in multicast optimization and IGMP configuration chapters

Question: 39

There are two types of routing entries on a PIM network. (S, G) routing entries are used to set up on a PIM network and are applicable to both PIM-DM and PIM-SM networks.

Answer: SPT

Explanation:

Routing Entries in PIM

(S, G) entries: Represent the shortest path tree (SPT) from a specific source (S) to a specific multicast group (G). These entries are created after data packets are transmitted directly from the source to the receivers.

(S, G) entries are used in both PIM-DM and PIM-SM modes.

PIM-DM and PIM-SM Modes

In PIM-DM, multicast traffic is initially flooded throughout the network and pruned where no receivers exist. (S, G) entries represent source-specific paths created after the flooding stage.

In PIM-SM, multicast traffic initially flows through a shared tree, but (S, G) entries are created when the network switches to the SPT for more efficient forwarding.

HCIP-Datacom-Core Reference

The functionality and application of (S, G) routing entries are detailed in multicast and PIM configuration chapters.

Question: 40

DRAG DROP

Match the following IPv4 multicast protocols with the corresponding functions.

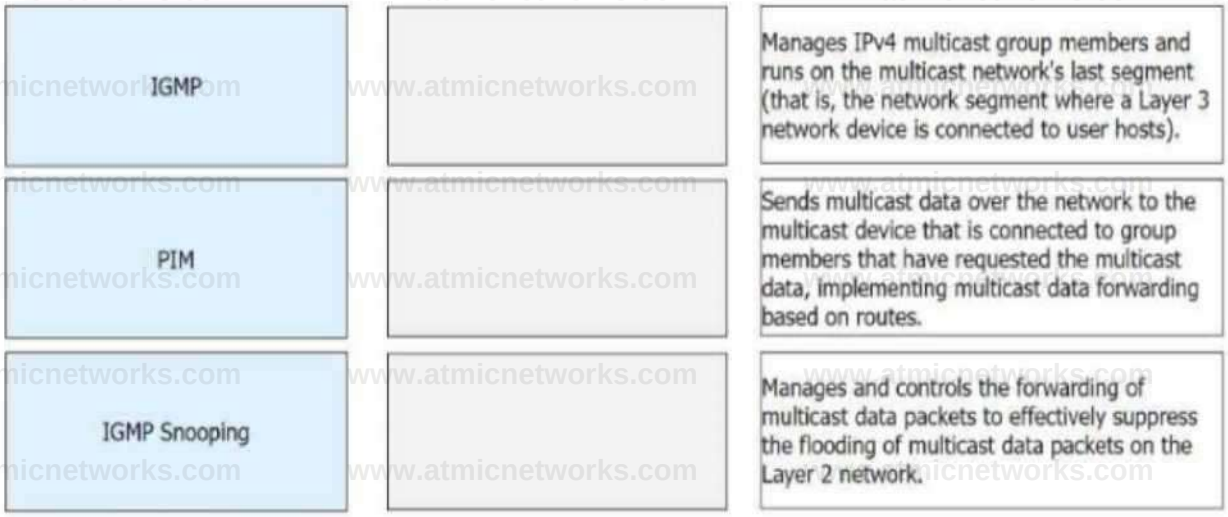

Answer:

Explanation:

IGMP: Manages IPv4 multicast group members and runs on the multicast network's last segment (that is, the network segment where a Layer 3 network device is connected to user hosts).

PIM: Sends multicast data over the network to the multicast device that is connected to group members that have requested the multicast data, implementing multicast data forwarding based on routes.

IGMP Snooping: Manages and controls the forwarding of multicast data packets to effectively suppress the flooding of multicast data packets on the Layer 2 network.

IGMP (Internet Group Management Protocol):

IGMP is used by hosts and adjacent multicast routers to establish and maintain multicast group memberships on a local subnet.

It allows a host to inform a multicast router of its desire to receive multicast traffic for a specific group.

It operates at Layer 3 and runs on the last segment of the network.

Reference: HCIP-Datacom-Core Technology Training Material (Multicast Fundamentals).

PIM (Protocol Independent Multicast):

PIM is a routing protocol designed for efficient routing of IP multicast traffic. It determines the paths over which multicast packets should be forwarded.

PIM is "protocol-independent" because it uses the underlying unicast routing table for RPF (Reverse Path Forwarding) checks.

It is responsible for distributing multicast data throughout the network.

Reference: HCIP-Datacom-Core Technology Training Material (Multicast Routing Protocols).

IGMP Snooping:

IGMP Snooping operates at Layer 2 and monitors IGMP traffic between hosts and routers.

It prevents multicast flooding by allowing a switch to forward multicast packets only to the ports that have joined specific multicast groups.

This significantly enhances efficiency in Layer 2 networks with multicast traffic.

Reference: HCIP-Datacom-Core Technology Training Material (IGMP Snooping Concepts).

Question: 41

Without a prior version check, an engineer configures IGMP snooping on a device and the version of IGMP snooping is earlier than the IGMP versions on user hosts. In this case, which of the following situations will occur?

A. Users cannot receive multicast data because the device forwards received IGMP Report messages only to router ports and does not generate group member ports or forwarding entries.

B. Users cannot receive multicast data, but the device generates forwarding entries after receiving IGMP Report messages.

C. The IGMP snooping version of the device is automatically degraded, and users can receive multicast data properly.

D. The IGMP versions of the hosts are automatically upgraded, and users can receive multicast data properly.

Answer: A

Explanation:

IGMP Version Mismatch

If the IGMP snooping version on the device is earlier than the IGMP version on user hosts, the device may fail to parse IGMP Report messages correctly. As a result, the device forwards these messages only to router ports without generating group member ports or forwarding entries.

Consequently, users cannot receive multicast data.

HCIP-Datacom-Core Reference

Multicast and IGMP snooping behaviors under mismatched conditions are described in the multicast configuration chapters.

Question: 42

The Neighbor Discovery Protocol (NDP) is an important basic protocol in the IPv6 protocol suite and plays an important role. Which of the following functions and features does it support?

A. Address resolution
B. Neighbor state tracing
C. Duplicate address detection
D. Redirection

Answer: A, C, D B

Explanation:

Neighbor Discovery Protocol (NDP):

NDP is a key protocol in the IPv6 protocol suite, replacing ARP (Address Resolution Protocol) in IPv4.

It operates using ICMPv6 (Internet Control Message Protocol for IPv6) and is critical for managing interactions between IPv6 nodes on the same link.

Reference: HCIP-Datacom-Core Technology Training Material (NDP Functions and Features).

NDP Features and Functions:

A . Address resolution:

NDP resolves IPv6 addresses into MAC addresses, similar to ARP in IPv4. It uses Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages for this purpose.

Reference: HCIP-Datacom-Core Technology Training Material (Address Resolution in IPv6).

B . Neighbor state tracing:

NDP tracks the state of neighbors to determine their reachability. It maintains a neighbor cache and uses NS/NA messages to verify whether neighbors are reachable.

Reference: HCIP-Datacom-Core Technology Training Material (Neighbor Reachability in IPv6).

C . Duplicate address detection (DAD):

NDP ensures that IPv6 addresses are unique within a network. Before assigning an address to an interface, DAD is used to verify that no other node is using the same address. This is done via NS messages.

Reference: HCIP-Datacom-Core Technology Training Material (IPv6 Address Assignment and DAD).

D . Redirection:

NDP provides redirection functionality to inform hosts of a better first-hop router for reaching a particular destination. It uses ICMPv6 Redirect messages for this purpose.

Reference: HCIP-Datacom-Core Technology Training Material (NDP Redirect Functionality).

Conclusion:

NDP supports all the mentioned functions and features: Address resolution, Neighbor state tracing, Duplicate address detection, and Redirection, making it indispensable for IPv6 networks.

Question: 43

IPv6 defines multiple types of addresses. Which of the following statements is false about these addresses?

A. Link-local addresses can be quickly generated using the EUI-64 method.

B. Anycast addresses can be used only as destination addresses.

C. Each interface can have multiple global unicast addresses with different network prefixes.

D. Manually configured link-local addresses have a higher priority than automatically generated ones.

Answer: D

Explanation:

IPv6 Address Behavior

Manually configured link-local addresses do not have a higher priority than automatically generated ones; they coexist and are equally preferred for local communication.

The other statements are true:

A . Link-local addresses can use the EUI-64 method for quick generation.

B . Anycast addresses are only used as destination addresses.

C . An interface can have multiple global unicast addresses with different prefixes.

HCIP-Datacom-Core Reference

IPv6 address types and priority behaviors are detailed in the IPv6 addressing sections.

Question: 44

There are various types of VPNs, which can be applied to different layers. Which of the following network layers does SSL VPN belong to?

A. Network layer

B. Application layer

C. Transport layer

D. Data link layer

Answer: B

Explanation:

SSL VPN and Its Functionality:

SSL VPN (Secure Sockets Layer Virtual Private Network) provides secure remote access to a network using SSL/TLS protocols.

SSL VPN operates at the Application Layer of the OSI model. It enables secure communication for applications like

web browsers, email clients, and file sharing.

Unlike IPsec VPN, which operates at the Network Layer, SSL VPN focuses on application-specific encryption and authentication.

Reference: HCIP-Security Training Material (VPN Basics and SSL VPN Configuration).

Question: 45

GRE is a VPN encapsulation technology that is widely used to transmit packets across heterogeneous networks. Which of the following statements is false about GRE?

A. GRE supports encryption and authentication.
B. GRE supports multicast transmission.
C. GRE is a Layer 3 VPN encapsulation technology.
D. GRE can work with other VPN protocols to better ensure data security.

Answer: A

Explanation:

GRE Characteristics

GRE does not inherently support encryption or authentication. It is a tunneling protocol for encapsulating packets, and data security features must be implemented using other protocols such as IPsec.

Other correct attributes of GRE include:

B . Supports multicast transmission.

C . Acts as a Layer 3 VPN encapsulation technology.

D . Can work with VPN protocols like IPsec for better security.

HCIP-Datacom-Core Reference

GRE features and limitations are discussed in VPN encapsulation technology chapters.

Question: 46

By default, some security zones are created when Huawei firewalls are enabled. Which of the following security zones is created by users?

A. DMZ

B. ISP

C. Trust

D. Local

Answer: A

Explanation:

By default, Huawei firewalls create security zones such as Trust, Untrust, and Local. The DMZ (Demilitarized Zone) is a security zone explicitly created by users. A DMZ is used to isolate an internal network from the external one, providing an additional layer of security by placing public-facing services (e.g., web servers) in this intermediary zone. This setup ensures that if a public-facing service is compromised, the internal network remains secure. Huawei Firewall configuration steps confirm this zoning principle, making DMZ creation an explicit user-driven action .

Question: 47

When receiving a packet that does not match any session table entry, the firewall discards the packet to prevent external attacks and ensure internal information security.

A. TRUE

B. FALSE

Answer: A

Explanation:

When a Huawei firewall receives a packet that does not match any existing session table entry, it discards the packet. This is part of the default firewall policy, which ensures that unrecognized traffic is treated as a potential security risk and blocked. This behavior is vital for preventing unauthorized access and mitigating external attacks. The feature aligns with Huawei's default security strategies as detailed in their firewall operation manuals .

Question: 48

GRE is a Layer 2 VPN encapsulation technology that encapsulates packets of certain data link layer protocols so that the encapsulated packets can be transmitted over an IP network.

A. TRUE
B. FALSE

Answer: B

Explanation:

GRE (Generic Routing Encapsulation) is not a Layer 2 VPN technology. Instead, it is a Layer 3 tunneling protocol used to encapsulate a wide variety of network layer protocols inside point-to- point connections. GRE is commonly used for creating VPN tunnels across IP networks, allowing for the transport of various types of payloads. This misunderstanding about GRE being a Layer 2 technology contradicts its definition and typical application .

Question: 49

When multiple access channels are set for the same access requirement, the insecure access channels are not used and secure access channels are selected in normal cases. Which of the following are secure access channels?

A. HTTPS
B. Telnet
C. SNMPv2
D. SFTP

Answer: A, D

Explanation:

Secure access channels include protocols that encrypt the transmitted data to protect against interception or unauthorized access. HTTPS (HyperText Transfer Protocol Secure) ensures data encryption over web communications, while SFTP (Secure File Transfer Protocol) provides secure file transfer by utilizing SSH for data encryption. Telnet and SNMPv2, on the other hand, lack robust encryption and are considered insecure. Huawei security standards highlight the importance of encrypted communication to prevent data leaks .

Question: 50

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, D B

Explanation:

For single-hop BFD (Bidirectional Forwarding Detection), the configuration must include the local and remote discriminators to uniquely identify the session endpoints. Additionally, enabling BFD globally is a prerequisite for initiating BFD sessions. Configuring multicast IP addresses is unnecessary for single-hop BFD, as it operates over direct links. Huawei's configuration guidelines specify these requirements to ensure effective deployment and operation of BFD.

Understanding BFD (Bidirectional Forwarding Detection):

BFD is a protocol used to detect link faults quickly between two routers.

Single-hop BFD operates on directly connected links and is commonly used for fast fault detection in routing protocols like OSPF and BGP.

Reference: HCIP-Datacom-Core Technology Training Material (BFD Basics and Configuration).

Mandatory Configurations for Single-Hop BFD:

A. Configure the remote discriminator of a BFD session:

The remote discriminator is used to uniquely identify the BFD session at the remote end. This is essential for session establishment.

B. Configure the local discriminator of a BFD session:

The local discriminator uniquely identifies the BFD session at the local end. This is required to establish a BFD session.

D. Enable BFD globally:

BFD must be enabled globally on the router for the protocol to operate and for session configurations to take effect.

Optional Configuration:

C. Configure a multicast IP address for BFD:

This is not required for single-hop BFD, as it operates over direct links using unicast communication. Multicast is used in other scenarios, like multi-hop BFD.

Conclusion:

The correct configurations for single-hop BFD are A, B, and D.

Question: 51

As shown in the figure, VRRP is associated with a BFD session. When a backup device detects a fault through BFD, the backup device immediately assumes the master role after the Master_Down_Timer timer expires.

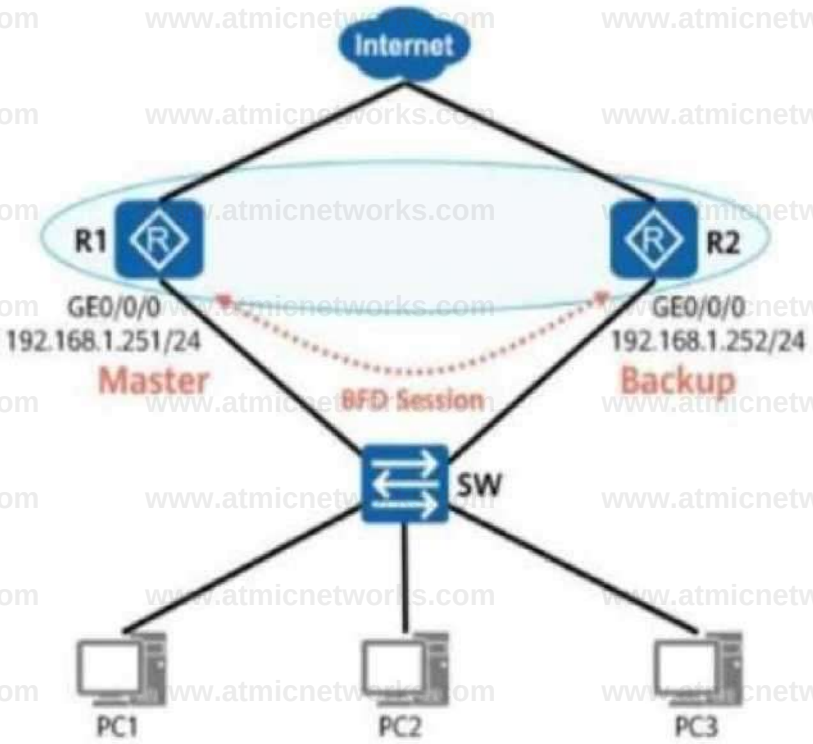

A. TRUE

B. FALSE

Answer: B

Explanation:

VRRP and BFD Association:

VRRP (Virtual Router Redundancy Protocol) is used to provide gateway redundancy by electing a master and backup router.

Associating VRRP with BFD (Bidirectional Forwarding Detection) allows faster detection of faults on the master device or the link between the master and the backup.

Reference: HCIP-Datacom-Core Technology Training Material (VRRP and BFD Association).

Master_Down_Timer Behavior:

Normally, when a VRRP backup device detects that the master is down (through missed VRRP advertisements), the Master_Down_Timer dictates the failover timing.

However, when VRRP is associated with BFD:

BFD detects faults immediately (sub-second detection).

The backup device does not wait for the Master_Down_Timer to expire. Instead, it immediately assumes the master role upon fault detection by BFD.

Reference: HCIP-Datacom-Core Technology Training Material (VRRP Failover Mechanism).

Why the Statement is FALSE:

The statement claims that the backup device assumes the master role after the Master_Down_Timer expires when a fault is detected by BFD.

This is incorrect because BFD bypasses the need for the Master_Down_Timer to expire. The backup device transitions to the master role immediately upon BFD detecting a fault.

Conclusion:

The correct behavior of VRRP when associated with BFD is immediate role assumption by the backup device upon fault detection, bypassing the Master_Down_Timer.

Therefore, the statement is FALSE.

Question: 52

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:

In asynchronous mode, two devices exchange BFD Control packets to monitor connectivity. However, this mode does support the echo function, which tests bidirectional paths by sending packets that loop back to the origin. This statement is incorrect. Demand mode, on the other hand, minimizes the control packet exchange, relying on periodic

echo packets if configured. Huawei's BFD configuration guides emphasize these distinctions .

Question: 53

BFD can quickly detect faults in channels at multiple network layers, ensuring high reliability. To which of the following layers does BFD belong?

A. Application layer
B. Data link layer
C. Physical layer
D. Network layer

Answer: D

Explanation:

BFD operates at the network layer, providing rapid fault detection for forwarding paths, such as IP routing and MPLS tunnels. It interacts closely with network protocols (e.g., OSPF, BGP) to ensure high reliability and quick response to faults. BFD does not belong to the application, data link, or physical layers but functions as a network layer diagnostic tool as confirmed in Huawei's network protocol training .

Question: 54

In VRRP networking, if VRRP is not configured to track an uplink interface and the uplink interface or link of the master device in a VRRP group fails, no switchover will be triggered. As a result, a traffic blackhole occurs.

A. TRUE
B. FALSE

Answer: A

Explanation:

If VRRP is not configured to track uplink interfaces, a failure in the master device's uplink or link will not trigger a switchover, resulting in a traffic blackhole. The VRRP mechanism relies on interface tracking to monitor connectivity and ensure role transitions upon faults. Without this configuration, no failover occurs, and traffic directed toward

the master device is lost .

Question: 55

In the SNMP management model, which of the following elements defines the attributes of a managed device?

A. MIB
B. Agent
C. Managed Object
D. NMS

Answer: A

Explanation:

In SNMP (Simple Network Management Protocol), the MIB (Management Information Base) defines the attributes of a managed device. It acts as a structured database of managed objects, each described by a unique OID (Object Identifier). These objects provide critical information and control interfaces for network management. Other elements like the agent (manages the MIB) and NMS (queries the agent) do not define attributes but interact with the MIB .

Question: 56

In inter-AC roaming scenarios, an AC can function as the mobility server of multiple mobility groups, but can be added only to one mobility group.

A. TRUE
B. FALSE

Answer: A

Explanation:

In inter-AC roaming scenarios, an AC (Access Controller) can serve as the mobility server for multiple mobility groups, enabling it to manage roaming among multiple groups. However, an AC can only belong to one specific mobility group. This constraint ensures that mobility management remains streamlined and avoids conflicts. Huawei WLAN mobility configuration guides validate this setup .

Question: 57

A large shopping mall configures a VLAN pool to prevent network performance deterioration caused by potentially large broadcast domains. A network engineer runs the display vlan pool name STA command to check information about the VLAN pool. The following command output is displayed:

<AC> display vlan pool name STA Name : STA Total : 2 Assignment: hash Threshold Notify Count: 3 Threshold Notify time(min): 3 VLAN ID : 2 4

Which of the following statements are true?

A. The VLANs with the IDs of 2 and 4 are added to the VLAN pool.
B. The total number of VLAN pools is 2.
C. The name of the VLAN pool is STA.
D. The VLAN pool uses the even VLAN assignment algorithm.

Answer: A, C

Explanation:

The command output confirms that the VLANs with IDs 2 and 4 are part of the VLAN pool named "STA." The VLAN pool concept is used to allocate VLANs dynamically to devices or subnets, which reduces broadcast domain size and improves network efficiency. The total number of VLAN pools is unrelated to this output, and the assignment algorithm (hash or even) is not explicitly mentioned in the output. Huawei VLAN pool management references align with this analysis .

Question: 58

STAs stay on different subnets before and after Layer 3 roaming. To enable the STAs to access the original network after roaming, network engineers need to configure user traffic to be forwarded to the original subnet over a CAPWAP tunnel.

A. TRUE
B. FALSE

Answer: A

Explanation:

In Layer 3 roaming scenarios, STAs (stations) remain on different subnets before and after roaming. To ensure that user traffic reaches the original subnet, it must be tunneled back over a CAPWAP tunnel. This technique prevents disruptions in communication and enables seamless roaming. Huawei's Layer 3 roaming design guidelines emphasize the use of CAPWAP tunnels for user traffic forwarding .

Question: 59

The typical characteristics of the AI era are that it focuses on data, explores data value, and improves AI efficiency. Therefore, the core requirement of AI for data center networks is speed, meaning low latency.

A. TRUE

B. FALSE

Answer: A

Explanation:

AI-driven networks focus on maximizing efficiency by minimizing latency to enhance data processing and decision-making speed. The AI era emphasizes leveraging large data volumes and ensuring rapid data center interconnectivity. This makes low latency a critical requirement for AI-ready data center networks, as highlighted in Huawei's AI and data center network optimization materials .

Question: 60

Which of the following statements about the forwarding plane of a switch is false?

A. Provides high-speed and non-blocking data channels.

B. Can encapsulate and decapsulate packets.

C. Can collect packet statistics.

D. Consists of main control boards and interface boards.

Answer: D

Explanation:

The forwarding plane of a switch consists of data forwarding hardware, such as line cards or forwarding engines, and is responsible for tasks like encapsulating/decapsulating packets, providing high-speed data channels, and collecting packet statistics. However, main control boards are part of the control plane, not the forwarding plane. This distinction ensures a separation of data forwarding and control functionalities .

Question: 61

A monitoring plane usually comprises the monitoring units of main control boards and interface boards. This plane can monitor the system environment independently. Which of the following environment monitoring functions can be provided by the monitoring plane?

A. Voltage monitoring

B. Temperature monitoring

C. Fan control

D. System power-on and power-off control

Answer: A B C D

Explanation:

The monitoring plane's primary role is environmental monitoring, ensuring the stability of the system. Functions like voltage monitoring, temperature tracking, fan speed control, and power management are standard. These are independent of the forwarding or control planes and critical for maintaining device health in network operations .

Question: 62

Which of the following statement regarding the display ospf peer command output is true?

```
<Huawei>display ospf peer

                OSPF Process 1 with Router ID 10.1.1.2
                    Neighbors
  Area 0.0.0.0 interface 10.1.1,2(GigabitEthernet1/0/0)'s neighbors

  Router ID: 10.1.1.1 Address: 10.1.1.1
    State: Full Mode:Nbr is Slave Priority: 1
    DR: 10.1.1.1 BDR: None MTU: 0
    Dead timer due in 38 sec
    Retrans timer interval: 5
    Neighbor is up for 00:00:04
    Authentication Sequence: [ 0 ]
```

A. Address: 10.1,1.1 Indicates that the local interface address is 10.1.1.1.

B. Through negotiation during DD packet exchange, the local end becomes the slave.

C. Router ID indicates that the local router ID Is 10.1.1.1.

D. The DR address is 10.1.1.1

Answer: B

Explanation:

The command output indicates that the OSPF neighbor state is Full and that the neighbor relationship has been established. The statement confirms that the negotiation process during the exchange of DD (Database Description) packets has determined the role of the router, with the local device becoming the slave in the Master-Slave relationship, which is critical for LSA synchronization. The other options do not align with the output or OSPF

principles

Question: 63

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 (Link State Acknowledgment) packets are used to acknowledge received LSAs and do not contain complete LSA information. Instead, they contain only the headers of LSAs being acknowledged. This contrasts with LS Update packets, which carry full LSA details. The misunderstanding of LS Ack functionality makes this statement false .

Question: 64

To prevent inter-area routing loops, OSPF does not allow advertising routing information between two non-backbone areas and allows sing routing information only within an area or between the backbone area and a non-backbone area. Therefore, each ABR must be connected to the backbone area.

A. TRUE

B. FALSE

Answer: A

Explanation:

OSPF (Open Shortest Path First) mandates that the backbone area (Area 0) serves as the central area for routing information exchange. To avoid inter-area routing loops, routing information is only exchanged between the backbone area and non-backbone areas, not directly between two nonbackbone areas. Consequently, each Area Border Router (ABR) must connect to the backbone area to facilitate this routing exchange .

Question: 65

Which of the following statements regarding OSPF route summarization is false?

A. OSPF supports two route summarization modes: ABR summarization and ASBR summarization.

B. Any router in OSPF can summarize routes.

C. Route summarization is the process of summarizing routes with the same prefix into one route and then advertising only the summarized route to other areas.

D. Route summarization can reduce routing information, decrease the routing table size, and improve router performance.

Answer: B

Explanation:

Route summarization in OSPF is restricted to specific routers: Area Border Routers (ABRs) and Autonomous System Boundary Routers (ASBRs). These devices perform summarization to reduce routing table size and minimize advertised routing information between areas or across AS boundaries. Regular routers cannot perform route summarization, making this statement false .

Question: 66

Which of the following statements is false?

A. If the current DR fails, the current BDR automatically becomes a new DR, and a BDR will be elected again.

B. A device with a higher router priority has a higher election priority.

C. When a router with the highest router priority joins an OSPF network, this router will become the new DR.

D. If two devices have the same router priority, the device with a larger router ID has a higher election priority.

Answer: C

Explanation:

The election of a Designated Router (DR) and Backup Designated Router (BDR) occurs during the OSPF network's initial setup or upon the failure of an existing DR. A new router joining an OSPF network does not automatically preempt the current DR, even if it has the highest priority. DR and BDR elections are not preemptive. The other statements are correct regarding OSPF DR/BDR election processes .

Question: 67

Which of the following statements regarding the LSA age field are true?

A. The unit of this field is seconds. In a LSDB, the LS age of a LSA increases with time.

B. If the LS age of a LSA has reached the LSRefreshTime (30 minutes), any router can regenerate an instance of this LSA again.

C. The unit of this field is seconds. In a LSDB, the LS age of a LSA decreases with time.

D. If the LS age of a LSA has reached the LSRefreshTime (30 minutes), the originator of this LSA needs to regenerate an instance of this LSA again.

Answer: A D

Explanation:

The LS age field in OSPF LSAs (Link State Advertisements) is measured in seconds and increments over time in the Link State Database (LSDB). If the LS age reaches the LSRefreshTime (30 minutes), the router that originated the LSA must regenerate it to keep the LSA valid in the network. Options C and B are incorrect since LS age does not decrease, and only the originator regenerates LSAs .

Question: 68

In the OSPF protocol, inter-area route calculation involves only Router LSA, Network LSA, and Summary LSA.

A. TRUE

B. FLASE

Answer: B

Explanation:

Inter-area route calculation in OSPF involves Summary LSAs (Type 3) for summarizing routes across areas and ASBR Summary LSAs (Type 4) for locating ASBRs. While Router LSAs and Network LSAs describe the internal structure of an area, they do not directly participate in inter-area route calculation. Thus, the statement that only Router, Network, and Summary LSAs are involved is false .

Question: 69

Which of the following statements regarding DR/BDR are false?

A. In a broadcast network, DR and BDR must be elected. A broadcast network without a DR or BDR cannot operate normally.

B. DR others listen on the multicast address 224.0.0.5.

C. All DR others establish neighbor relationships with DR and BDR only.

D. DR others listen on the network address 224.0.0.6.

Answer: D

Explanation:

In OSPF, DR and BDR listen on the multicast address 224.0.0.6, while all OSPF routers (including DR others) listen on 224.0.0.5. DR others establish neighbor relationships with DR and BDR, but they do not listen on 224.0.0.6. This makes Option D incorrect, as only DR and BDR use 224.0.0.6 .

Question: 70

When two routers exchange LSDB information using DD packets, a master/slave relationship is formed first, the router with a larger router ID is the master, and determine the MS bit.

A. TRUE

B. FALSE

Answer: A

Explanation:

When OSPF routers exchange LSDBs using Database Description (DD) packets, they first establish a master-slave relationship. The router with the larger Router ID becomes the master, setting the MS (Master/Slave) bit in the DD packets. This hierarchical relationship ensures an orderly exchange of routing information .

Question: 71

Which of the following TLVs is used by ISIS to describe the IP address of an interface?

A. 129

B. 131

C. 128

D. 132

Answer: A

Explanation:

In the IS-IS protocol, TLV 129 (IPv4 interface address) is used to describe the IP address of an interface. Each TLV type carries specific information, and TLV 129 specifically relates to interface IP addresses .

Question: 72

OSPF supports area authentication and interface authentication. If both authentication modes are configured, Interface authentication takes preference over area authentication.

A. TRUE

B. FALSE

Answer: A

Explanation:

OSPF supports both area-level and interface-level authentication. When both are configured, interface authentication takes precedence over area authentication. This hierarchy ensures that specific interface-level configurations override the broader area-level settings when applicable .

Question: 73

A local router runs IS-IS and its command output is shown in the following figure. Which of the following statements is true?

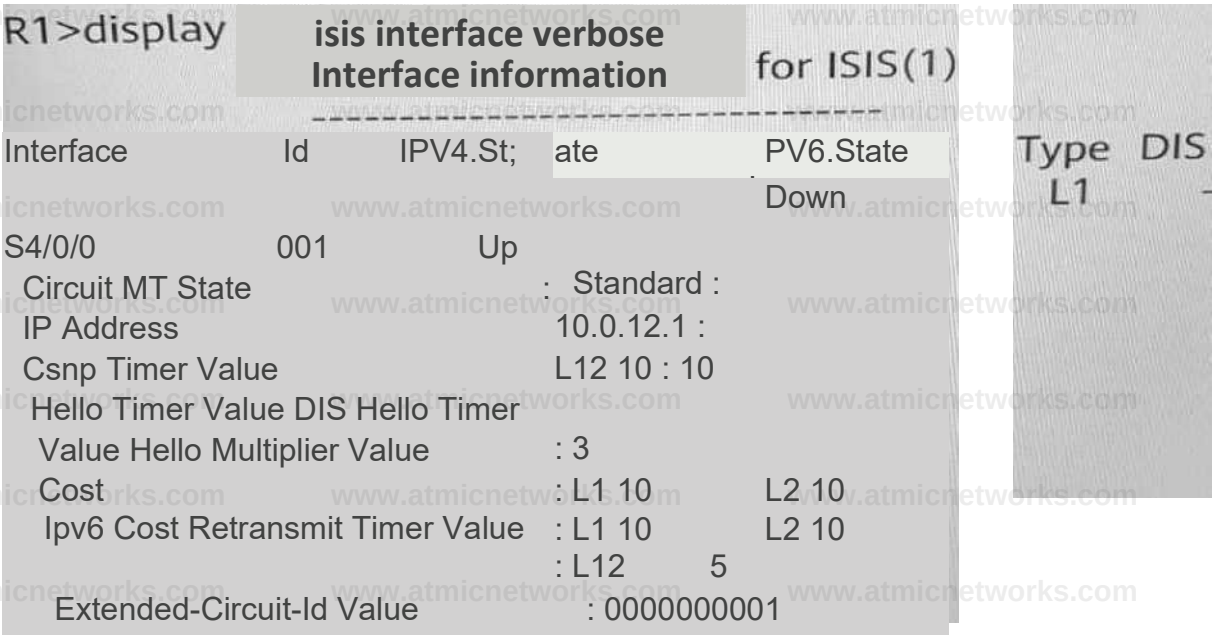

A. The circuit level of S4/0/0 is Level-1.

B. S4/0/0 supports IPv6.

C. S4/0/0 sends IIH packets at the interval of 30s.

D. The cost of S4/0/0 is 20.

Answer: D

Explanation:

From the command output, the interface S4/0/0 has its cost value explicitly listed under the Cost section as 20 for Level-1 circuits. This means that all traffic routed through this interface will incur this cost in the IS-IS metric calculation. The other options (such as circuit level and IPv6 support) are either not correct or not supported by the provided output .

Question: 74

Four routers run ISIS and have established adjacencies. The area IDs and router levels are marked in the following figure. If route leaking is configured on R3, which of the following is the cost of the route from R4 to 10.0.2.2/32?

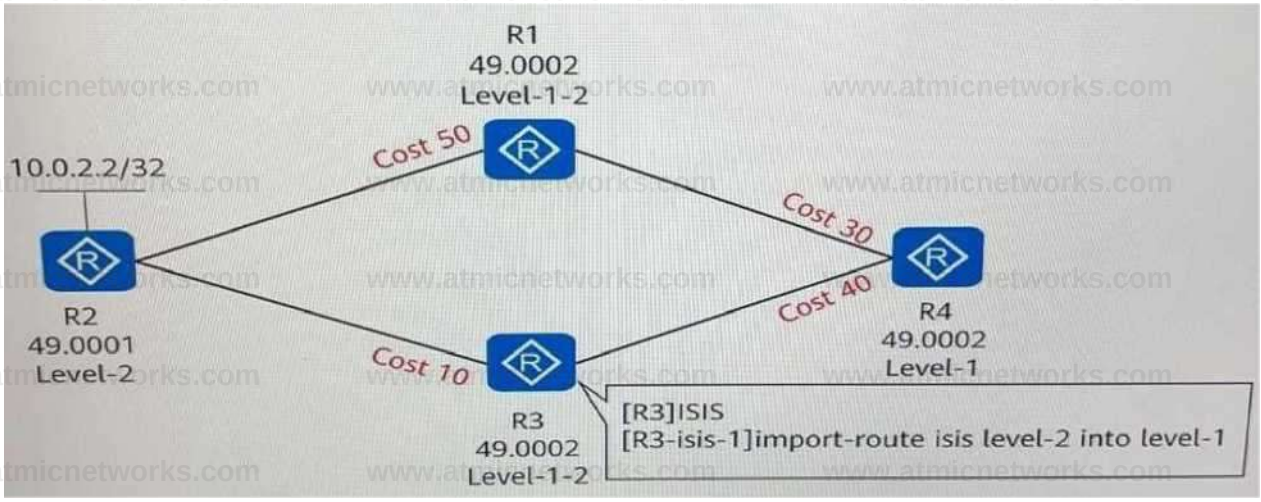

A. 80

B. 50

C. 40

D. 30

Answer: A

Explanation:

In the given topology, the route from R4 to 10.0.2.2/32 traverses R3, which performs route leaking from Level-2 to Level-1. The cost is calculated as follows: R4 to R3 (40) + R3 to R2 (10) + R2 to the destination (30), resulting in a total cost of 80 .

Question: 75

Which of the following is the default interval at which the DIS on a broadcast IS_IS network sends

CSNPs.

A. 30

B. 3.3

C. 10

D. 40

Answer: A

Explanation:

The Designated Intermediate System (DIS) in an IS-IS broadcast network sends Complete Sequence Number PDUs (CSNPs) at a default interval of 30 seconds. This interval ensures periodic synchronization of the Link State Database (LSDB) among IS-IS neighbors .

Question: 76

See the following figure. RTA, RTB, RTC, and RTD are in the same AS and establish IBGP peer relationships through direct links. RTB an RTC are route reflectors (RRs), RTA and RTC are the RR clients of RTB, and RTB and RTD are the RR clients of RTC. If RTA advertises the route 10.1.1.0/24 to the BGP process, the Originator ID of the BGP route received by RTD is----------.

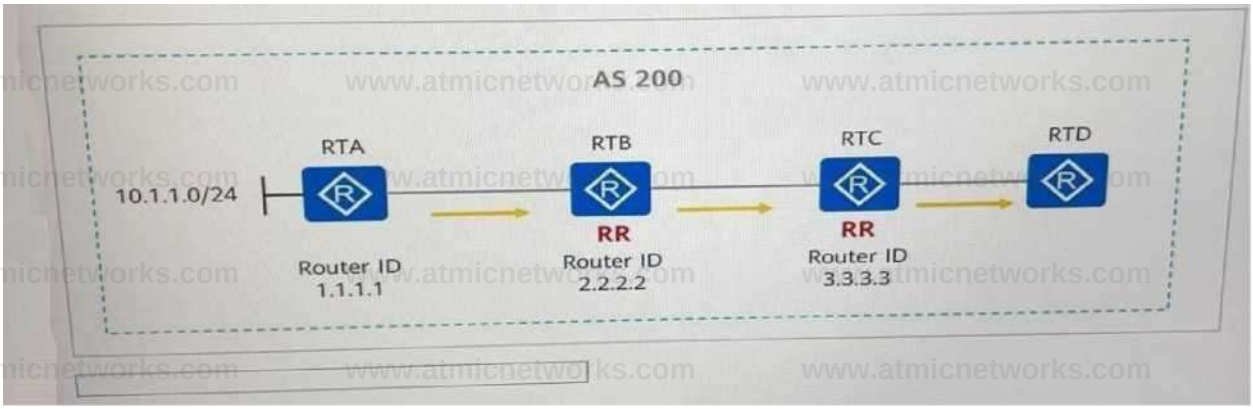

Answer: 1.1.1.1

Explanation:

The Originator ID in a BGP route represents the Router ID of the device that originally advertised the route into the BGP domain. Since RTA originates the route 10.1.1.0/24 into the BGP process, the Originator ID in the route received by RTD will be set to 1.1.1.1 .

Question: 77

After which of the following parameters are modified does an IS-IS neighbor relationship need to reestablished?

A. The cost of an IS-IS interface is changed.

B. The IP address of an ISIS interface is changed.

C. The level of an ISIS interface is changed.

D. The interval at which an IS-IS interface sends IIH packets is changed.

Answer: C

Explanation:

When the level of an IS-IS interface is changed, the adjacency must be re-established because IS-IS adjacencies are formed based on matching levels (Level-1, Level-2, or both). Other changes, such as cost or hello intervals, do not disrupt the existing adjacency but might impact routing metrics or timing .

Question: 78

To inject IGP routes into BGP routes, you can only use the network command.

A. TRUE

B. FALSE

Answer: B

Explanation:

IGP routes can be injected into BGP using multiple methods, not just the network command. The import-route command can also be used to redistribute IGP routes into BGP. The network command requires the route to exist in the routing table, while import-route allows direct redistribution .

Question: 79

In a route-policy, which of the following BGP attributes can be used in apply clauses?

A. MED

B. AS_Path

C. Tag

D. Local-Preference

Answer: AD

Explanation:

In a route-policy, attributes such as MED (Multi-Exit Discriminator) and Local-Preference can be used in apply clauses to influence BGP route selection. Attributes like AS_Path and Tag are typically matched or filtered but not directly applied in the apply clause .

Question: 80

According to BGP route selection rules, the route with the higher Local_Pref is preferred.

A. TRUE

B. FALSE

Answer: A

Explanation:

According to BGP route selection rules, a higher Local_Pref value indicates a more preferred route. This rule is used to influence traffic flow within an autonomous system (AS), giving preference to routes with a higher Local_Pref for outgoing traffic .

Question: 81

If a router ID is configured in both the system via and the BGP view, BGP uses the router ID configured in the BGP view because the BGP view takes precedence over the system view.

A. TRUE

B. FALSE

Answer: A

Explanation:

When the router ID is configured in both the system view and the BGP view, the router ID in the BGP view takes precedence because BGP-specific configurations override global settings. This ensures that BGP operates with the most relevant configurations .

Question: 82

Which of the following attribute must be carried in an Update message?

A. Local-pref

B. Prefval

C. MED

D. AS-Path

Answer: D

Explanation:

In BGP, the AS-Path attribute is mandatory and must be included in all UPDATE messages. It records the autonomous systems a route has traversed. Attributes like Local-Preference, MED, and Prefval are optional and may not always be included. The AS-Path is crucial for loop prevention and route selection .

Question: 83

In BGP, the origin attribute of the routes imported using the import-route command is incomplete.

A. TRUE

B. FALSE

Answer: A

Explanation:

When routes are imported into BGP using the import-route command, the origin attribute is set to incomplete by default. This indicates that the route's origin is not known or is from an external source, as opposed to being explicitly learned through an IGP (IGP) or EGP (EGP) protocol .

Question: 84

Which of the following statements regarding the BGP error display of a router false?

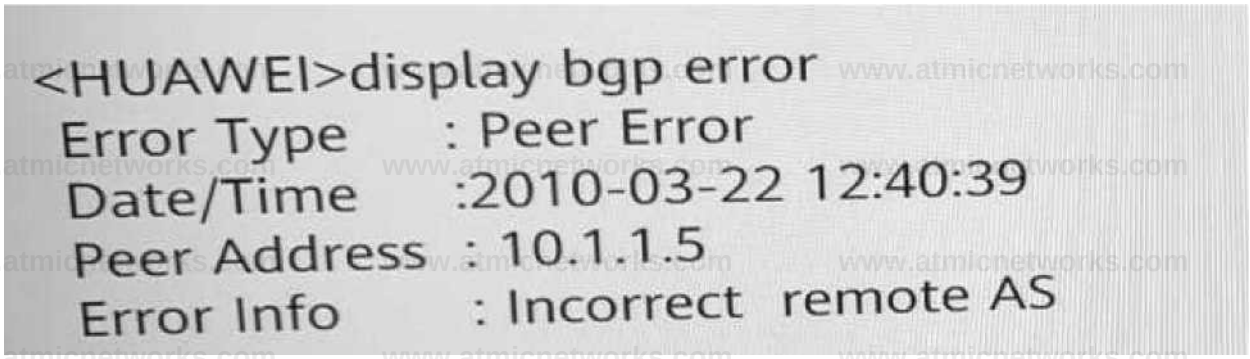

A. The error may be caused by the incorrect neighbor address.

B. The neighbor address of this router is 10.1.1.5.

C. Error Type indicates that the BGP error is caused by the neighbor relationship error.

D. The error occurred at 12:40:39 on March 22, 2010.

Answer: A

Explanation:

The error indicates "Incorrect remote AS," which refers to a mismatch in the autonomous system (AS) numbers during BGP peer configuration. This issue is unrelated to the neighbor address. The neighbor address (10.1.1.5) and timestamp (2010-03-22 12:40:39) provided in the output are correct, and the error type confirms a neighbor relationship issue .

Question: 85

Which of the following statements regarding BGP route advertisement are false?

A. All the BGP routes learned from all BGP peer will be advertised to other BGP peers.

B. Only the optimal routes preferred by BGP can be advertised to other BGP peers.

C. By default, the routes learned from an IBGP peer will not be forwarded to other IBGP peers.

D. Only the routes learned from IGPs can be advertised to other BGP peers.

Answer: A

Explanation:

BGP does not advertise all routes learned from peers to other peers. By default, only the best (optimal) routes selected by BGP are advertised to other peers. Additionally, routes learned from IBGP peers are not forwarded to other IBGP peers unless a route reflector or confederation setup is used. The statement claiming all learned routes are advertised is false .

Question: 86

Which of the following statements regarding the summary automatic command and BGP route summarization is false?

A. After this command is configured, BGP sends only the summarized routes to peers

B. This command is used to implement automatic summarization. Automatic summarization takes precedence over manual summarization

C. This command enables automatic summarization for the locally imported routes

D. After this command is configured, BGP summarizes routes based on natural network segments

Answer: B

Explanation:

The summary automatic command in BGP is used to enable automatic summarization of routes, particularly for classful networks. However, manual summarization takes precedence over automatic summarization when both are configured. This makes the statement claiming automatic summarization takes precedence false .

Question: 87

Which of the following routing protocols support the default route through command configuration?

A. BGP

B. IS-IS

C. OSPF

D. ICMP

Answer: A BC

Explanation:

BGP, IS-IS, and OSPF support the configuration of default routes through specific command configurations. For example, BGP uses the default-originate command, IS-IS uses a default-route- advertise method, and OSPF can advertise default routes using default-information originate. ICMP, however, is not a routing protocol and does not support this functionality .

Question: 88

When a routing policy is used to filter routes, which of the following route prefixes will be denied by the IP prefix below?

```
[HUAWEI]ip ip-prefix aa index 10 permit 1.1.1.1 24 greater-equal 26 less-equal 32
```

A. 1.1.1.1/26

B. 1.1.1.2/16

C. 1.1.1.1/32

D. 1.1.1.1/24

Answer: B D

Explanation:

The prefix list permits the prefix 1.1.1.1/24 with a mask length greater than or equal to /26 and less than or equal to /32. This means only prefixes with the base 1.1.1.1 and mask lengths between /26 and /32 will be permitted. Routes like 1.1.1.2/16 and 1.1.1.1/24 do not meet the mask length criteria, so they will be denied. Conversely, routes like 1.1.1.1/26 and 1.1.1.1/32 satisfy the condition and will be permitted .

Question: 89

The filter-policy 2000 export command is run in an ISIS process. Which of the following statements about the functions

of a filter policy is false?

A. Is used together with the route import function to advertise some imported external routes to neighbors.

B. If this command is not run, the device by default advertises all routes that ISIS Imports form external routing protocols.

C. Controls the advertisement of routes generated by the device.

D. Filters LSPs to be advertised.

Answer: D

Explanation:

The filter-policy command in IS-IS is used to control the advertisement and filtering of routing information. However, it does not directly filter LSPs (Link State Packets); instead, it controls route advertisements. The false statement here is that it filters LSPs .

Question: 90

Preferences of routing protocols determine the sequence In which a router selects a route among routes to the same destination /earned through different routing protocols.

A. TRUE

B. FALSE

Answer: A

Explanation:

The preference of routing protocols determines the selection order of routes when a router receives multiple routes to the same destination from different routing protocols. For instance, OSPF has a higher preference over RIP, so OSPF routes will be selected first. Huawei routing preferences confirm this behavior .

Question: 91

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

```
[HUAWEI]ip ip-prefix 1 permit 11.1.0.0 16
[HUAWEI]route-policy Set-cost permit node 10 [HUAWEI-route-policy]if-
  match ip-prefix 1 [HUAWEI-route-policy]apply cost 300 [HUAWEI-route-
```

policy]quit
[HUAWEI]route-policy Set-cost permit node 20 [HUAWEI-route-policy]apply cost 200
[HUAWEI-route-policy] quit

A. The route 11.1.0.0/16 is permitted by node 10, and its tost 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-policy configuration specifies that node 10 matches the IP prefix 11.1.0.0/16 and applies a cost of 300 to this route. Once a route matches a node, it is no longer processed by subsequent nodes, making option D incorrect. Option B is also incorrect because only specific matching routes

have their costs changed .

Question: 92

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 designed to match both the IP address and its prefix length, making it suitable for filtering routing information rather than data packets. Options A and D are incorrect because IP prefix filters do not work directly on data packets, and they can match both prefix numbers and lengths .

Question: 93

RSTP provides different functions in different scenarios. Which of the following statements Is false?

A. After TC-BPDU attack defense function is enabled, you can set the number of times the switch processes TC BPDUs within a certain period

B. The role of the designated port that is enabled with root protection cannot be changed

C. If the edge port on the switch enabled with BPDU protection receives RST BPDU, the switch sets the edge port as a non-edge-port and triggers STP calculation

D. When the designated port enabled with root protection receives optimal RST BPDUs, the port enters the Discarding state and does not forward packets. If the port does not receive optimal RST BPDUs within a certain period of time, the port will automatically restore to the Forwarding state

Answer: B

Explanation:

Root protection ensures that a designated port does not become a root port by discarding superior BPDUs. However, if superior BPDUs are no longer received, the port can return to its original

forwarding state, meaning its role can change. This makes option B incorrect, while the other options correctly describe RSTP behavior .

Question: 94

Compared with STP, RSTP defines the different port states. Which of the following statements regarding discarding and learning states are true?

A. The port In discarding or learning state does not forward data frames.

B. The port in discarding state does not learn MAC addresses table.

C. The port in discarding or learning state does not learn MAC addresses

D. The port In learning stole docs not learn MAC addresses table.

Answer: A B

Explanation:

In RSTP, a port in the discarding state neither forwards data frames nor learns MAC addresses. A port in the learning state does not forward data frames but starts learning MAC addresses. Therefore, options A and B are correct, while C and D are incorrect due to misunderstandings of these states .

Question: 95

As shown in the figure, SWA, SWB, SWC, and SWD run the Rapid Spanning Tree Protocol (RSTP).

Which of the following statements are true?

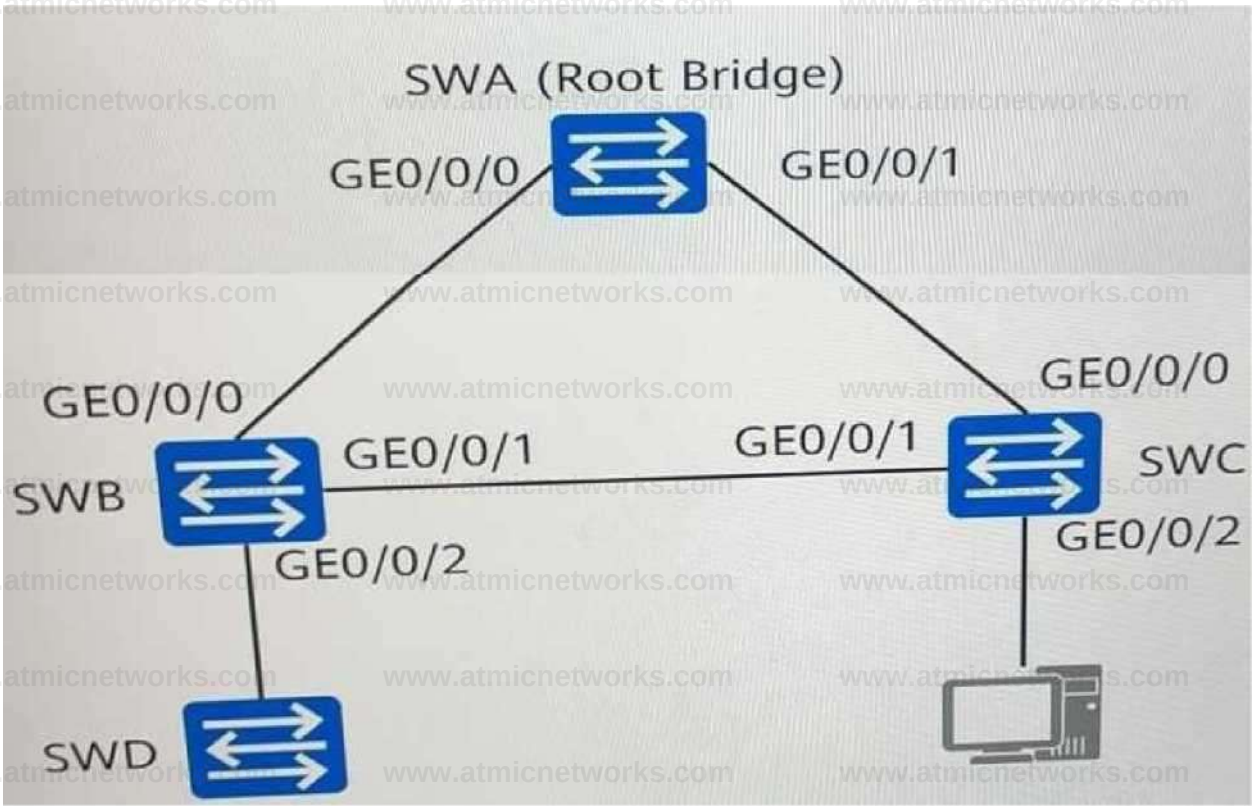

A. Once receiving OPOUs, the edge port re-partlci pates in the calculation of the spanning tree.

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: B CD

Explanation:

Edge ports are used in RSTP to bypass the spanning tree calculation for ports connected to end devices. By configuring ports like SWD's GE0/0/2 and SWC's GE0/0/2 as edge ports, they transition directly to the Forwarding state upon activation, ensuring faster convergence. Statement A is incorrect because receiving BPDU invalidates the edge port status .

Question: 96

A switch runs MSTP. The configuration is shown in the figure. What is the role of this switch in MSTI 1?

[SWA] display stp brief

MSTID	Port	Role	STP State	Protection
0	GigabitEthernetO/O/12	DESI	FORWARDING	NONE
0	GigabitEthernetO/O/15	ROOT	FORWARDING	NONE
1	GigabitEthernetO/O/12	DESI	FORWARDING	NONE
1	GigabitEthernetO/O/1 5	DESI	FORWARDING	*NONE*
2	GigabitEthernetO/O/1 2	DESI	FORWARDING	*NONE*
2	GigabitEthernetO/O/1 5	ROOT	FORWARDING	*NONE*

A. Uncertain

B. Root switch

C. Non-root switch

D. Secondary root switch

Answer: A

Explanation:

Based on the MSTP configuration shown, the role of the switch in MSTI 1 cannot be determined without additional details about the topology or priority values of other switches in the instance. The role could be a root switch, secondary root, or non-root, depending on these factors .

Question: 97

Which of the following statements regarding IGMPv1 and IGMPv2 are true?

A. 1GMPv2 supports only general query.

B. 1GMPv2 defines the Leave message type.

C. IGMPv1 does not define the IGMP Leave message type.

D. IGMPv1 supports general query.

Answer: BCD

Explanation:

IGMPv1 supports general queries but does not define a Leave message. IGMPv2 introduces the Leave message type, which allows hosts to notify routers when they wish to leave a multicast group, improving efficiency. These distinctions align with the functionality of IGMP versions .

Question: 98

IGMPV3 not only supports IGMPv1 General Query and IGMPv2 Group-Specific Query, and also IGMPv3 Source/Group-Specific Query.

A. TRUE

B. FALSE

Answer: A

Explanation:

IGMPv3 builds on IGMPv1 and IGMPv2 by supporting all previous query types, including General Query and Group-Specific Query, while introducing Source/Group-Specific Query. This enhancement allows finer control over multicast group membership .

Question: 99

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

```
<RTA> display igmp interlace
Interface information of VPN-Instance*, public net
 GigabitttthernetO/O/1(192.168.1.1):
    IGMP is enabled
    Current IGMP version is 2
    IGMP state: up
     IGMP group policy: none
    IGMP limit: -                                    ,"
      Value of query interval for IGMP (negotiated)
      Value of query interval for IGMP (configured): 60 s
       Value of other querier timeout for IGMP: Os

          Value of maximum query response time for IGMP. 10 s Querier for IGMP: 192.168.1.1
(this router)
```

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: A

Explanation:

The information provided indicates that the configured query interval is 60 seconds, which applies to general queries. Group-specific queries are sent at a shorter interval, usually derived from the

maximum response time. Therefore, the statement about a 60-second interval for group-specific queries is false. Other statements regarding maximum response time (10s), IP address, and IGMP version (V2) are correct .

Question: 100

Which of the following PIM protocol packets have unicast destination addresses.

A. Register Stop

B. Bootstrap

C. Graft

D. Assert

Answer: C A

Explanation:

In the PIM protocol, Register Stop and Graft messages are sent with unicast destination addresses, typically to specific neighbors or RP (Rendezvous Point) routers. Other messages like Bootstrap and Assert use multicast addresses for broader dissemination across the network .

Question: 101

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: BCD

Explanation:

Understanding DHCPv6:

DHCPv6 (Dynamic Host Configuration Protocol for IPv6) is used to assign configuration parameters to IPv6 clients.

It supports both stateful and stateless modes:

Stateful: Assigns IPv6 addresses and other configuration parameters.

Stateless: Assigns only configuration parameters (e.g., DNS server) without providing IPv6 addresses.

Reference: HCIP-Datacom-Core Technology Training Material (IPv6 Address Configuration and DHCPv6).

Analysis of Parameters:

A . Gateway address:

False. DHCPv6 does not assign the default gateway. Instead, the default gateway is advertised by routers using Router Advertisement (RA) messages in IPv6.

B . DNS server address:

True. DHCPv6 can assign the IPv6 address of DNS servers to the client.

C . IPv6 address/prefix:

True. In stateful mode, DHCPv6 assigns IPv6 addresses and prefixes to clients.

D . SNTP server address:

True. DHCPv6 can assign the address of an SNTP (Simple Network Time Protocol) server to synchronize time on the client.

Reference: HCIP-Datacom-Core Technology Training Material (DHCPv6 Configuration and Features).

Conclusion:

The correct parameters that a DHCPv6 server can assign to a client are: B. DNS server address, C. IPv6 address/prefix, D. SNTP server address.

Question: 102

Which of the following statements about multicast packet forwarding is true?

A. If a multicast data packet fails the RPF check, the packet must have been received through a sub- optimal interface. However, this interface still receives and forwards the multicast traffic downstream.

B. IGMP snooping cannot control the scope of multicast traffic flooding on a Layer 2 network.

C. The source address of a multicast packet is a unicast address.

D. In multicast transmission, the destination address of a packet can be the unicast address of a host.

Answer: C

Explanation:

A stateful inspection firewall tracks the state of network connections and only matches the initial packet against its rule set. Subsequent packets in the same connection are matched in the state table. Contrary to this, UDP packets can be inspected by correlating them with connection states, and packets in a single connection are always correlated .

Question: 103

Which of the following statements regarding the stateful inspection firewall is true?

A. When the stateful inspection firewall checks packets, packets of one same connection are not correlated.

B. Because UDP is a connectionless protocol, so the stateful inspection firewall cannot match UDP packets with the status table.

C. The stateful inspection firewall only needs to match the first data packet against a rule, and the subsequent packets of the connection are matched directly in the state table.

D. The stateful inspection firewall needs to match the rules for each incoming packet.

Answer: C

Explanation:

A stateful inspection firewall tracks the state of network connections and only matches the initial packet against its rule set. Subsequent packets in the same connection are matched in the state table. Contrary to this, UDP packets can be inspected by correlating them with connection states, and packets in a single connection are always correlated .

Question: 104

Compress the 2001:0DBB:B8:0000:C030:0000:0000:09A0:CDEF address. (if the answer contains letters, capitalize them.)

Answer: 2001:DBB:B8:0:C030::9A0:CDEF

Explanation:

The IPv6 address 2001:0DBB:B8:0000:C030:0000:0000:09A0:CDEF can be compressed by removing leading zeros in each segment and collapsing consecutive zero groups into ::. The result is 2001:DBB:B8:0:C030::9A0:CDEF .

Question: 105

A Huawei firewall by default creates security zones named untrust, dmz, and local. (Use Lowercase letters.)

Answer: TRUST

Explanation:

By default, Huawei firewalls create security zones named untrust, dmz, trust, and local. These zones facilitate security policies for inbound, outbound, and inter-zone traffic control .

Question: 106

Which of the following attacks is not the network layer attack?

A. IP spoofing attack

B. ICMP attack

C. Smurf attack

D. ARP spoofing attack

Answer: D

Explanation:

ARP spoofing is a Layer 2 attack, as it targets ARP tables in switches or end devices. All other attacks, including IP spoofing, ICMP attacks, and Smurf attacks, are network layer (Layer 3) attacks .

Question: 107

Which of the following statements regarding the firewall zone security level is false?

A. The configured security level cannot be changed.

B. Two zones cannot be configured with the same security level.

C. The default security level of the new zone is 1.

D. Only the security level of the user-defined zone can be configured.

Answer: A

Explanation:

In Huawei firewalls, the security level of a zone can be reconfigured by an administrator, making the statement that it cannot be changed false. Other statements accurately describe security level restrictions or defaults .

Question: 108

ON a stateful inspection Firewall where there is no session table, when the status detection mechanism is enabled and the second packet (CYN+ACK) of 3-way hadshakes reaches the firewall. Which of the following statements is true?

A. If the firewall security policy permits packets to pass, the session table is created.

B. By default, when status detection is disabled and the permit policy is configured packets can pass.

C. Packets must pass the firewall, and a session table is established.

D. If the firewall security policy permits packets to pass, the packets can pass the firewall.

Answer: D

Explanation:

In a stateful inspection firewall, if the status detection mechanism is enabled, it tracks and validates the state of connections using the session table. If there is no session table and a SYN+ACK packet reaches the firewall, it checks the security policy. If the policy explicitly permits the packet, it will pass through the firewall, but no session table will be created without the initial SYN packet. The other options are either incorrect or misrepresent the behavior of stateful inspection .

Question: 109

Which of the following configurations are not mandatory when an administrator configures VRRP?

A. Preemption mode

B. Preemption delay

C. Virtual router priority

D. Virtual IP address

Answer: AB

Explanation:

When configuring VRRP, the virtual router priority and virtual IP address are mandatory because they define the VRRP instance and its role within the group. Preemption mode and preemption delay are optional configurations that influence how routers take over the master role when priorities change. These are not mandatory for the VRRP protocol to function .

Question: 110

DRAG DROP

Drag the following VRRP states to the corresponding working mechanisms.

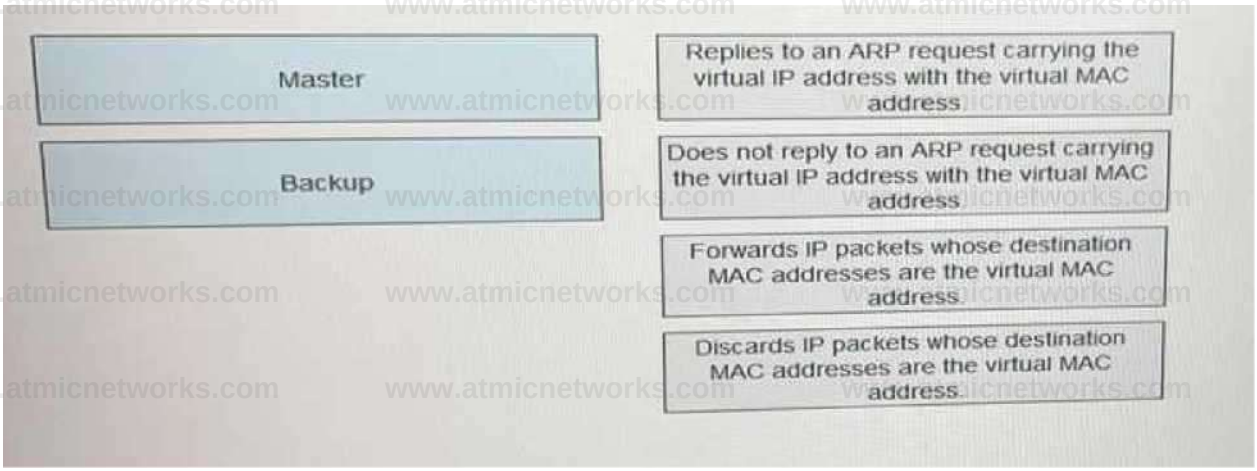

Answer:

Explanation:

Master

Backup

Master

Backup

Question: 111

The Interface \P address and VRRP virtual IP address can be the same.

A. TRUE

B. FALSE

Answer: A

Explanation:

BFD control packets are encapsulated in UDP packets, and the destination port number for multi-hop BFD control packets is 4784. This is a standardized port for multi-hop BFD operation .

Question: 112

BFD control packets are encapsulated in UDP packets for transmission. What is the destination port number of multi-hop BFD control packets?

A. 4784

B. 3784

C. 5784

D. 2784

Answer: A

Explanation:

BFD control packets are encapsulated in UDP packets, and the destination port number for multi-hop BFD control packets is 4784. This is a standardized port for multi-hop BFD operation .

Question: 113

There are two BFD operating modes. In mode, the local end sends BFD Control packets at specified intervals, and the remote end checks whether the local end periodically sends BFD Control packets. (Use lowercase letters.)

Answer: asynchronous

Explanation:

In asynchronous mode, the local device sends BFD control packets at predefined intervals, and the remote device monitors the receipt of these packets to detect connectivity issues. This is one of the two operating modes in BFD, the other being demand mode .

Question: 114

As shown in the following figure, a new AP is deployed In dual link MSB networking (load balancing mode). Which AC will the connect to?

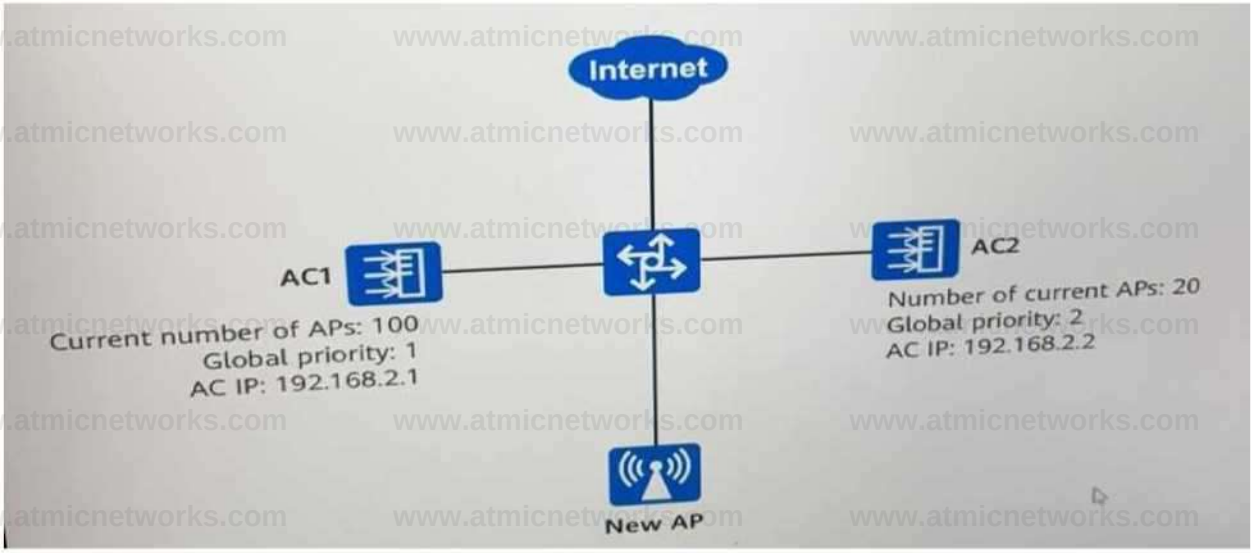

A. AC1

B. Random access

C. None

D. AC2

Answer: D

Explanation:

In dual-link MSB networking with load balancing, the new AP will connect to the AC with the lower number of connected APs. In this case, AC2 has only 20 connected APs, while AC1 has 100. Therefore, the new AP will connect to AC2 .

Question: 115

An engineer sets the CAPWAP heartbeat detection interval to 20 of the active link before an active/standby switchover occurs?

A. 75s

B. 20s

C. 60s

D. 90s

Answer: C

Explanation:

The CAPWAP heartbeat detection interval is typically multiplied by a predefined factor to determine the switchover time.

For instance, in Huawei devices, a common default is three missed heartbeats before a switchover is triggered. With a detection interval of 20 seconds, the switchover occurs after 60 seconds (20 x 3). This ensures that transient network issues do not cause unnecessary switchovers

Question: 116

DRAG DROP

What is the correct procedure for a VLAN pool to assign terminals to VLANs.

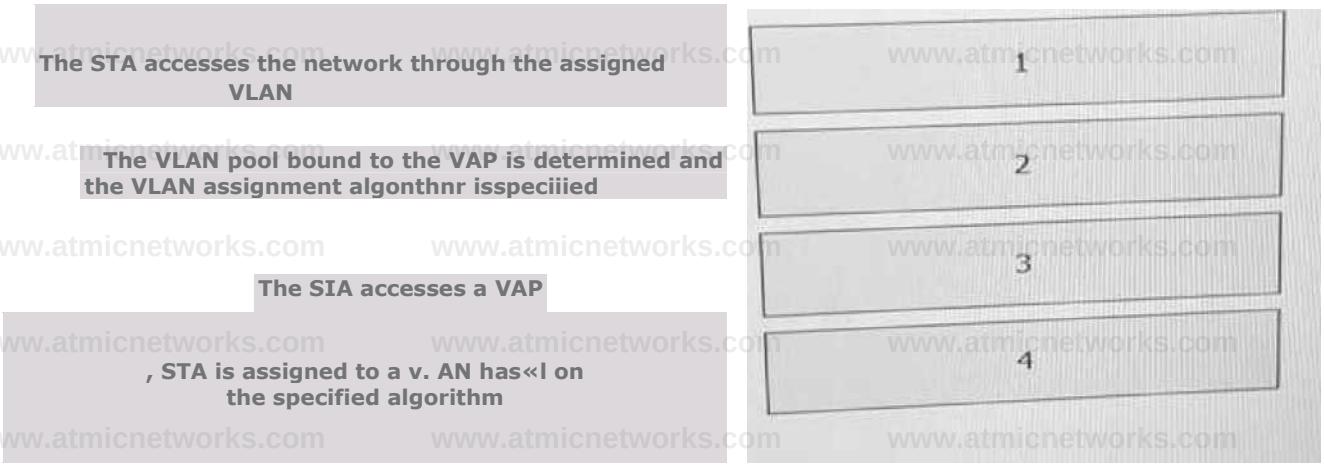

Answer:

Explanation:

Question: 117

Which of the following statements about WLAN roaming are true?

A. APs for roaming do not need to have overlapping signal coverage.

B. APs for roaming must be in the same extended service set (ESS).

C. APs for roaming must have overlapping signal coverage.

D. APs for roaming must be in the same basic service set (BSS).

Answer: A

Explanation:

Question: 118

Which Of the following IEEE 802.11 standards is also known as Wi-Fi 6?

A. 802.11ac

B. 802.11n

C. 802.11ax

D. 802.11b

Answer: C

Explanation:

Wi-Fi 6 corresponds to the IEEE 802.11ax standard. It offers improved efficiency, higher data rates, and better performance in dense environments compared to its predecessors like 802.11ac (Wi-Fi 5) and 802.11n (Wi-Fi 4) .

Question: 119

The traffic limiting policy feature only supports the number of connections initiated by the specified IP or the number of connections received.

A. TRUE

B. FALSE

Answer: B

Explanation:

The traffic limiting policy feature supports not only limiting the number of connections initiated by or received by an IP address but can also apply other traffic metrics such as bandwidth or packet rates. This expanded capability makes the statement false .

Question: 120

Which of the following methods is usually used by a network administrator to configure a newly purchased device for the first time?

A. Telnet

B. SNMP

C. Login through the Console port

D. FTP

Answer: C

Explanation:

For first-time configuration of a newly purchased network device, administrators typically use the console port. This direct connection is secure and independent of network configuration, making it suitable for initial setups. Telnet, SNMP, or FTP require prior IP configuration, which is not feasible during the initial setup phase .

Question: 121

A forwarding information database (FIB) can directly guide packet forwarding on a router.

A. TRUE

B. FALSE

Answer: A

Explanation:

The Forwarding Information Base (FIB) is derived from the router's routing table and contains the exact forwarding entries used to route packets. It directly guides packet forwarding decisions by matching incoming packets with the correct output interface or next-hop address .

Question: 122

Which of the following statements regarding OSPF route summarization commands are true?

A. The asbr-summary command is executed in the OSPF view.

B. The abr-summary command is executed in the OSPF area view.

C. Advertise is the default parameter of the abr-summary command. That is, if not-advertise is not specified in the abr-summary command, the advertise parameter takes effect by default.

D. Not-advertise is the default parameter of the abr-summary command. That is, if advertise is not specified in the abr-summary command, the not-advertise parameter takes effect by default.

Answer: A, B, C

Explanation:

The asbr-summary command is executed in the global OSPF view to summarize external routes. The abr-summary command is executed in the OSPF area view to summarize routes between areas. By default, the advertise parameter is enabled unless explicitly overridden by the not-advertise option .

Question: 123

See the command output of a router below. Which of the following statements is true?

<R2>display ospf interface GigabitEthernet 0/0/0 verbose

OSPF Process 1 with Router ID 10.0.2.2

Interface: 10.0.12.2 (GigabitEthernet 0/0/0)

Cost: 1 State: BDR Type: Broadcast MTU: 1500

Priority: 1

Designated Router: 10.0.12.1

Backup Designated Router: 10.0.12.2

Timers: Hello 10 , Dead 40 , Poll 120 , Retransmit 5 , Transmit Delay 1

A. The router ID is 10.0.12.2.
B. The interface cost is 1500.
C. This router is a BDR.
D. The interface IP address is 10.1.12.1.

Answer: C

Explanation:

The output shows that the router with the IP 10.0.12.2 is the Backup Designated Router (BDR), as indicated in the State: BDR field. The router ID is 10.0.2.2, the interface cost is 1, and the IP address of the interface is 10.0.12.2. Hence, C is the correct answer .

Question: 124

Which of the following LSAs are advertised only within a single area?

A. Network LSA
B. Router LSA
C. Summary LSA
D. AS External LSA

Answer: A, B

Explanation:

Router LSAs (Type 1) and Network LSAs (Type 2) are advertised only within a single OSPF area, as they describe the topology within the area. In contrast, Summary LSAs (Type 3) and AS External LSAs (Type 5) are used for inter-area and external routing, respectively .

Question: 125

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 routers refresh and flood LSAs in the Link-State Database (LSDB) every 30 minutes by default to maintain topology consistency. OSPF does not rely on TCP; it uses its own acknowledgment mechanism. It also uses the Dijkstra algorithm, not Bellman-Ford. The flooding interval for LSU packets is shorter, typically 5 seconds .

Question: 126

Which of the following statements regarding the OSPF protocol is false?

A. Each OSPF router uses only one Router-LSA to describe the local active connection status of an area.

B. Routing information can be advertised only between backbone and non-backbone areas and cannot be advertised directly between non-backbone areas.

C. Router-LSA describes four connection types: P2P, TransNet, SubNet, and virtual link.

D. Link State ID in a Type 3 LSA indicates the router ID of an ABR.

Answer: D

Explanation:

The Link State ID in a Type 3 LSA (Summary LSA) indicates the network or subnet being summarized, not the router ID of an ABR. All other statements accurately describe OSPF functionality, making D the false statement .

Question: 127

This configuration is part of RTA configuration. Which of the following statements regarding the configuration are true?

[RTA] ospf 100

[RTA-ospf-100]silent-interface GigabitEthernet 1/0/0

A. RTA cannot establish a neighbor relationship with the neighbor that this interface is directly connected to.

B. GigabitEthernet 1/0/0 is prohibited from sending OSPF packets.

C. Direct routes of GigabitEthernet 1/0/0 can still be advertised.

D. This interface cannot send Hello packets.

Answer: A, B, D

Explanation:

The silent-interface command disables the sending of OSPF packets, including Hello packets, on the specified interface, which prevents the establishment of OSPF neighbor relationships. However, the interface can still advertise directly connected routes through other interfaces, making Option C incorrect .

Question: 128

In the OSPF protocol, intra-area route calculation involves only Router LSA, Network LSA, and Summary LSA.

A. TRUE
B. FALSE

Answer: B

Explanation:

Intra-area route calculation in OSPF involves only Router LSAs (Type 1) and Network LSAs (Type 2). Summary LSAs (Type 3) are used for inter-area routes and do not participate in intra-area route calculation. Therefore, the statement is false .

Question: 129

Which of the following statements regarding OSPF multi-instance is false?

A. Route exchange between different OSPF processes is similar to route exchange between different routing protocols.
B. An interface of a router belongs to only a certain OSPF process.
C. The OSPF process IDs must be the same when OSPF neighbor relationships are established between different routers.
D. Multiple OSPF processes can run on the same router, and they are independent of each other.

Answer: C

Explanation:

OSPF process IDs are locally significant and do not need to match between routers to form neighbor relationships. The other statements correctly describe the behavior of OSPF multi-instances, making Option C false .

Question: 130

In an OSPF routing domain, two routers must be specified as one DR and one BDR in a broadcast or NBMA network with at least two routers.

A. TRUE

B. FALSE

Answer: A

Explanation:

In an OSPF broadcast or NBMA network with at least two routers, one router must act as the DR (Designated Router), and another as the BDR (Backup Designated Router) to ensure efficient communication and reduce the number of adjacencies. This is a requirement for OSPF operation in such network types .

Question: 131

Four routers run IS-IS and have established adjacencies. The area IDs and router levels are marked in the following figure. R1 and R2 are connected through a PPP link, and R3 is the DIS. Which of the following statements are true?

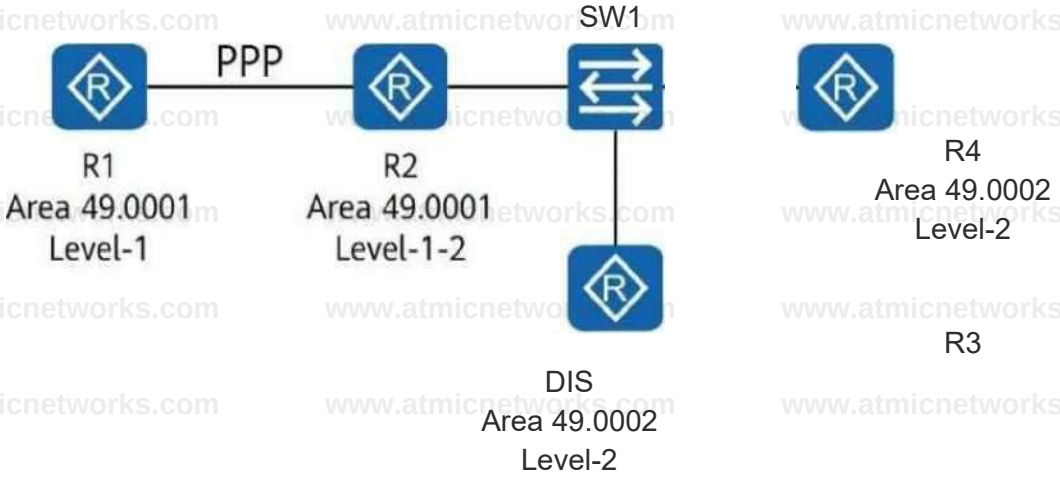

A. If R2 sends a Level-2 LSP, R3 needs to send a PSNP for acknowledgment.

B. R3 periodically sends CSNPs to implement Level-2 LSDB synchronization.

C. R2 sends an LSP to R3 and R4 in unicast mode.

D. If R1 sends an LSP, R2 needs to send a PSNP for acknowledgment.

Answer: B, D

Explanation:

R3, as the DIS in the Level-2 domain, periodically sends Complete Sequence Number Protocol Data Units (CSNPs) to ensure LSDB synchronization among Level-2 routers.

R1 and R2 are in a Level-1 domain, and IS-IS requires the receiving router (R2) to acknowledge an LSP from R1 by sending a Partial Sequence Number Protocol Data Unit (PSNP).

IS-IS LSPs in a broadcast network are sent using multicast, not unicast, making option C incorrect.

For Level-2 LSPs, acknowledgments are not provided using PSNP by R3; CSNPs are sufficient for synchronization .

Question: 132

Four routers run IS-IS and have established adjacencies. The area IDs and router levels are marked in the following figure. Which of the following statements is true?

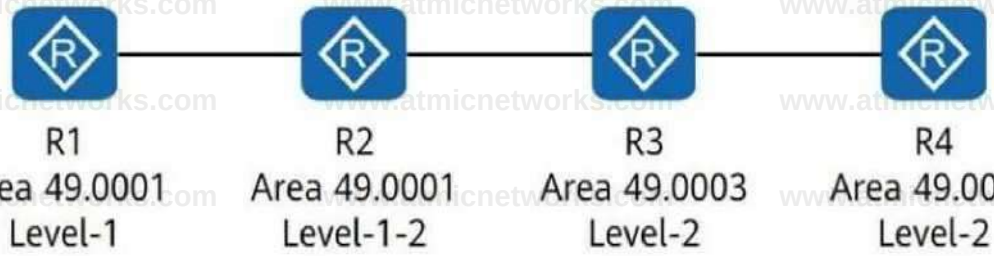

A. The LSDB of R2 does not containthe LSP of R4.
B. The LSDB of R1 does not containthe LSP of R4.
C. The LSDB of R2 does not containthe LSP of R3.
D. The LSDB of R3 does not containthe LSP of R4.

Answer: B

Explanation:

R1 is a Level-1 router in Area 49.0001, while R4 is in Area 49.0004. Level-1 routers only maintain LSDBs for their own area and do not contain LSPs from other areas, such as R4's LSP.

R2 is a Level-1-2 router, so it maintains LSDBs for both Area 49.0001 (Level-1) and Level-2 domains, which include LSPs from R4.

R3 and R4, as Level-2 routers, exchange LSPs with each other within the Level-2 domain.

Thus, the LSDB of R1 does not include R4's LSP .

Question: 133

See the following figure. All routers on the network run IS-IS and are in area 49.0001. By referring to the LSDB of R1, the Level-2 DIS is. (Enter the device name, for example, R1.)

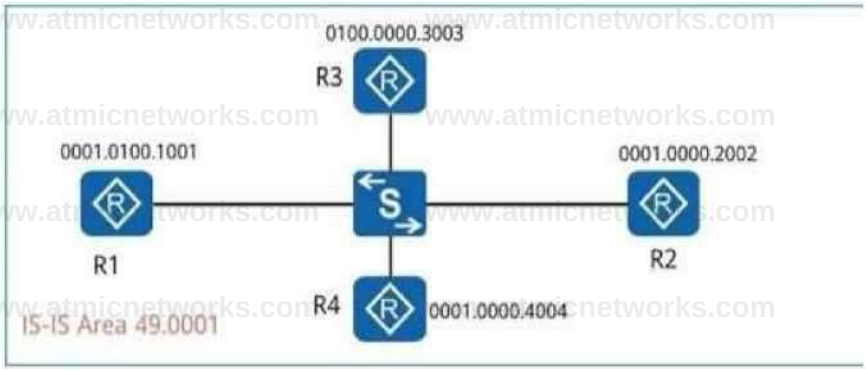

<R1 display isis lsdb
Database information for ISIS(1)

Level-1 Link State Database

LSPID	Seq Num	Checksum	Holdtime	Length	ATT/P/OL
0001.0000.4004.00-00	0x00000008	0xb701	1186	68	0/0/0
0001 0000.2002.00-00	0x00000008	0xb701	1186	68	0/0/0
0001.0100.1001.00-00*	0x00000005	0x2 f9d	1187	68	0/0/0
0001.0100.1001.01-00*	0x00000001	0xa79e	1110	55	0/0/0

<R1>display isis Lsdb
Database information for ISIS(1)

Answer: R3

Level-1 Link State Database

LSPID	Seq Num	Checksum	Holdtime	Length	ATT/P/OL
0001.0000.4004.00-00	0x00000008	0xb701	1186	68	0/0/0
0001.0000.2002.00-00	0x00000008	0xb701	1186	68	0/0/0
0001.0100.1001.00-00*	0x00000005	0x2f9d	1187	68	0/0/0
0001.0100.1001 01-00*	0x00000001	0xa79e	1110	55	0/0/0

Level-2 Link State Database

LSPID	Seq Num	Checksum	Holdtime	Length	ATT/P/OL
0001.0000.2002.00-00	0x00000008	0xb701	1188	68	0/0/0
0001.0100.1001.00-00*	0x00000006	0x2d9e	1187	68	0/0/0
0001.0100.1001.01-00*	0x00000005	OxdObO	1191	66	0/0/0
0100.0000.3003.00-00	0x00000005	0xfe53	1185	56	0/0/0

Explanation:

Understanding the LSDB and Level-2 DIS Election in IS-IS:

In IS-IS, the Designated Intermediate System (DIS) is elected for both Level-1 and Level-2 on broadcast networks to manage

the link-state database (LSDB) and reduce the number of LSAs exchanged.

The election of the DIS is based on the highest priority. If the priority is the same, the router with the highest System ID becomes the DIS.

Reference: HCIP-Datacom-Core Technology Training Material (IS-IS DIS Election).

Analyzing the LSDB:

From the LSDB of R1, the Level-2 link-state database includes the following entries:

LSPID 0100.0000.3003.00-00: This is R3, and it is present in the Level-2 LSDB with the sequence number and checksum details.

Other routers (R1, R2, R4) are present but do not have the characteristics of the Level-2 DIS in this topology.

Based on the System IDs, R3 (0100.0000.3003) has the highest System ID, making it the Level-2 DIS.

Conclusion:

The Level-2 DIS on the network is R3, based on the election rules and the LSDB information shown in the figure.

Question: 134

R1, R2, R3, and R4 run IS-IS, and the DIS priorities of their interfaces are shown in the following figure. If all these devices are started simultaneously, will be elected as the DIS. (Enter the device

name, for example, R1.)

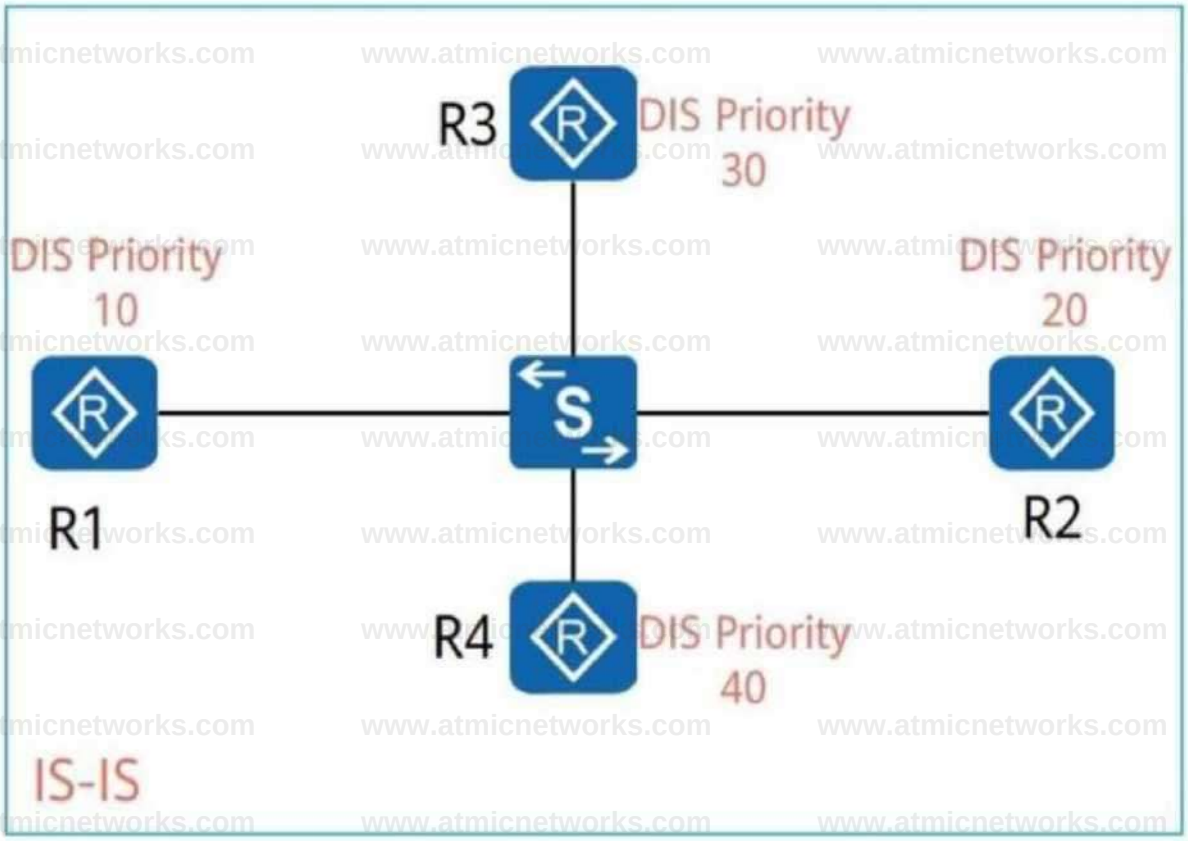

Answer: R4

Explanation:

DIS Election in IS-IS:

In IS-IS, the Designated Intermediate System (DIS) is elected on broadcast networks to reduce the number of LSAs exchanged and maintain the LSDB.

The DIS is elected based on priority:

The router with the highest priority is elected as the DIS.

If priorities are equal, the System ID is used as a tie-breaker, and the router with the highest System ID becomes the DIS.

Reference: HCIP-Datacom-Core Technology Training Material (IS-IS DIS Election Mechanism).

Given DIS Priorities:

R1: Priority = 10

R2: Priority = 20

R3: Priority = 30

R4: Priority = 40

R4 has the highest DIS priority of 40, which makes it the clear choice for DIS election.

Scenario Details:

Since all devices are started simultaneously, the DIS election process will follow the priorities without requiring a tie-breaker (System ID).

Conclusion:

The device with the highest DIS priority is R4, and it will be elected as the DIS.

Question: 135

See the network shown in the following figure.

R1 and R2 run IS-IS and establish an adjacency. IS-IS is enabled on Loopback0 of R2 but disabled on Loopback3 of R2. The configurations shown in the figure are performed in the IS-IS process of R2. Which of the following statements are true?

A. The routing table of R1 contains the route 10.0.2.3/32.
B. The routing table of R1 does not contain the route 10.0.2.3/32.
C. The routing table of R1 does not contain the routes 10.0.2.2/32 and 10.0.2.3/32.
D. The routing table of R1 contains the route 10.0.2.2/32.

Answer: B, D

Explanation:

In the provided configuration, IS-IS is enabled on Loopback0 but not on Loopback3 of R2. As a result, R2 will advertise the route 10.0.2.2/32 (from Loopback0) to R1 through IS-IS. However, the route 10.0.2.3/32 will not be advertised because IS-IS is not enabled on that loopback interface. The

import-route direct command does not override this behavior .

Question: 136

Which of the following statements regarding the BGP error display of a router is false?

A. The error occurred at 11:40:39 on March 22, 2010.
B. The neighbor address of this router is 10.1.1.2.
C. The error may be caused by the incorrect peer AS number.
D. Error Type indicates that the BGP error is caused by the neighbor relationship error.

Answer: A

Explanation:

The provided BGP error display shows the error timestamp as 12:40:39 on March 22, 2010. The error message confirms a mismatch in the peer AS number, as indicated in the "Incorrect remote AS" error info. This eliminates options related to other causes, and the timestamp in Option A is incorrect .

Question: 137

Which of the following statements regarding the display bgp routing-table command output is true?

```
<HUAWEI>display bgp routing-table

 BGP Local router ID is 192.168.2.1
 Status codes: * - valid, > - best, d - damped, h - history, i - internal, s -
               suppressed, S - Stale Origin : i - IGP, e - EGP, ? - incomplete

 Network            NextHop MED LocPrf PrefVal Path/Ogn

 *> 192.168.1.0/24  10.1.1.1  0                0       100i
```

A. The route to the destination address 192.168.1.0 is learned through AS 200.
B. The MED value of the route to the destination address 192.168.1.0 is 100.
C. The route to the destination address 192.168.1.0 is not the optimal route in the BGP routing table.
D. The route to the destination address 192.168.1.0 is injected into the BGP routing table using the network command.

Answer: D

Explanation:

The Origin attribute in the display bgp routing-table output is marked as i, indicating that the route was injected into the BGP routing table using the network command. The other options are incorrect because the AS path is not displayed, the MED is 0, and the route is marked as the best (>), meaning it is the optimal route .

Question: 138

Multiple BGP processes cannot be configured on the same router.

A. TRUE

B. FALSE

Answer: B

Explanation:

On Huawei devices, multiple BGP processes can be configured on the same router, but they are typically used for specific use cases such as multi-instance deployments or virtual routers. This flexibility is supported by Huawei's routing protocols .

Question: 139

Which of the following parameters are not mandatory during the configuration of a BGP peer?

A. password

B. Peer IP address

C. as-number

D. description

Answer: A, D

Explanation:

The mandatory parameters for BGP peer configuration are the peer IP address (peer ip-address) and AS number (as-number). A password is optional for MD5 authentication, and a description is an optional comment for reference. These are not required for establishing the peer relationship .

Question: 140

Which of the following statements regarding the MED value in BGP are true?

A. According to BGP route selection rules, the MED value has a lower priority than AS_Path, Preferred-Value, Local-Preference, and Origin.

B. The default MED value of BGP routes is 0.

C. By default, BGP can compare the MED values of routes from different ASs.

D. By default, if there is no MED value in routes, the value 0 is used. If the bestroute med-none-as- maximum command is configured, the maximum MED value 4294967295 is used.

Answer: A, D

Explanation:

The MED (Multi-Exit Discriminator) is an optional, non-transitive attribute used to influence the exit path. MED is considered after attributes such as Local-Preference, AS-Path, and Origin during route selection. By default, if a MED value is missing, it is treated as 0 unless the bestroute med-none-as- maximum command is configured to treat it as 4294967295 .

Question: 141

When two BGP peers support different Hold Time, they will negotiate to support the shortest Hold Time interval they can support.

A. TRUE

B. FALSE

Answer: A

Explanation:

When two BGP peers negotiate the Hold Time, the shorter value between the two configured Hold Time values is selected. This ensures compatibility and maintains the stability of the BGP session .

Question: 142

In the VRP, by default, the routes imported by BGP will be automatically summarized.

A. TRUE

B. FALSE

Answer: B

Explanation:

By default, BGP does not automatically summarize imported routes. Route summarization must be explicitly configured using the summary or similar commands .

Question: 143

Which of the following statements regarding Local-Preference in BGP is true?

A. Local-Preference affects traffic that enters an AS.

B. Local-Preference can be transmitted between ASs.

C. The default Local-Preference value is 100.

D. Local-Preference is a well-known mandatory attribute.

Answer: C

Explanation:

The Local-Preference attribute is used within an AS to influence outbound traffic paths. The default value is 100, and it is a well-known discretionary attribute, meaning it is not mandatory and does not travel across AS boundaries .

Question: 144

Which of the following scenarios is not suitable for deploying interface PBR?

A. A core switch needs to forward traffic between the intranet and extranet to an AC device that connects to the core switch in off-path mode.

B. A core switch needs to forward the traffic between the intranet and extranet to a security detection device that connects to the core switch in off-path mode.

C. A device needs to modify the next-hop IP address for locally originated traffic.

D. On an enterprise network with multiple ISP outbound interfaces, each internal network segment accesses the Internet through a particular ISP outbound interface.

Answer: C

Explanation:

Policy-Based Routing (PBR) on interfaces is used to forward traffic based on specific policies rather than the routing table. However, PBR is not suitable for modifying the next-hop of locally originated traffic. Other scenarios listed are typical use cases for interface PBR .

Question: 145

Which of the following attributes cannot be directly referenced in an apply clause of a route-policy?

A. community

B. IP-prefix

C. tag

D. origin

Answer: B

Explanation:

Attributes such as community, tag, and origin can be directly referenced in the apply clause of a route-policy. However, IP-prefix is not an attribute but a prefix list used for matching, and it cannot be directly applied .

Question: 146

A route-policy can have multiple nodes, and each node can have multiple if-match and apply clauses. Which of the following statements are false?

A. The operator between if-match clauses under a node is AND.
B. The operator between nodes is AND.
C. The operator between if-match clauses under a node is OR.
D. The operator between nodes is OR.

Answer: B, C

Explanation:

The operator between if-match clauses under a node is AND, meaning all conditions must be satisfied. However, the operator between nodes is OR, meaning a route matching any node will pass the route-policy .

Question: 147

Which of the following statements regarding routing policy and policy-based routing are true?

A. A routing policy is used to control import, advertisement, and receiving of routing information.
B. Policy-based routing is used to control import, advertisement, and receiving of routing information.
C. Policy-based routing is used to control packet forwarding without following routes in the routing table.
D. A routing policy is used to control packet forwarding without following routes in the routing table.

Answer: A, C

Explanation:

Routing policies are used to control how routes are imported, advertised, or received. Policy-Based Routing (PBR) is used to forward packets based on policies, bypassing the routing table. Routing
policies are not used for packet forwarding .

Question: 148

Access control lists can be classified into which types as follows?

A. Basic ACL
B. User-defined ACL
C. Advanced ACL
D. Layer 2 ACL

Answer: A, , C, D

Explanation:

Access Control Lists (ACLs) can be classified into several types, including Basic ACL (filters based on source IP), Advanced ACL (filters based on multiple parameters like source and destination IP, protocol, and ports), and Layer 2 ACL (filters based on MAC addresses). User-defined ACLs are not a standard classification .

Question: 149

STP ensures a loop-free network but has a slow network topology convergence speed, affecting communication quality. RSTP has made some improvements based on STP. Which of the following improvements is not included?

A. If a port does not receive configuration BPDUs from the upstream device within four Hello intervals, the switch considers that the negotiation with the neighbor fails.

B. The non-root switch running RSTP sends configuration BPDUs at the interval specified by the Hello timer, which is performed independently by each device.

C. RSTP reduces five port states to three port states based on user traffic forwarding and MAC address learning.

D. RSTP deletes three port states and adds two port roles.

Answer: A

Explanation:

RSTP improves convergence by introducing changes to the BPDU handling mechanism and port roles. However, RSTP does not rely on a fixed timeout of four Hello intervals for negotiation failure. Instead, RSTP uses a more dynamic mechanism to detect topology changes. The other options describe valid RSTP improvements .

Question: 150

As shown in the figure, SWA, SWB, and SWC run the Rapid Spanning Tree Protocol (RSTP). What are the roles of SWB's GE0/0/2 and SWC's GE0/0/1?

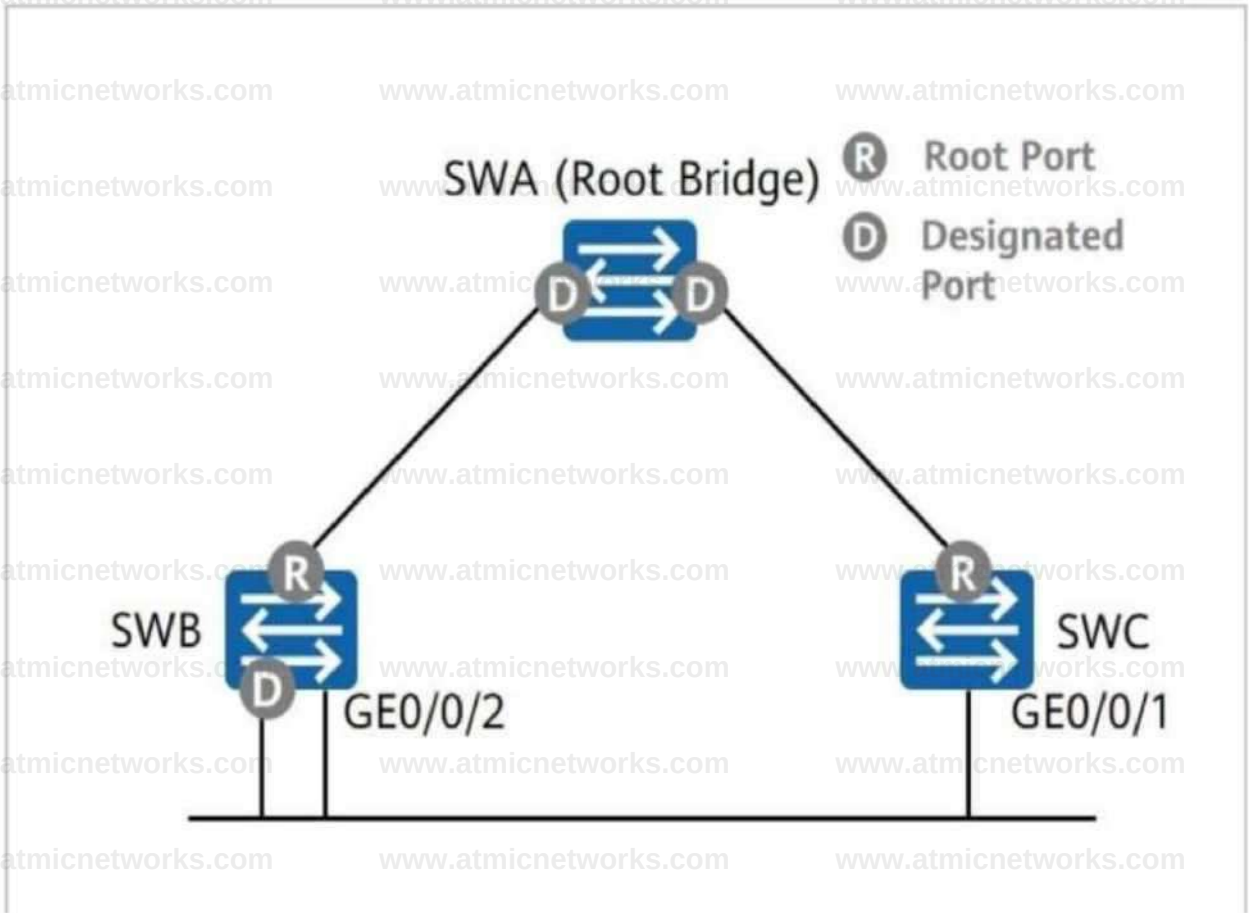

A. Backup port, alternative port B. Alternative port, backup port C. Backup port, root port
D. Root port, designated port

Answer: D

Explanation:

In the topology, SWB's GE0/0/2 is the Root Port because it has the least cost path to the Root Bridge (SWA). SWC's GE0/0/1 is the Designated Port because it forwards traffic to SWA while maintaining the shortest path to the Root Bridge. This setup ensures efficient traffic flow in the spanning tree topology .

Question: 151

On a network, some switches are enabled with RSTP and some switches are enabled with STP. What will happen?

A. A Huawei switch changes from RSTP to STP. After the STP-enabled switch is removed from the network, the RSTP-enabled switch can be moved back to the RSTP mode.

B. A Huawei switch changes its mode from STP to RSTP. After the RSTP-enabled switch is removed from the network,

the STP-enabled switch can be moved back to the RSTP mode.

C. STP and RSTP are compatible with each other, but the rapid convergence of RSTP is unavailable.

D. STP and RSTP calculation are performed independently.

Answer: A, C

Explanation:

When STP and RSTP coexist in a network, RSTP switches will downgrade their operation to STP to ensure compatibility. However, this disables RSTP's rapid convergence features, as the protocol behaves like STP to maintain interoperability .

Question: 152

Which of the following statements about stack split is false?

A. After a stack splits, the MAC addresses of the two stacks change immediately.

B. If the master and standby switches are still in the same stack after the stack splits, the slave switches separated from the original master and standby switches re-elect the master and standby switches due to protocol packet timeout.

C. If the master and standby switches are in two stacks after the split, the stack where the original master switch resides updates the topology and selects a new standby switch. The original standby switch becomes the master switch in the new stack and a new standby switch is elected.

D. After a stack splits, multiple stacks have the same IP address and MAC address. To prevent network faults, the stacks perform MAD detection. The stack that fails the MAD detection shuts down all physical ports except reserved ports.

Answer: A

Explanation:

When a stack splits, the MAC addresses of the newly formed stacks do not change immediately. Instead, they retain their original MAC addresses until new ones are reassigned during topology updates. This delayed update prevents immediate conflicts and allows MAD (Multiple Active Detection) to handle any resulting issues .

Question: 153

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-DM (Protocol Independent Multicast - Dense Mode), the SPT (Shortest Path Tree) establishment process involves flooding to propagate multicast traffic across the network, prune messages to stop sending multicast traffic to unwanted interfaces, and graft messages to re-establish forwarding when needed. State-refresh is not directly involved in SPT establishment but helps maintain the prune state .

Question: 154

Which of the following statements regarding multicast MAC addresses are false?

A. One multicast MAC address maps to 32 multicast IP addresses.

B. One multicast MAC address maps to only one multicast IP address.

C. The high 24 bits of the multicast MAC address are 0x01005E, the 25th bit is fixed to 1, and the lower 23 bits of the MAC address map to the lower 23 bits of the multicast IP address.

D. A multicast MAC address identifies receivers of a multicast group on the data link layer.

Answer: B

Explanation:

One multicast MAC address maps to 32 multicast IP addresses due to the limited number of bits used in the mapping process. Specifically, the lower 23 bits of the MAC address map to the lower 23 bits of the multicast IP address, while the high 24 bits of the MAC address are fixed as 0x01005E. Thus, Option B is false .

Question: 155

By default, a router interface sends PIM Hello messages at an interval of seconds.

Answer: 30 seconds

Explanation:

By default, a router interface sends PIM (Protocol Independent Multicast) Hello messages every 30 seconds to maintain neighbor relationships. This interval can be adjusted using configuration commands .

Question: 156

Which three transmission modes are supported for IPv4 packets?

A. Anycast

B. Broadcast

C. Unicast

D. Multicast

Answer: B, C, D

Explanation:

IPv4 supports three main transmission modes:

Unicast: One-to-one communication.

Broadcast: One-to-all communication within a network.

Multicast: One-to-many communication to a group of interested receivers.

Anycast is not a native IPv4 mode but is introduced in IPv6 .

Question: 157

ICMPv6 messages are classified as error or informational messages.

A. TRUE
B. FALSE

Answer: A

Explanation:

ICMPv6 messages are classified into two types: error messages (e.g., Destination Unreachable, Time Exceeded) and informational messages (e.g., Echo Request, Echo Reply). This classification is fundamental to ICMPv6's operation .

Question: 158

Compress the 2001:0DB8:0000:C030:0000:0000:09A0 address.

Answer: 2001:DB8:0:C030::9A0

Explanation:

To compress the IPv6 address 2001:0DB8:0000:C030:0000:0000:09A0, the following rules are applied:

Remove leading zeros in each hextet (e.g., 0DB8 becomes DB8).

Replace contiguous blocks of zeros with :: (only once in the address).

Thus, the compressed form is 2001:DB8:0:C030::9A0 .

Question: 159

A firewall receives a packet that PC1 sends to PC2. Which of the following statements are true?

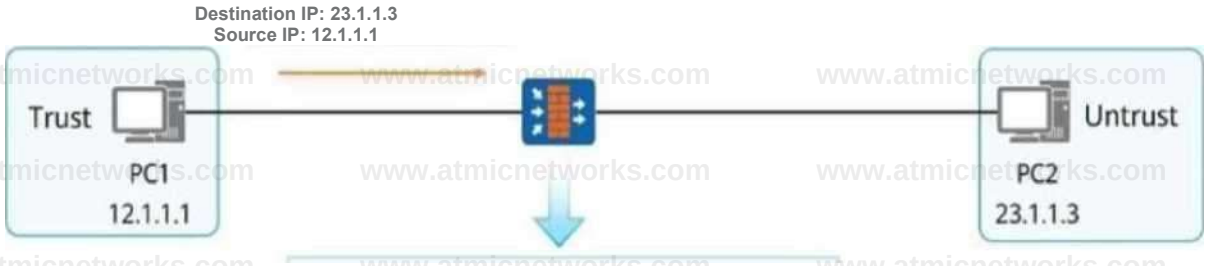

security-policy

rule name 1 source-zone trust destination-zone untrust source-address 12.1.1 2 mask 255 255.255 255 action permit

rule name 2 action deny

A. The packet does not match any security policy.

B. The packet matches security policy rule 1, and the firewall forwards the packet.

C. No source or destination security zone is specified in security policy rule 2, indicating that any security zone is a match.

D. The packet matches security policy rule 1, and the firewall discards the packet.

Answer: B

Explanation:

The security policy specifies that traffic originating from 12.1.1.2 and destined for the untrust zone is permitted. Since the source address of the packet (12.1.1.1) does not match this rule, the packet matches the default implicit deny rule. However, rule 1 does not deny all other traffic explicitly, so the packet is forwarded based on further configurations .

Question: 160

Which of the following protocols are multi-channel protocols?

A. H.323

B. FTP

C. Telnet

D. SMTP

Answer: A, B

Explanation:

Multi-channel protocols such as H.323 and FTP use separate control and data channels. H.323 uses different channels for call signaling and media streaming, while FTP uses a control channel for commands and a data channel for file transfers. Telnet and SMTP are single-channel protocols .

Question: 161

As shown in the figure, the stateful inspection firewall forwards the packet because the packet matches the session status of the firewall.

A. TRUE
B. FALSE

Answer: A

Explanation:

A stateful inspection firewall tracks the state of active connections. If a packet matches an existing session in the firewall's state table, it is allowed to pass. The diagram indicates that the TCP packet matches the session state, so the firewall forwards it .

Question: 162

ASPF enables the firewall to support multi-channel protocols such as FTP and to define security policies for complex applications.

A. TRUE
B. FALSE

Answer: A

Explanation:

Application Specific Packet Filtering (ASPF) is a feature that allows the firewall to understand and handle multi-channel protocols such as FTP, H.323, and SIP. ASPF inspects the control channel to dynamically create temporary rules for the data channels, enabling the firewall to secure and manage complex application protocols. This feature ensures that appropriate security policies are applied to these multi-channel applications .

Question: 163

DRAG DROP

To enable PC1 to access PC2 using the IP address of the firewall's GE0/0/2, you need to configure NAPT for source address translation. In the following figure, fill in the blank to complete the command.


```
[FWl]nat address-group inner
[FW1-address-group-inner] mode _ V
[FW1 -address-group-inner] section 0 10.0.12.1 10.0.12.1
[FW1-address-group-inner] quit
[FW1]nat-policy
[FW1-policy-nat] rule name 1
[FW1-policy-nat-rule-1] source-zone 0
[FW1 - policy-nat-rute-l]
[FW1-policy-nat-rule-1]
source-address O 24
destination-zone untrust
[FW1 policy nat-rule-1] action source-nat address-group _ O
[FW1-policy-nat-rule-1] quit
```

Answer:

Explanation:

pat

trust

12.1.1.0

inner

Question: 164

What is the default sending interval of BFD packets?

A. 10s
B. 5s
C. 100ms
D. 1000ms

Answer: C

Explanation:

The default sending interval for BFD (Bidirectional Forwarding Detection) control packets is 100 milliseconds. This interval ensures rapid detection of link faults, providing fast failover and minimizing downtime in network operations .

Question: 165

Which of the following statements about BFD operating modes are true?

A. In demand mode, once a BFD session is set up, the system no longer periodically sends BFD Control packets.
B. Asynchronous mode does not support the Echo function.
C. 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 time, the BFD session is declared down.
D. Asynchronous mode is the primary BFD operating mode.

Answer: A, C, D

Explanation:

In asynchronous mode, BFD control packets are exchanged periodically to detect faults.

In demand mode, once the session is established, BFD stops sending periodic control packets and relies on external mechanisms to verify connectivity.

Asynchronous mode is the most commonly used mode. However, the Echo function is supported in asynchronous mode, making Option B incorrect .

Question: 166

Which command is used to configure the VRRP preemption delay?

A. vrrp vrid 1 timer delay 20
B. vrrp vrid 1 preempt-delay 20
C. vrrp vrid 1 preempt-mode timer delay 20
D. vrrp vrid 1 preempt-timer 20

Answer: B

Explanation:

The correct command for configuring the VRRP (Virtual Router Redundancy Protocol) preemption delay is vrrp vrid 1 preempt-delay 20. This command sets the delay before a higher-priority VRRP router preempts the master role, ensuring stable operation during network transitions .

Question: 167

DRAG DROP

BFD for OSPF is deployed on a directly connected link. If the physical link is disconnected, drag the following BFD processes to the corresponding sequence numbers.

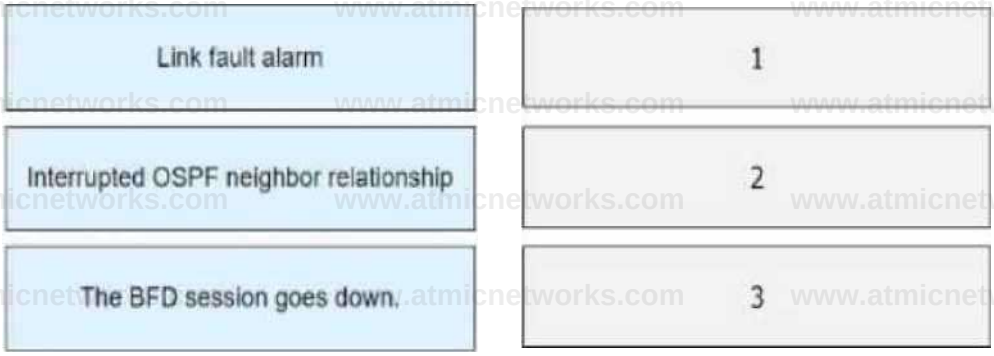

Answer:

Explanation:

Link fault alarm

The BFD session goes down.

Interrupted OSPF neighbor relationship

Question: 168

The VRID of the VRRP virtual router is 3 and the virtual IP address is 100.1.1.10. What is the virtual MAC address?

A. 01-00-5E-00-01-64

B. 01-00-5E-00-01-03

C. 00-00-5E-00-01-64

D. 00-00-5E-00-01-03

Answer: D

Explanation:

The virtual MAC address for a VRRP virtual router is determined by the formula 00-00-5E-00-01- [VRID in hexadecimal]. For VRID 3, the hexadecimal equivalent is 03. Thus, the virtual MAC address is 00-00-5E-00-01-03 .

Question: 169

During the DHCP interaction process, the DHCP server and client exchange various types of packets. Which of the following packets is not sent from the client to the server?

A. DHCP Release

B. DHCPNAK

C. DHCP Request

D. DHCP Discover

Answer: B

Explanation:

The DHCPNAK message is sent by the DHCP server to the client to indicate that the requested configuration is not valid. Messages like DHCPDiscover, DHCPRequest, and DHCPRelease are initiated by the client in the DHCP process .

Question: 170

On the network shown in the following figure, the management VLAN is VLAN 10, and the AP is configured to obtain an IP address on the network segment 10.1.1.0/24 through DHCP. The AP, AC, and AR function as the DHCP client, DHCP relay agent, and DHCP server, respectively. Which of the following configurations for the DHCP relay agent and DHCP server are correct?

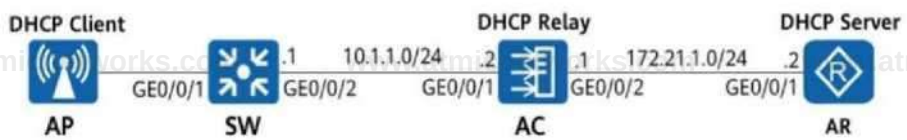

A. [AC] dhcp server group AP

[AC-dhcp-server-group-AP] dhcp-server 172.21.1.2

[AC-dhcp-server-group-AP] quit

[AC] interface Vlanif 10

[AC-Vlanif 10] dhcp select relay

[AC-Vlanif 10] dhcp relay server-select AP

[AC-Vlanif 10] quit

B. [AR] ip pool AP

[AR-ip-pool-AP] network 10.1.1.0 mask 24

[AR-ip-pool-AP] gateway-list 10.1.1.2

[AR-ip-pool-AP] excluded-ip-address 10.1.1.1

[AR-ip-pool-AP] quit

[AR] interface GigabitEthernet 0/0/1

[AR-GigabitEthernetO/O/1] dhcp select global

[AR-GigabitEthernet0/0/l] quit

[AR] ip route-static 10.1.1.0 255.255.255.0 172.21.1.1

C. [AR] ip pool AP

[AR-ip-pool-AP] network 172.21.1.0 mask 24

[AR-ip-pool-AP] gateway-list 172.21.1.2

[AR-ip-pool-AP] excluded-ip-address 172.21.1.1

[AR-ip-pool-AP] quit

[AR] interface GigabitEthernet 0/0/1

[AR-GigabitEthernetO/0/1] dhcp select global

[AR-GigabitEthernetO/0/1] quit

[AR] ip route-static 10.1.1.0 255.255.255.0 172.21.1.1

D. [AC] dhcp server group AP

[AC-dhcp-server-group-AP] dhcp-server 10.1.1.2

[AC-dhcp-server-group-AP] quit

[AC] interface Vlanif 10

[AC-Vlanif 10] dhcp select relay

[AC-Vlanif 10] dhcp relay server-select AP

[AC-VlaniflO] quit

Answer: A B

Explanation:

Option A correctly configures the DHCP relay agent (AC) to forward requests to the DHCP server using VLAN 10.

Option B correctly configures the DHCP server (AR) with the correct IP pool (10.1.1.0/24), gateway, and excluded IP addresses. The static route to 10.1.1.0 through 172.21.1.1 ensures proper routing .

Question: 171

As shown in the figure, data traffic is forwarded in tunnel mode along the path of STA -> HAP -> HAC - > upper-layer network before roaming. What is the flow direction of data traffic after Layer 3 roaming?

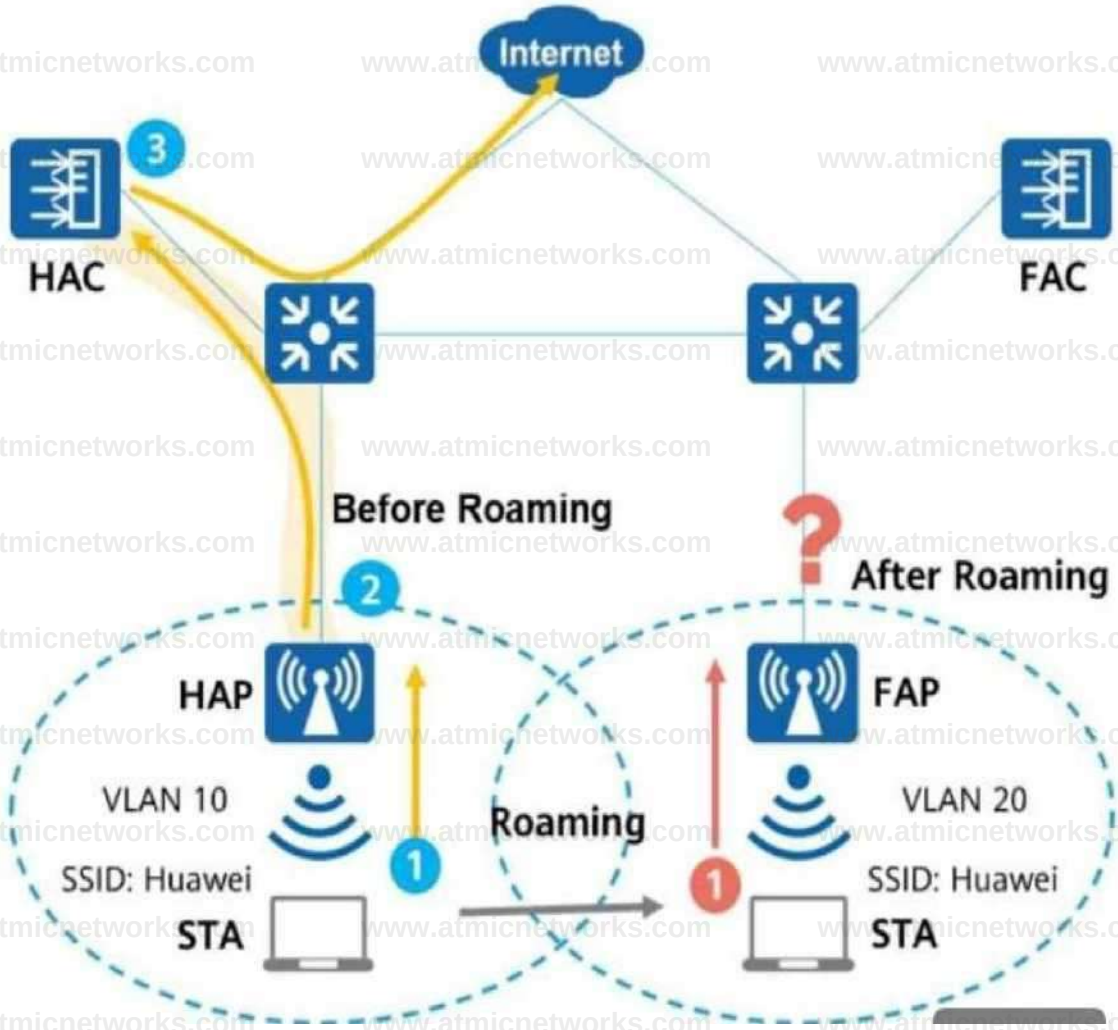

A. STA -> FAP -> FAC -> HAC -> HAP -> HAC -> Upper-layer network

B. STA -> FAP -> FAC -> HAC -> Upper-layer network

C. STA -> FAP -> FAC -> Upper-layer network

D. STA -> FAP -> FAC -> HAC -> HAP -> Upper-layer network

Answer: B

Explanation:

After Layer 3 roaming, the STA's traffic is forwarded from the Foreign AP (FAP) to the Foreign AC (FAC), and then to the Home AC (HAC) before reaching the upper-layer network. This ensures that the STA remains on the same subnet without breaking existing sessions .

Question: 172

After HSB is configured, the HSB channel fails to be established and cannot back up information on the active device to the standby device. What are the possible causes for this HSB function failure?

A. The TCP channel is not established.

B. The retransmission count and interval for HSB service packets are different on the two devices.

C. The two ACs have the same priority.

D. The source IP address and port number of the local end are different from the destination IP address and port number of the remote end.

Answer: A, B, D

Explanation:

High-availability Stateful Backup (HSB) ensures synchronization between active and standby devices. If the HSB channel fails to establish, the following causes are possible:

TCP channel not established: The HSB channel relies on a TCP connection. If the TCP session cannot be established, the backup process cannot start.

Mismatched retransmission settings: If the retransmission count and interval for HSB service packets are inconsistent between the two devices, the synchronization process fails.

Incorrect source and destination configurations: If the source IP address and port number on one device do not match the destination IP address and port number on the other, the channel cannot be established.

Option C is incorrect because the priority of the two ACs affects active/standby roles but does not directly impact HSB channel establishment .

Question: 173

Which of the following statements about WLAN roaming are false?

A. APs with which a STA is associated before and after roaming can work on different channels.

B. During roaming, the STA sends a Reassociation Request frame containing home AP information to the foreign AP.

C. During roaming, the STA sends an Association Request frame containing home AP information to the foreign AP.

D. APs with which a STA is associated before and after roaming must work on the same channel.

Answer: C, D

Explanation:

During WLAN roaming, the STA sends a Reassociation Request frame to the foreign AP to maintain connectivity. Association frames are used during initial connections, not roaming.

APs with which a STA is associated before and after roaming can work on different channels, as roaming often occurs across channels in environments like multi-SSID or multi-channel setups. Options C and D are false, as they contradict standard WLAN roaming procedures .

Question: 174

The native AC function allows an agile switch to integrate AC capabilities, achieving wired and wireless

convergence. The agile switch centrally manages wired and wireless service traffic.

A. TRUE
B. FALSE

Answer: A

Explanation:

The native AC (Access Controller) function in Huawei's agile switches enables wired and wireless convergence. This integration allows the agile switch to act as both a wired network switch and a wireless controller, managing all traffic centrally. This architecture simplifies network management and improves efficiency, making the statement true .

Question: 175

Depending on the type of algorithms used, routing protocols are classified into distance-vector protocols and link-state protocols. Which of the following routing protocols are link-state protocols?

A. RIP
B. OSPF
C. IS-IS
D. BGP

Answer: B, C

Explanation:

Comprehensive and Detailed Step-by-Step
Distance-Vector Protocols:
Distance-vector protocols determine the best path based on the distance to a destination and periodically share the entire routing table with neighbors.
Example: RIP is a distance-vector protocol.
Link-State Protocols:
Link-state protocols share information about directly connected links, enabling routers to build a complete map of the network.
OSPF (Open Shortest Path First) and IS-IS (Intermediate System to Intermediate System) are link-state protocols that use Dijkstra's algorithm to calculate the shortest path.
Border Gateway Protocol (BGP):
BGP is neither a distance-vector nor a link-state protocol. Instead, it is classified as a path-vector protocol that selects the best path based on attributes such as AS-path, next-hop, and local preference.
Reference:
HCIA-Datacom Study Guide, Chapter: Routing Protocol Basics
Huawei Networking Fundamentals, Section: Classification of Routing Protocols

Question: 176

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:

Comprehensive and Detailed Step-by-Step

Routing Table Structure:

Routers maintain a local core routing table and separate protocol routing tables for OSPF, BGP, etc. The local core routing table stores optimal routes selected based on metrics, protocol preferences, and administrative distances.

Fields in the Core Routing Table:

Destination Address (A): Indicates the destination network or host.

Routing Protocol Preference (B): Determines the priority of routing protocols (e.g., OSPF > RIP).

Routing Protocol (D): Specifies the protocol (e.g., OSPF, BGP) that contributed the route.

Inbound Interface: This is stored in the protocol-specific routing table but not in the local core table.

Reference:

HCIA-Datacom Study Guide, Chapter: Routing Table Management

Huawei Routing Fundamentals

Question: 177

On an office network of an enterprise, OSPF is enabled on two directly connected routers. During adjacency establishment, the state machine stays in the 2-way state. Which of the following statements are not possible causes?

A. Hello time values are different.
B. Area IDs are different.
C. Router IDs conflict.
D. The N-bit and E-bit of Hello packets are the same.

Answer: C, D

Explanation:

Comprehensive and Detailed Step-by-Step

OSPF State Machine:

OSPF adjacency progresses through several states: Down > Init > 2-Way > ExStart > Exchange > Loading > Full.

The 2-Way state indicates that routers have exchanged Hello packets but are not forming full adjacencies.
Possible Causes for Staying in the 2-Way State:
Hello Time Values Are Different (A): Mismatched Hello/Dead intervals prevent adjacency establishment.
Area IDs Are Different (B): OSPF routers must belong to the same area to establish adjacency.
Router IDs Conflict (C): While Router ID conflicts cause other issues, they do not result in the 2-Way state.
N-bit and E-bit Are the Same (D): Matching N-bit/E-bit values are normal and do not cause problems.
Reference:
HCIA-Datacom Study Guide, Chapter: OSPF Adjacency Establishment
Huawei OSPF Neighbor State Analysis

Question: 178

A campus network uses OSPF for network communication. The display ospf peer command is run on a router, and the command output is as follows:

vbnet
Copy
Edit
<R2> display ospf peer
OSPF Process 1 with Router ID 10.0.2.2
Area 0.0.0.0 interface 10.0.235.2 (GigabitEthernet0/0/1)'s neighbors
Router ID: 10.0.5.5 Address: 10.0.235.5
State: Full Mode: Nbr is Master Priority: 1
DR: 10.0.235.5 BDR: 10.0.235.3 MTU: 0
Dead timer due in 40 sec

Area 0.0.0.0 interface 10.0.24.2 (Serial1/0/1)'s neighbors
Router ID: 10.0.4.4 Address: 10.0.24.4
State: Full Mode: Nbr is Master Priority: 1
DR: None BDR: None MTU: 0
Dead timer due in 35 sec
Which of the following statements are true about the device?

A. The IP address of the DR on the network where R2 and its neighbor 10.0.5.5 reside is 10.0.235.5.
B. R2 has established adjacencies with 10.0.5.5 and 10.0.4.4.
C. The DR/BDR election fails on the network where R2 and its neighbor 10.0.4.4 reside.
D. Router IDs of the two OSPF neighbors of R2 are 10.0.5.5 and 10.0.4.4.

Answer: A, B, D

Explanation:
Comprehensive and Detailed Step-by-Step
Analysis of the OSPF Neighbor State:

Router ID 10.0.5.5:
On interface 10.0.235.2 (GigabitEthernet0/0/1), the neighbor state is Full, indicating that adjacency is established.

The Designated Router (DR) is 10.0.235.5, which is the same as the neighbor's address.

Router ID 10.0.4.4:

On interface 10.0.24.2 (Serial1/0/1), the neighbor state is also Full, meaning adjacency is established.

The DR/BDR election status shows None, indicating that the interface is on a point-to-point link where DR/BDR election does not apply.

Validation of Each Option:

Option A: Correct. The DR IP address is explicitly shown as 10.0.235.5.

Option B: Correct. The neighbor state for both 10.0.5.5 and 10.0.4.4 is Full, confirming adjacency is established.

Option C: Incorrect. DR/BDR elections are irrelevant on point-to-point links like Serial1/0/1.

Option D: Correct. The Router IDs of R2's neighbors are explicitly listed as 10.0.5.5 and 10.0.4.4. Reference:

HCIA-Datacom Study Guide, Chapter: OSPF Neighbor Relationships

Huawei OSPF Neighbor and Adjacency Details

Question: 179

OSPF has multiple types of routes with varying priorities. Which of the following types of routes has the lowest priority when they have the same prefix?

A. Type 1 external route
B. Inter-area route
C. Type 2 external route
D. Intra-area route

Answer: C

Explanation:

Comprehensive and Detailed Step-by-Step

OSPF Route Preference Hierarchy:

OSPF uses the following preference order for route types when the prefix is the same: Intra-area routes (D): Routes within the same OSPF area have the highest priority.

Inter-area routes (B): Routes between different OSPF areas are preferred after intra-area routes.

Type 1 external routes (A): External routes redistributed into OSPF, considering the internal OSPF cost to the ASBR.

Type 2 external routes (C): External routes redistributed into OSPF, ignoring the internal OSPF cost to the ASBR.

Lowest Priority:

Type 2 external routes (C) are given the lowest priority because they represent external information that

does not consider internal OSPF costs.
Reference:

HCIA-Datacom Study Guide, Chapter: OSPF Route Preference
Huawei OSPF Routing Table Selection Rules

Question: 180

An enterprise administrator configures route summarization on ASBRs to reduce the number of inter-area Type 3 LSAs, which in turn reduces the routing table size and improves device resource utilization.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step

OSPF Route Summarization:

Route summarization combines multiple specific routes into a single summary route. This reduces the number of LSAs advertised between OSPF areas, especially Type 3 LSAs.
Benefits of Route Summarization:

Reduces the size of the OSPF database and routing table.
Minimizes the processing overhead on routers by decreasing the number of routing entries.
Improves the scalability of the OSPF network.
ASBR Function:

The ASBR advertises external routes into OSPF and can summarize routes at area boundaries or during redistribution.
Reference:

HCIA-Datacom Study Guide, Chapter: OSPF Route Summarization
Huawei Routing Protocol Optimization

Question: 181

OSPF networks are classified into broadcast, P2P, P2MP, and NBMA networks. Which of the following types of networks use the default Hello time (30s)?

A. P2P
B. Broadcast

C. P2MP

D. NBMA

Answer: C, D

Explanation:

Comprehensive and Detailed Step-by-Step

Default OSPF Hello Timer:

The OSPF Hello timer is used to send periodic Hello packets and maintain neighbor relationships.

Default Hello timers:

Broadcast and P2P: 10 seconds

NBMA and P2MP: 30 seconds

Network Types Using 30s Hello Time:

P2MP (C): OSPF sends Hello packets every 30 seconds by default.

NBMA (D): OSPF also uses 30 seconds for NBMA networks.

Reference:

HCIA-Datacom Study Guide, Chapter: OSPF Network Types

Huawei OSPF Timer Configuration

Question: 182

On an IS-IS network, two directly connected routers establish a neighbor relationship through the three-way handshake mechanism by default, regardless of whether the network type is P2P or broadcast.

A. TRUE

B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step

IS-IS Three-Way Handshake:

IS-IS uses a three-way handshake to establish neighbor relationships.

This process is applied on both point-to-point (P2P) and broadcast network types to confirm bidirectional communication.

P2P Networks:

Routers directly exchange IS-IS Hello (IIH) packets to establish adjacency.

Broadcast Networks:

IS-IS also uses the three-way handshake on broadcast networks to ensure that the Designated Intermediate System (DIS) election and adjacency formation are correct.

Reference:

HCIA-Datacom Study Guide, Chapter: IS-IS Adjacency Establishment

Huawei IS-IS Network Configuration

Question: 183

On an IS-IS network, areas are divided by router, and a router can belong to only one are a. Therefore, an IS-IS router only needs to maintain the LSDB of its area.

A. TRUE

B. FALSE

Answer: B

Explanation:

Comprehensive and Detailed Step-by-Step

IS-IS Area Division:

IS-IS divides areas by routers, not by links. A router is assigned to a specific IS-IS area, but it can belong to multiple levels:

Level 1: Maintains the LSDB (Link State Database) for the local area only.

Level 2: Maintains the LSDB for the backbone area and exchanges routes between areas.

Behavior of a Level 1-2 Router:

A router configured as Level 1-2 must maintain two separate LSDBs:

One for intra-area routes (Level 1).

One for inter-area routes (Level 2).

Incorrect Statement:

The statement that an IS-IS router only needs to maintain the LSDB of its area is incorrect because Level 1-2 routers maintain LSDBs for both levels.

Reference:

HCIA-Datacom Study Guide, Chapter: IS-IS Area Structure

Huawei IS-IS Fundamentals

Question: 184

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:

Comprehensive and Detailed Step-by-Step
Analysis of BGP Attributes:
Original Next Hop: The output shows that the original next hop of the route is 172.16.1.2. (A is correct).
Local Router ID: The local BGP router ID is 10.1.1.1, as displayed. (B is correct).
Local AS Number: The local AS number is 100, as displayed. (C is correct).
Incorrect Option:
Local_Pref: Although Local_Pref is an important attribute, this route is marked as invalid due to "IP unreachable" and thus cannot be preferentially selected.
Reference:
HCIA-Datacom Study Guide, Chapter: BGP Routing Attributes
Huawei BGP Routing Table Analysis

Question: 185

A large number of routes typically exist in a BGP routing table, and transmitting such extensive routing information brings a heavy burden to a device. In order to address this problem, it is necessary to filter those routes to be advertised. You can configure a device to advertise only necessary routes or those that its peers require.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step
BGP Route Filtering:
BGP provides flexible policies for route filtering to optimize routing table size and reduce network resource usage.

By configuring route filters, a device can advertise only required routes to peers, thereby improving efficiency.
Purpose of Filtering:
Reduces the routing table size.
Decreases CPU and memory usage on routers.
Limits unnecessary route propagation.
Reference:
HCIA-Datacom Study Guide, Chapter: BGP Policy Control

Question: 186

In special scenarios, when advertising routes to an IBGP peer, a BGP device needs to set the next hop to its IP address to prevent blackhole routes. Which of the following commands can be run in this case?

A. peer next-hop-local
B. peer mpls-local-ifnet
C. peer private-nexthop
D. peer next-hop-invariable

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step

Blackhole Route Issue in BGP:

In IBGP, the next hop of a route learned from an EBGP peer is not changed by default. If an IBGP peer cannot reach the next hop, blackhole routes may occur.

To prevent this, the peer next-hop-local command is used to set the next hop to the local router's IP address.

Other Commands:

peer mpls-local-ifnet: Used in MPLS networks to set the next hop as the local interface address.

peer private-nexthop: Allows advertising private IP addresses as the next hop.

peer next-hop-invariable: Ensures the next hop is not modified during route advertisement, which is the opposite of what is required here.

Reference:

HCIA-Datacom Study Guide, Chapter: BGP Next Hop Behavior

Huawei BGP Configuration Commands

Question: 187

After BGP initiates a TCP connection, the ConnectRetry timer is disabled if the TCP connection is successfully established. If the TCP connection fails to be established, the device tries to reestablish the TCP connection when the ConnectRetry timer expires.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step

BGP ConnectRetry Timer:

The ConnectRetry timer is used when establishing a TCP connection between BGP peers.

If the connection is successful, the timer is disabled.

If the connection fails, the timer restarts, and the router attempts to reconnect after the timer expires.
Behavior:
This timer prevents excessive connection attempts and ensures efficient resource utilization. Reference:
HCIA-Datacom Study Guide, Chapter: BGP Session Establishment
Huawei BGP Timers and Session Management

Question: 188

Which of the following statements is false about the default processing of the next hop address when a BGP device advertises a route?

A. When advertising a locally originated route to an IBGP peer, a BGP device sets the next hop address to the IP address of its interface connected to the peer.
B. When advertising a non-labeled route received from an EBGP peer to an IBGP peer, a BGP device changes the next hop address to the IP address of its interface connected to the IBGP peer.
C. When advertising a route to an EBGP peer, a BGP device sets the next hop address to the IP address of its interface connected to the peer.
D. A BGP device does not change the next hop address of a route if the route is received from an IBGP peer and is to be sent to another IBGP peer.

Answer: B

Explanation:

Comprehensive and Detailed Step-by-Step
Default Next Hop Behavior in BGP:
Locally Originated Routes (A): The next hop is set to the advertising router's interface IP address.
Routes Received from EBGP (B): The next hop is not changed when these routes are advertised to IBGP peers. (This makes Option B false.)
Routes Advertised to EBGP (C): The next hop is set to the interface IP address of the advertising router.
Routes Advertised Between IBGP Peers (D): The next hop is not modified by default.
Incorrect Statement:
Option B is incorrect because the next hop of a non-labeled route received from an EBGP peer is not changed when it is advertised to IBGP peers.
Reference:
HCIA-Datacom Study Guide, Chapter: BGP Next Hop Behavior
Huawei BGP Routing Principles

Question: 189

After a BGP peer relationship is established between two ends, changing the router ID of one end resets the BGP peer relationship.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step

BGP Router ID Behavior:

The router ID uniquely identifies a BGP device in the network.

If the router ID is changed, it is equivalent to creating a new BGP instance, and the existing BGP session is reset to establish a new peer relationship.

Impact of Router ID Change:

BGP peering relies on stable identifiers. Changing the router ID disrupts the TCP session and requires reestablishing the relationship.

Reference:

HCIA-Datacom Study Guide, Chapter: BGP Peer Relationships

Huawei BGP Configuration Best Practices

Question: 190

By default, if no router ID is configured but multiple loopback interface addresses are configured, BGP selects the largest loopback interface address as the router ID.

A. TRUE

B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step

BGP Router ID Selection Rules:

If a router ID is not manually configured:

The largest loopback address is selected as the router ID.

If no loopback interfaces exist, the largest active physical interface address is selected.

Default Behavior:

This ensures that a stable and unique router ID is chosen automatically, even without manual configuration.

Reference:

HCIA-Datacom Study Guide, Chapter: BGP Router ID

Huawei Routing Device Configuration

Question: 191

In BGP, the Origin attribute is used to identify the origin of a route. Which of the following statements are true about the Origin attribute?

A. If a route is imported to BGP by the originator using the network command, the Origin attribute of the BGP route is displayed as i in the BGP routing table.

B. If a route is learned through EGP, the Origin attribute of this BGP route is displayed as ? in the BGP

routing table.

C. If multiple routes carry the same destination address but different Origin attributes and all other route selection conditions are the same, BGP selects the optimal route according to the Origin attribute in the following order: IGP > EGP > Incomplete.

D. This attribute is a well-known mandatory attribute.

Answer: A, C, D

Explanation:

Comprehensive and Detailed Step-by-Step

Definition of the Origin Attribute:

The Origin attribute is a well-known mandatory attribute used to indicate the source of a route. It has three possible values:

i (IGP): The route was injected into BGP using the network command.

e (EGP): The route was learned from an EGP protocol (not commonly used now).

? (Incomplete): The route was redistributed into BGP from another routing protocol.

Analysis of Each Option:

Option A: Correct. When the network command is used, the Origin is set to i (IGP).

Option B: Incorrect. Routes learned from EGP have the Origin attribute set to e (EGP), not ? (Incomplete).

Option C: Correct. BGP prefers routes with the Origin attribute in the order: IGP > EGP > Incomplete.

Option D: Correct. The Origin attribute is a well-known mandatory attribute.

Reference:

HCIA-Datacom Study Guide, Chapter: BGP Attributes

Huawei BGP Attribute Comparison

Question: 192

The Origin attribute is used to define the origin of BGP path information. There are three types of Origin attributes. Which of the following lists the Origin attributes in descending order of priority?

A. IGP > EGP > Incomplete

B. Incomplete > IGP > EGP

C. EGP > IGP > Incomplete

D. Incomplete > EGP > IGP

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step

BGP Origin Attribute Priority:

The Origin attribute determines the source of the route and affects the selection of the best path.

The priority order is as follows:

IGP (i): Highest priority. Routes injected using the network command.

EGP (e): Lower priority. Routes learned from the EGP protocol.

Incomplete (?): Lowest priority. Routes redistributed into BGP from other protocols.

Correct Option:

A: IGP > EGP > Incomplete is the correct order of priority.

Reference:
HCIA-Datacom Study Guide, Chapter: BGP Path Selection
Huawei BGP Routing Rules

Question: 193

A router uses an advanced ACL to filter dat
a. The ACL configuration is shown below. Which of the following statements is false about the configuration?
csharp
CopyEdit
[Huawei] acl 3001
[Huawei-acl-adv-3001] rule permit icmp source 192.168.1.3 0 destination 192.168.2.0 0.0.0.255

A. The ACL permits ICMP packets from host 192.168.1.3 to host 192.168.2.200.
B. The ACL denies ICMP packets from host 192.168.1.2 to network segment 192.168.2.0/24.
C. The ACL permits IP packets from host 192.168.1.3 to network segment 192.168.2.0/24.
D. The ACL is a numbered ACL whose number is 3001.

Answer: C

Explanation:

Comprehensive and Detailed Step-by-Step
Understanding the ACL Configuration:
ACL Number: 3001 is an advanced ACL.
Rule Description: The rule permits ICMP packets with the following criteria:
Source IP: 192.168.1.3 (host-specific).
Destination IP: 192.168.2.0/24 (entire subnet).
Analysis of Each Option:
Option A: True. ICMP packets from 192.168.1.3 to 192.168.2.200 (within the subnet) are permitted.
Option B: True. Packets from 192.168.1.2 (not matching the source) are denied by default.
Option C: False. The ACL only permits ICMP packets, not general IP packets.
Option D: True. The ACL is identified by the number 3001.
Reference:
HCIA-Datacom Study Guide, Chapter: ACL Filtering Rules
Huawei Advanced ACL Configuration

Question: 194

An ACL can be used to match routes or data, but cannot be used to match both the IP address prefix length and mask length.

A. TRUE
B. FALSE

Answer: B

Explanation:

Comprehensive and Detailed Step-by-Step

ACL Capabilities:

ACLs are versatile tools that can match:

Data packets: Using IP addresses, protocols, ports, etc.

Routes: ACLs can be used in route filtering, matching both IP prefix length and mask length.

Correct Statement:

ACLs can match both the IP address prefix length and mask length in route filtering applications. Reference:

HCIA-Datacom Study Guide, Chapter: ACL Applications

Huawei Route Filtering Using ACLs

Question: 195

A route-policy consists of one or more nodes. What is the maximum number of nodes in a routepolicy?

A. 4096
B. 65535
C. 256
D. 1024

Answer: D

Explanation:

Comprehensive and Detailed Step-by-Step

Route Policy Overview:

A route-policy is a policy-based routing tool that filters and modifies routing information.

It consists of one or more nodes, and each node can specify match conditions and apply actions to routes.

Maximum Number of Nodes:

The maximum number of nodes supported in a route-policy is 1024, allowing flexibility in route filtering and control.

Reference:

HCIA-Datacom Study Guide, Chapter: Route-Policy Configuration

Huawei Route-Policy Configuration Guidelines

Question: 196

A company has a stack consisting of three switches that are running properly. The master switch restarts due to a fault. Given this, which of the following statements are true?

A. Before the restart of the original master switch is complete, the original slave switch is specified as the new standby switch.
B. After the original master switch restarts, it becomes the new master switch.
C. After the original master switch restarts, the original standby switch becomes a new slave switch.
D. Before the restart of the original master switch is complete, the original standby switch becomes the new master switch.

Answer: B, D

Explanation:

Comprehensive and Detailed Step-by-Step

Switch Stack Overview:

In a switch stack, roles are assigned as master, standby, and slave.

The master handles configuration and control, while the standby serves as the backup master.

Behavior During Master Restart:

D: If the master switch restarts, the standby switch immediately takes over as the new master.

B: After the original master switch restarts, it becomes the master again (default behavior).

Incorrect Options:

A: The slave switch is not promoted to standby before the master switch restart completes.

C: The original standby switch remains the master after the restart.

Reference:

HCIA-Datacom Study Guide, Chapter: Switch Stack Management

Huawei Switch Stack Role Transition Rules

Question: 197

Link aggregation is a common network technology. Which of the following are advantages of link aggregation?

A. Load balancing
B. Improved reliability
C. Increased link bandwidth
D. Route backup

Answer: A, B, C

Explanation:

Comprehensive and Detailed Step-by-Step

Definition of Link Aggregation:

Link aggregation combines multiple physical links into a single logical link to improve bandwidth and reliability.

Advantages of Link Aggregation:

Load Balancing (A): Traffic is distributed across all aggregated links to improve performance.

Improved Reliability (B): If one link fails, traffic is rerouted through other links.

Increased Link Bandwidth (C): Bandwidth is effectively the sum of all aggregated links.

Incorrect Option:

Route Backup (D): Route backup is a feature of routing protocols, not link aggregation.

Reference:

HCIA-Datacom Study Guide, Chapter: Link Aggregation Fundamentals

Huawei Link Aggregation Configuration

Question: 198

On an STP network, only the designated port processes inferior BPDUs. On an RSTP network, a port with any

role processes inferior BPDUs.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step
STP and RSTP Behavior:
STP: Only the designated port processes inferior BPDUs. Inferior BPDUs are received from downstream switches and typically indicate topology changes.
RSTP: All ports, regardless of role (root, designated, or alternate), process inferior BPDUs to ensure faster convergence.
Behavior Difference:
RSTP improves convergence by allowing any port to process inferior BPDUs, unlike STP, which restricts this function to designated ports.
Reference:
HCIA-Datacom Study Guide, Chapter: STP vs RSTP
Huawei STP and RSTP Configuration Comparison

Question: 199

The Internet Assigned Numbers Authority (IANA) allocates Class D addresses to IPv4 multicast. An IPv4 address is 32 bits long, and the four most significant bits of a Class D address are 1110.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step
IPv4 Address Classes:
Class D: Reserved for multicast and defined by the first four bits as 1110.
Address range: 224.0.0.0 to 239.255.255.255.
Structure:
Class D does not include a network/host division, as it is designed for group communication rather than device identification.
Reference:
HCIA-Datacom Study Guide, Chapter: IPv4 Address Classes
Huawei Multicast Configuration Guidelines

Let me continue with the next set of questions!

Question: 200

In IGMPv1, querier selection depends on a multicast routing protocol, such as PIM. In IGMPv2 and IGMPv3, the interface with the largest IP address acts as the querier.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step

IGMPv1 Querier Selection:

In IGMPv1, there is no built-in querier election mechanism. Querier selection depends on the multicast routing protocol (e.g., PIM).

IGMPv2 and IGMPv3 Querier Selection:

In IGMPv2 and IGMPv3, the router interface with the largest IP address in the multicast group is selected as the querier.

Correct Statement:

The statement accurately describes the querier selection mechanisms in IGMP versions.

Reference:

HCIA-Datacom Study Guide, Chapter: IGMP Querier Mechanism

Huawei Multicast Protocol Comparison

Question: 201

IGMP has three versions. Different versions support different features. Which of the following features is supported by all versions?

A. Group-Specific Query message
B. Leave message
C. Report message
D. Specifying a multicast source

Answer: C

Explanation:

Comprehensive and Detailed Step-by-Step

Overview of IGMP Versions:

IGMPv1: The simplest version, supports basic group membership management using Report messages but lacks Leave messages.

IGMPv2: Adds features like Leave messages and Group-Specific Queries.

IGMPv3: Supports specifying multicast sources for Source-Specific Multicast (SSM).

Feature Supported by All Versions:

Report messages are fundamental to IGMP functionality and are supported in all three versions.

Features Not Supported by All Versions:
Group-Specific Queries (A): Supported in IGMPv2 and IGMPv3.
Leave Messages (B): Introduced in IGMPv2.
Specifying Multicast Sources (D): Introduced in IGMPv3.
Reference:
HCIA-Datacom Study Guide, Chapter: IGMP Overview
Huawei Multicast Protocol Details

Question: 202

In IPv6, interface IDs can be manually configured, automatically generated by the system, or generated based on the IEEE EUI-64 standard.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step
Interface ID in IPv6:
The interface ID is the last 64 bits of an IPv6 address, and it uniquely identifies an interface within a subnet.
It can be generated in three ways:
Manually Configured: Explicitly assigned by an administrator.
System-Generated: The operating system assigns a random or pseudo-random value.
EUI-64 Standard: Generated based on the MAC address of the interface.
Correct Statement:
IPv6 supports all three methods for generating interface IDs.
Reference:
HCIA-Datacom Study Guide, Chapter: IPv6 Address Structure
Huawei IPv6 Address Generation Methods

Let me continue with the next batch of questions!

Question: 203

Which of the following statements are true about the packet filtering firewall?

A. The packet filtering firewall can analyze associated packets to improve security.
B. The packet filtering firewall supports per-packet detection.
C. The packet filtering firewall can check application-layer data.
D. The packet filtering firewall filters data packets based on ACLs.

Answer: B, D

Explanation:

Comprehensive and Detailed Step-by-Step

Packet Filtering Firewall:

A packet filtering firewall inspects packets at the network and transport layers based on predefined rules (e.g., ACLs).

Correct Statements:

Option B: Packet filtering is performed on a per-packet basis.

Option D: ACLs define the filtering rules for traffic.

Incorrect Statements:

Option A: Packet filtering firewalls do not analyze associated packets. For this, a stateful firewall is required.

Option C: Application-layer inspection is not supported. This functionality is provided by applicationlayer firewalls.

Reference:

HCIA-Datacom Study Guide, Chapter: Firewalls

Huawei Firewall Packet Filtering Configuration

Question: 204

Typically, a protocol that occupies two ports during communication is called a multi-channel protocol. For such protocols, the ASPF function must be enabled on the firewall to ensure smooth setup of the data channel and reduce the risk of attacks. Which of the following protocols is not a multi-channel protocol?

A. FTP
B. SIP
C. SMTP
D. H.323

Answer: C

Explanation:

Comprehensive and Detailed Step-by-Step

Multi-Channel Protocols:

Multi-channel protocols use separate channels for control and data communication, requiring application-specific packet filtering (ASPF) to track sessions.

Examples:

FTP (A): Uses separate control and data channels.

SIP (B): Uses multiple ports for signaling and media.

H.323 (D): A VoIP protocol using multiple channels.

Single-Channel Protocol:

SMTP (C): A single-channel protocol for email communication, which does not require ASPF.

Reference:

HCIA-Datacom Study Guide, Chapter: Firewall ASPF Configuration

Huawei Multi-Channel Protocol Handling

Question: 205

If the interval for two consecutive packets of a TCP session reaching the firewall is longer than the aging time of the session, the firewall deletes the session information from the session table to ensure network

security.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step
Firewall Session Table:
A session table is used to track active sessions. If a session remains idle for longer than the configured aging time, it is removed to free resources and enhance security.
TCP Session Timeout:
If the interval between two packets exceeds the session timeout, the firewall deletes the session information, requiring the session to be re-established.
Reference:
HCIA-Datacom Study Guide, Chapter: Firewall Session Management
Huawei Firewall TCP Session Timeout Configuration

Question: 206

When a packet passes through a firewall, the firewall creates a session connection for the packet to guide subsequent forwarding of the packet. However, the firewall does not create session entries for all packets. For which of the following packets does the firewall not create session entries when the packet reaches the firewall?

A. ICMP ping packet
B. GRE packet
C. Subsequent fragment
D. UDP packet

Answer: C

Explanation:

Comprehensive and Detailed Step-by-Step
Session Creation in Firewalls:
Firewalls create session entries for packets requiring stateful inspection (e.g., TCP, UDP, ICMP, and GRE).
Subsequent fragments of large packets do not require new sessions. Instead, they are processed based on the session created for the first fragment.
Correct Option:
C (Subsequent Fragment): Does not trigger a new session entry as it belongs to an existing session.
Reference:
HCIA-Datacom Study Guide, Chapter: Firewall Session Handling
Huawei Fragmentation Processing in Firewalls

Question: 207

BFD provides fast fault detection independent of media and routing protocols. To use this mechanism to detect link connectivity, devices at both ends must support this feature.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step
Bidirectional Forwarding Detection (BFD):
BFD is a lightweight protocol that provides rapid fault detection on links independent of the underlying routing protocol or media.
Requirements for BFD:
Both devices at the endpoints of the link must support and configure BFD for it to operate.
If one device does not support BFD, the feature cannot be used for link detection.

Reference:
HCIA-Datacom Study Guide, Chapter: BFD Overview
Huawei BFD Configuration Guidelines

Question: 208

During routine maintenance, an enterprise administrator runs a command to check VRRP group information.
Which of the following statements is false about the command output?
yaml
CopyEdit
<HUAWEI> display vrrp verbose
Vlanif100 | Virtual Router 1 State: Master
Virtual IP: 10.1.1.100
Master IP: 10.1.1.2
PriorityRun: 120 PriorityConfig: 120
DR: None BDR: None MTU: 0
Preempt: YES Delay Time: 20s
Remain: -
Track: YES Priority Reduced: 20
Auth Type: MD5
BFD-session State: UP

A. Preemption is enabled for the VRRP group.
B. This VRRP group is an mVRRP group.
C. Authentication is enabled for the VRRP group.
D. The ID of the VRRP group is 1.

Answer: B

Explanation:

Comprehensive and Detailed Step-by-Step

Analysis of VRRP Configuration:

Option A (Preemption): Correct. The output explicitly states that preemption is enabled.

Option C (Authentication): Correct. Authentication is enabled using MD5.

Option D (Group ID): Correct. The VRRP group ID is explicitly stated as 1.

False Statement:

Option B (mVRRP Group): The output does not indicate this is an mVRRP (Multicast VRRP) group. This feature must be explicitly configured and is not enabled by default.

Reference:

HCIA-Datacom Study Guide, Chapter: VRRP Configuration Details

Huawei VRRP Command Reference

Question: 209

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:

Comprehensive and Detailed Step-by-Step

Single-Hop BFD:

Single-hop BFD is used for detecting link failures in directly connected devices.

It uses UDP port 3784 for communication.

Multi-Hop BFD:

For multi-hop scenarios, BFD uses UDP port 4784 to ensure end-to-end connectivity.

Incorrect Options:

TCP ports 3784 and 4784 (C and D): BFD does not use TCP.

Reference:

HCIA-Datacom Study Guide, Chapter: BFD Ports and Functions

Huawei BFD Configuration Details

Question: 210

When deploying a VRRP network, an enterprise administrator sets the virtual IP address to 192.168.1.254 and VRID to 1. Which of the following is the virtual MAC address after the network becomes stable?

A. 0000-5e01-0101
B. 0000-5e01-0254
C. 0000-5e00-0101
D. 0000-5e00-0254

Answer: C

Explanation:

Comprehensive and Detailed Step-by-Step

VRRP Virtual MAC Address Format:

The VRRP virtual MAC address is generated using the format: 0000-5e00-01XX, where XX represents the VRID in hexadecimal. Calculation:

VRID = 1 → Hexadecimal = 01.

Virtual MAC = 0000-5e00-0101.

Correct Option:

C (0000-5e00-0101).

Reference:

HCIA-Datacom Study Guide, Chapter: VRRP MAC Address Format

Huawei VRRP Configuration Guidelines

Let me continue with the next set of questions!

Question: 211

After the administrator of an enterprise deploys a DHCP server, employees complain that their clients cannot obtain IP addresses from the DHCP server. Which of the following may cause this problem?

A. Multiple DHCP servers are configured.
B. STP is enabled on the DHCP server.
C. The DHCP function is disabled by default, and the administrator forgets to enable the DHCP function.
D. DHCP clients and the DHCP server are on different network segments, and no DHCP relay agent is configured on the network.

Answer: C, D

Explanation:

Comprehensive and Detailed Step-by-Step

DHCP Configuration Issues:

Option C: The DHCP function is disabled by default on Huawei devices. If the administrator forgets to enable it, clients cannot receive IP addresses.

Option D: DHCP uses broadcast messages, which do not traverse routers. If the clients and DHCP server are on different subnets, a DHCP relay agent must be configured.

Other Options:

Option A: Multiple DHCP servers can coexist as long as their IP address pools do not overlap. This is unlikely to cause the problem.

Option B: STP does not affect DHCP operations unless it delays port activation, which is uncommon in this scenario.

Reference:

HCIA-Datacom Study Guide, Chapter: DHCP Configuration and Troubleshooting

Huawei DHCP Relay Agent Configuration

Question: 212

iMaster NCE-Campus can be used as an authentication server on a WLAN to authenticate STAs.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step

iMaster NCE-Campus Overview:

iMaster NCE-Campus is Huawei's network management platform for managing and controlling campus networks.

It supports user authentication, including STA (Station) authentication, in WLAN environments.

Authentication Modes Supported:

iMaster NCE-Campus integrates with AAA (RADIUS) servers to perform authentication for wireless clients.

Correct Statement:

The platform can function as an authentication server for WLAN STAs.

Reference:

HCIA-Datacom Study Guide, Chapter: WLAN Management with iMaster NCE-Campus

Huawei iMaster NCE-Campus Product Overview

Question: 213

On an enterprise WLAN where Portal authentication is deployed, an AC functions as an access device and communicates with a Portal server using the Portal protocol. Which of the following statements are true about the Portal protocol?

A. The HTTP or HTTPS protocol can be used as the Portal access or authentication protocol.
B. By default, the access device processes Portal protocol packets through port 2000.
C. By default, the device uses the destination port number 50100 to proactively send packets to the Portal server.
D. Portal protocol packets are transmitted over TCP.

Answer: A, D

Explanation:

Comprehensive and Detailed Step-by-Step

Portal Authentication:

Portal authentication uses a web-based mechanism where clients are redirected to a login page.

Communication between the access device and Portal server uses HTTP or HTTPS protocols.

Correct Statements:

Option A: HTTP or HTTPS is used for Portal authentication.

Option D: Portal protocol packets are transmitted over TCP for reliable communication.

Incorrect Statements:

Option B: Port 2000 is not the default for processing Portal packets.
Option C: The default destination port for communication with the Portal server is not 50100.
Reference:
HCIA-Datacom Study Guide, Chapter: WLAN Portal Authentication
Huawei WLAN Authentication Protocols

Question: 214

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 dat
a. 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, C, D

Explanation:

Comprehensive and Detailed Step-by-Step
HSB (Hot Standby Backup) Service:
HSB is used on WLAN controllers to back up critical information in real time to a backup device, ensuring service continuity in case of failures.
Backed-Up Information:
CAPWAP Tunnel Information (A): Essential for maintaining connections between APs and controllers.
AP Entries (C): Information about connected APs is backed up.
User Data Information (D): Includes client authentication and session details.
Incorrect Option:
DHCP Address Information (B): Not part of HSB's responsibilities, as DHCP is handled separately.
Reference:
HCIA-Datacom Study Guide, Chapter: WLAN Redundancy with HSB
Huawei HSB Configuration Guide

Question: 215

A wide area network (WAN) is a remote network that connects local area networks (LANs) or metropolitan area networks (MANs) in different areas for communication purposes. It is typically used to interconnect campus networks or data center networks.

A. TRUE
B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Step-by-Step
WAN Definition:

A WAN is a large-scale network spanning a wide geographical area, connecting smaller networks such as LANs and MANs.

Use Cases:

WANs are commonly used for interconnecting campus networks, branch offices, and data centers over long distances.

Correct Statement:

The statement accurately describes the purpose and scope of WANs.

Reference:

HCIA-Datacom Study Guide, Chapter: WAN Fundamentals

Huawei Networking Basics

Question: 216

On the campus OSPF network shown in the following figure, the interfaces connecting the five routers are GE interfaces, and their costs are not changed. The import-route command is run on R1. After the network converges, the route to the server at 192.168.1.0/24 is queried on R2. Which of the following is the cost of this route?

```
[R1]ospf 1
[RI -ospf-1]import-route direct type 1 cost 2
```

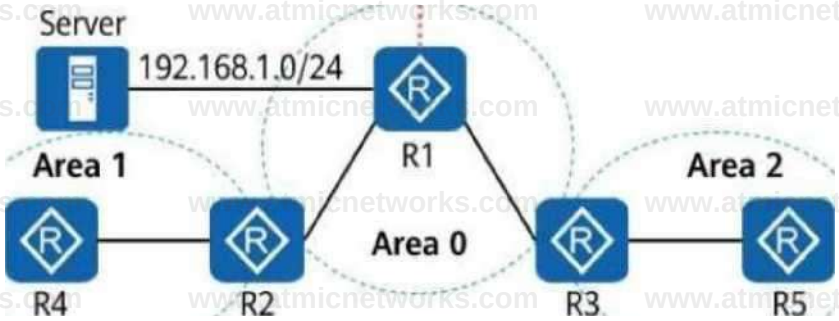

A. 2
B. 4
C. 3
D. 1

Answer: C

Explanation:

Understanding the Network Setup:

The OSPF network consists of three areas (Area 1, Area 0, and Area 2).

The server's subnet 192.168.1.0/24 is directly connected to R1 and imported into OSPF using the command:

import-route direct type 1 cost 2

The type 1 specifies the route as an OSPF Type 1 external route, meaning both the external cost and the internal OSPF cost will be included when calculating the total cost.

The external cost specified for the imported route is 2.

Cost Calculation to R2:

The interfaces between the routers are GE interfaces, and the default cost for GE interfaces is 1.

The path from R2 to the server goes through the following hops:

R2 → R1: Cost = 1 (intra-area link).

R1 → Server: External cost = 2 (specified in the import-route command).

Total cost to R2:

Intra-area cost (1) + External cost (2) = 3

Verification of Each Option:

Option A (2): Incorrect. This only accounts for the external cost, ignoring the internal OSPF cost.

Option B (4): Incorrect. This overestimates the cost by adding an extra hop.

Option C (3): Correct. The total cost is 3 (1 for the intra-area link + 2 for the external cost).

Option D (1): Incorrect. This ignores the external cost of the imported route.

Reference:

HCIA-Datacom Study Guide, Chapter: OSPF Route Types and Cost Calculation

Huawei OSPF Cost Configuration Details

Question: 217

On a broadcast IS-IS network shown in the following figure, a DIS needs to be elected to create and update pseudonodes. Which of the following routers is elected as the DIS?

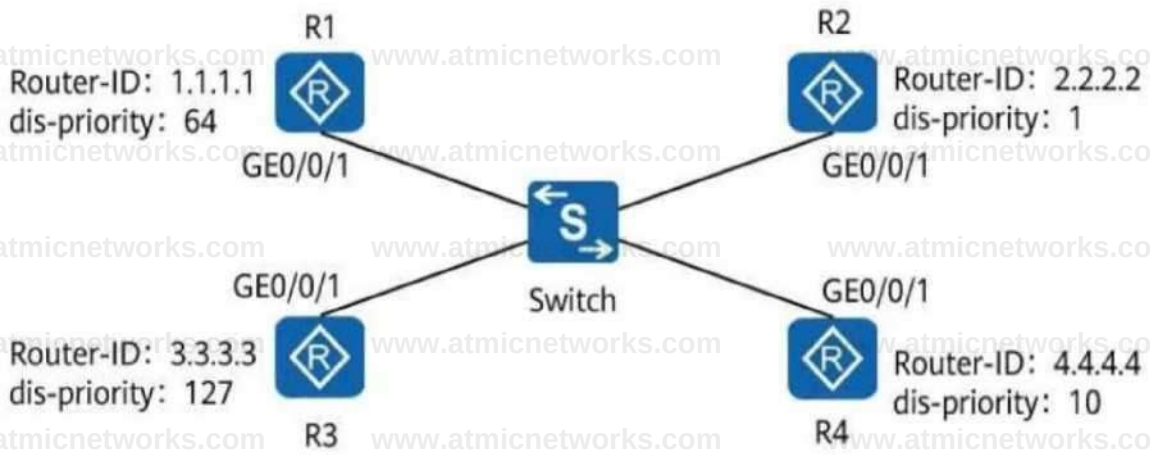

A. R1

B. R3

C. R4

D. R2

Answer: B

Explanation:

Understanding DIS Election in IS-IS:
On an IS-IS broadcast network, a Designated Intermediate System (DIS) is elected to create and update pseudonodes for efficient communication.
The election is based on the following criteria:
DIS Priority: The router with the highest priority is elected as the DIS.
Router ID (tie-breaker): If priorities are equal, the router with the highest Router ID is elected.
Analyzing the DIS Priorities and Router IDs:
R1: Priority = 64, Router ID = 1.1.1.1
R2: Priority = 1, Router ID = 2.2.2.2
R3: Priority = 127, Router ID = 3.3.3.3
R4: Priority = 10, Router ID = 4.4.4.4
Among the routers, R3 has the highest DIS priority (127), making it the DIS.
Correct Option:
Option B (R3): Correct, as R3 has the highest priority.
Option A (R1): Incorrect, as its priority is lower (64).
Option C (R4): Incorrect, as its priority is lower (10).
Option D (R2): Incorrect, as its priority is the lowest (1).
Reference:
HCIA-Datacom Study Guide, Chapter: IS-IS DIS Election
Huawei IS-IS Configuration and DIS Role Details

Question: 218

On the OSPF network shown in the figure, an adjacency has been established between R1 and R2. An engineer configures the commands in the figure on R2. In this case, which of the following routing entries may exist in the routing table of R2?

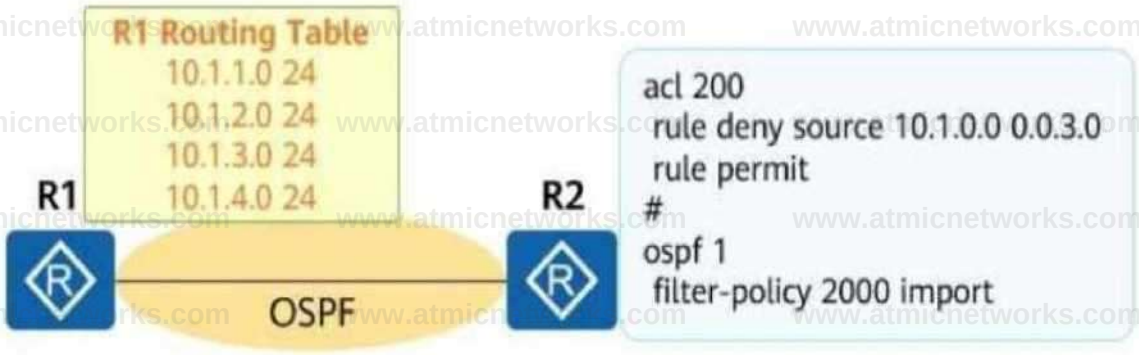

A. 10.1.4.0/24

B. 10.1.3.0/24

C. 10.1.2.0/24

D. 10.1.1.0/24

Answer: ABC

Explanation:

Configuration Analysis:

On R2, the following configuration has been applied:

acl 200

rule deny source 10.1.0.0 0.0.3.0

rule permit

#

ospf 1

filter-policy 2000 import

This configuration uses ACL 200 to filter routes during import into the OSPF routing table on R2.

Rule deny source 10.1.0.0 0.0.3.0: Blocks routes in the range 10.1.0.0/24 to 10.1.3.0/24 (inclusive).

Rule permit: Allows all other routes to be imported.

Impact of ACL 200 on Route Import:

10.1. 0.0/24 to 10.1.3.0/24: These subnets are explicitly denied by ACL 200 and will not appear in R2's routing table.

10.1.4. 0/24 and beyond: These subnets are permitted by the rule permit statement and will be imported into R2's routing table.

Routing Table Entries on R2:

Option A (10.1.4.0/24): Exists in R2's routing table because it is permitted.

Option B (10.1.3.0/24): Does not exist because it is denied by ACL 200.

Option C (10.1.2.0/24): Does not exist because it is denied by ACL 200.

Option D (10.1.1.0/24): Does not exist because it is denied by ACL 200.

Correct Options:

A (10.1.4.0/24)

Reference:

HCIA-Datacom Study Guide, Chapter: OSPF Route Filtering

Huawei ACL Configuration for Route Policies

Question: 219

On the OSPF network shown in the figure, R1, R2, and R3 run OSPF, and R1 advertises four VPN routes to OSPF. A filter-policy needs to achieve the following goal: R1's and R3's routing tables contain the routes to 192.168.3.0/24, but R2's routing table does not. Which of the following filterpolicies cannot meet this requirement?

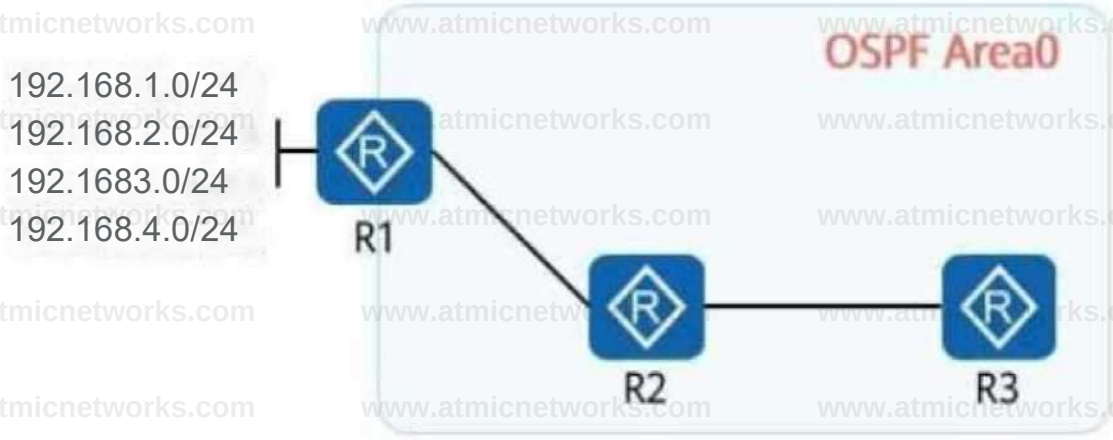

A. A filter-policy onR2 for filtering received routes
B. A filter-policy on R2 for filtering the routes to be advertised
C. A filter-policy on R1 for filtering the routes to be imported
D. A filter-policy on R1 for filtering the imported routes to be advertised

Answer: C

Explanation:
Goal Analysis:
Requirement:
The route to 192.168.3.0/24 must exist in R1's and R3's routing tables.
The route must not exist in R2's routing table.
This requires filtering to ensure the route is either:
Blocked on R2's routing table (via filtering on R2), or
Blocked before it is advertised to R2.
Analysis of Each Option:
Option A (Filter-policy on R2 for filtering received routes):
Applying a filter-policy on R2 to filter received routes will block the route from entering R2's routing table but still allow it to propagate to R3.
This meets the requirement.
Option B (Filter-policy on R2 for filtering the routes to be advertised):
Blocking the advertisement of routes from R2 to other routers does not affect the routes received by R2 itself.
This does not meet the requirement but does not affect the propagation to R3.
This is valid if the received route is blocked.
Option C (Filter-policy on R1 for filtering the routes to be imported):
If the route is filtered on R1 during the import phase, the route will not exist in R1's routing table and thus cannot be advertised to either R2 or R3.
This fails to meet the requirement because the route must exist in R1's and R3's routing tables.
Option D (Filter-policy on R1 for filtering the imported routes to be advertised):
Filtering routes on R1 before advertising to R2 will prevent R2 from receiving the route but allow R1 to advertise the route to R3.
This meets the requirement.
Correct Option:
C (Filter-policy on R1 for filtering the routes to be imported): This will prevent the route from existing in both

R1 and R3, violating the stated requirement.

Reference:

HCIA-Datacom Study Guide, Chapter: OSPF Route Filtering

Huawei OSPF Configuration and Filtering Methods