

HCIA-Datacom V1.0 Exam

Huawei H12-811 V1.0

Version Demo

Total Demo Questions: 72

Total Premium Questions: 724

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Data Communication Network Basics	25
Topic 2, Huawei Router and Switch Basics	252
Topic 3, Network Layer Protocols and IP Routing Basics	184
Topic 4, IP Services and Applications	80
Topic 5, WLAN Basics	59
Topic 6, WAN Basics	19
Topic 7, Network Security Basics	51
Topic 8, IPv6 Basics	36
Topic 9, Huawei Datacom Solutions	17
Topic 10, Mix Questions	1
Total	724

QUESTION NO: 1

If the broadcast address of a network is 172.16.1.255, what might its network address be? ()

- A. 172.16.2.0
- B. 172.16.1.128
- C. 172.16.1.1
- D. 172.16.1.253

ANSWER: B

Explanation:

172.16.1.128 might be the network address because it forms a valid subnet whose directed broadcast address is 172.16.1.255. With a /25 prefix length (subnet mask 255.255.255.128), the final octet is divided into two ranges: 0–127 and 128–255. The network beginning at 172.16.1.128 therefore has all host bits set to 1 at 172.16.1.255, making that address its broadcast address. Its usable host-address range is 172.16.1.129 through 172.16.1.254. This result follows the IPv4 subnetting rule that a network address has all host bits set to 0, while the directed broadcast address for that same subnet has all host bits set to 1. CIDR prefix notation and the associated address-block boundaries are defined in [RFC 4632](#). Huawei's IPv4 addressing documentation also describes using subnet masks to distinguish network and host portions of an IPv4 address: [Huawei IPv4 Addresses documentation](#).

QUESTION NO: 2

Which of the following is not included in the VRP system login method? ()

- A. Telnet
- B. Netstream
- C. SSH
- D. web

ANSWER: B

Explanation:

Netstream is not included in the VRP system login methods. On Huawei VRP-based devices, administrators can access and manage the command-line interface through local console access and remote management protocols such as Telnet and SSH/STelnet. Devices may also provide web-based management through the HTTP or HTTPS server, depending on the product model and enabled services. These mechanisms authenticate a user and establish an administrative management session with the device.

NetStream, by contrast, is a traffic-flow statistics and monitoring feature. It collects information about network flows—such as source and destination addresses, ports, protocol information, packet counts, and byte counts—and can export those records to a collector for analysis. It therefore supports traffic visibility, accounting, capacity analysis, and security operations rather than interactive administrator login. Huawei's VRP documentation groups user access methods under management and login configuration, while NetStream is documented as a flow-analysis capability. See Huawei's [VRP configuration documentation](#) and the [Huawei NetStream configuration documentation](#) for further details.

QUESTION NO: 3

ForwardD. Delay exists when an STP port transitions between which of the following states?

- A. Forwarding-Di sableD. Learning-Forwarding
- B. Learning-Forwarding

- C. Disabled-Blocking
- D. Blocking-Listening
- E. Listening-Learning

ANSWER: B E

Explanation:

In classic IEEE 802.1D STP, Forward Delay is the timer used while a designated port progresses toward forwarding traffic after topology calculation. The port first enters the listening state, where it participates in STP processing but does not learn source MAC addresses or forward user frames. It remains in listening for one Forward Delay interval before moving to learning. In the learning state, the switch builds its MAC address table from received frames but still does not forward user traffic. It remains there for a further Forward Delay interval before entering forwarding. Therefore, the relevant state transitions are *Listening-Learning* and *Learning-Forwarding*. These timed stages provide time for BPDUs to propagate and for switches to converge on a loop-free topology before data forwarding begins. Huawei documentation describes STP port-state transitions and the role of the forward-delay timer in preventing transient Layer 2 loops; see the [Huawei STP configuration documentation](#). The underlying STP behavior is also standardized in [IEEE 802.1D](#).

QUESTION NO: 4

To ensure that the same data flow is forwarded on the same physical link, which load sharing method does the Eth-Trunk adopt? ()

- A. Flow-based load sharing
- B. Packet-based load sharing
- C. Load sharing based on application layer information
- D. Load sharing based on data packet inbound interface

ANSWER: A

Explanation:

Flow-based load sharing is correct because an Eth-Trunk selects a member interface by applying a hash calculation to packet fields that identify a flow, such as source and destination MAC addresses, IP addresses, protocol information, or port numbers, depending on the configured load-balancing mode. Packets with the same relevant header-field values produce the same hash result and are therefore mapped to the same physical member link throughout the flow. This preserves in-order packet delivery, which is important for upper-layer protocols and applications that expect packets from one conversation to arrive sequentially. Eth-Trunk load balancing distributes multiple independent flows across available member interfaces while maintaining this per-flow forwarding consistency. If a member link becomes unavailable, the device recalculates forwarding using the remaining active members, allowing the trunk to retain connectivity and share traffic efficiently. Huawei's Eth-Trunk documentation describes link aggregation and its traffic-load-sharing behavior, while Huawei's VRP configuration documentation provides the relevant Ethernet interface and trunk configuration concepts. See [Huawei product documentation](#) and [Huawei VRP documentation](#).

QUESTION NO: 5

Which of the following message types does DHCP contain?

- A. DHCP DISCOVER
- B. DHCP REQUEST
- C. DHCP OFFER
- D. DHCP ROLLOVER

ANSWER: A B C

Explanation:

DHCP uses defined protocol messages to let a client obtain and maintain IPv4 configuration dynamically. **DHCP DISCOVER** is the initial client message used to locate available DHCP servers, typically sent as a broadcast when the client has no usable address. A server that can provide configuration responds with **DHCP OFFER**, which proposes an address lease and associated parameters such as the subnet mask, default gateway, DNS server, and lease duration. The client then sends **DHCP REQUEST** to indicate the offered address it is selecting and to formally request that configuration. This exchange forms the first three stages of the commonly described Discover, Offer, Request, Acknowledgment process. The DHCP protocol also defines additional messages for lease confirmation, renewal, release, decline, and server notifications, but the listed three are foundational DHCP message types used in normal address allocation. RFC 2131 specifies these DHCP message types and the client-server message flow. Huawei documentation likewise describes Discover, Offer, Request, and Acknowledgment as the standard DHCP address-assignment interaction. See [RFC 2131](#) and [Huawei DHCP Overview](#).

QUESTION NO: 6

Regarding the corresponding relationship between the access control list number and type, the following description is correct ()

- A. The basic access control list number range is 1000-2999
- B. The number range of the Layer 2 access control list is 4000-4999
- C. The number range of the interface-based access control list is 1000-2000
- D. Advanced access control list number range is 3000-4000

ANSWER: B


Explanation:

The number range of the Layer 2 access control list is 4000-4999. On Huawei VRP devices, the ACL number identifies both an ACL instance and its category. Numbered Layer 2 ACLs use the 4000 through 4999 range and are intended for traffic classification and filtering based on Layer 2 information, such as source and destination MAC addresses, VLAN-related fields, and protocol information carried in Ethernet frames. This numbering convention lets the device interpret the ACL type correctly when it is created with a numbered ACL command and ensures that the available rule-matching fields are appropriate for Layer 2 forwarding scenarios.

Using the documented number range is important in operational configurations because an ACL number is not merely an arbitrary label: it selects the ACL category and its supported rule structure. Huawei's product documentation portal provides the VRP configuration references and ACL command documentation used for particular device and software versions. See the [Huawei Enterprise documentation portal](#) and the [Huawei technical documentation site](#) for the applicable VRP ACL configuration guide.

QUESTION NO: 7

As shown in the figure below, the network. All switches enable the STP protocol. Disable SWA's G0/0/2 port to configure the BPDU sending function. SWC's G0/0/1 reconverges to become the root port. Which of the following statements is correct about this process? () (Multiple Choice Questions)



- A. SWB forwards TCN BPDU to SWA
- B. SWB sends configuration BPDU with TCA set to SWC
- C. SWC sends TCN BPDU to SWB
- D. SWA sends configuration BPDU with TC set

ANSWER: A B C D

Explanation:

When SWA's G0/0/2 link is disabled, SWC loses its original path toward the root bridge. After STP reconverges and SWC's G0/0/1 becomes its root port, this port-role transition constitutes a topology change. In classic IEEE 802.1D STP, SWC notifies the root bridge by transmitting a Topology Change Notification (TCN) BPDU through its newly selected root port to SWB. SWB acknowledges the notification toward SWC by sending a Configuration BPDU with the Topology Change Acknowledgment (TCA) flag set, and it relays the TCN upstream on its own root-port path to SWA. Once SWA, the root bridge, receives the TCN, it sets the Topology Change (TC) flag in Configuration BPDUs sent from its designated ports for the topology-change propagation interval. This causes bridges receiving those BPDUs to use the shortened forwarding-database aging behavior required during convergence, reducing the chance that stale MAC-address entries lead to temporary misforwarding. Therefore, all four listed BPDU actions are part of the normal 802.1D STP topology-change notification and acknowledgment sequence. See Huawei's [STP documentation](#) and the IEEE [802.1Q bridging standard](#).

QUESTION NO: 8

The router uses the sub-interface as the gateway of the host, and the P address of the gateway is 10.0.12.2. Which of the following commands can achieve this requirement? ()

- A. Interface GigabitEthernet0/0/0.20
dot1q termination vid 10
ip address 10.0.12.2 255.255.255.0
arp broadcast enable
- B. Interface GigabitEthernet0/0/0.10
dot1q termination vid 10
ip address 10.0.12.2 255.255.255.0
arp broadcast enable
- C. Interface GigabitEthernet0/0/0.10
dot1q termination vid 20
ip address 10.0.12.2 255.255.255.0
arp broadcast enable
- D. Interface GigabitEthernet0/0/0.20
dot1q termination vid 20
ip address 10.0.12.2 255.255.255.0
arp broadcast enable

ANSWER: C D

Explanation:

The configurations using `GigabitEthernet0/0/0.10 with dot1q termination vid 20` and `GigabitEthernet0/0/0.20 with dot1q termination vid 20` can both provide the gateway address `10.0.12.2/24` for hosts carried in VLAN 20. On Huawei VRP devices, a Layer 3 Ethernet sub-interface can terminate traffic for a specified 802.1Q VLAN through the `dot1q termination vid` command. The IP address configured on that sub-interface then acts as the routed default gateway for hosts in the associated VLAN/subnet. Enabling broadcast ARP processing allows the sub-interface to receive and respond to host ARP broadcasts, which is necessary for hosts to resolve the gateway's MAC address before sending IPv4 traffic to it.

The numeric suffix of the sub-interface name, such as `.10` or `.20`, is locally significant and does not itself determine the VLAN that is terminated. The effective VLAN association is established by `dot1q termination vid 20`. Therefore, either sub-interface number is valid when the VLAN termination, IPv4 gateway address, mask, and ARP broadcast handling are configured as shown. Huawei's official documentation portal provides VRP configuration and command references at [Huawei Enterprise Documentation](#); additional Huawei enterprise product documentation is available through [Huawei InfoSupport Documentation](#).

QUESTION NO: 9

In the STP protocol, which of the following factors will affect the election of the root switch? (*)

- A. IP address of the switch
- B. Switch interface ID
- C. Switch interface bandwidth port
- D. Switch priority
- E. MAC address of the switch

ANSWER: D E

Explanation:

In STP, the root switch is elected by comparing each participating switch's Bridge ID (BID). The BID is constructed from the configurable switch priority together with the switch's MAC address (with the extended system ID incorporated in implementations using VLAN-aware STP). Therefore, **Switch priority** and **MAC address of the switch** directly determine the result of root-bridge election. STP selects the device with the numerically lowest BID as the root switch. Because priority is the higher-order, administrator-configurable part of the BID, it is normally the value network engineers adjust when they need to intentionally place the root switch at a particular point in the Layer 2 topology. If two switches retain the same priority, their MAC addresses provide the tie-breaking portion of the BID, and the numerically lower MAC address is elected root. This deterministic comparison ensures that all switches converge on a single root bridge. Huawei STP documentation describes the root bridge as the bridge with the smallest bridge ID and explains the role of bridge priority in that election. See [Huawei STP Configuration Guide](#) and the [IEEE 802.1D standard page](#).

QUESTION NO: 10

The G0/0/1 interface configuration information of router A is shown in the figure. Which of the following statements are correct? () (Multiple Choice Questions)

```
#
acl number 3000
 rule 5 deny 17
 rule 10 deny 89
 rule 15 deny 6
#
interface GigabitEthernet0/0/1
 traffic-filter inbound acl 3000
#
```

- A. The interface will not forward the received FTP packets
- B. This interface can establish OSPF neighbor relationships with other routers
- C. This interface will not forward ICMP packets
- D. This interface will not forward received SNMP packets

ANSWER: A D

Explanation:

The correct statements are “The interface will not forward the received FTP packets” and “This interface will not forward received SNMP packets.” The interface configuration shown applies an inbound traffic-filter that references an ACL. The ACL contains deny rules matching FTP and SNMP traffic, so packets matching those rules are discarded as they enter G0/0/1 and therefore cannot be forwarded by the router. FTP control traffic conventionally uses TCP port 21, while standard SNMP requests use UDP port 161; ACL rules can match packets by protocol and Layer 4 port number. Because the filter is applied in the inbound direction, the word “received” is important: the policy is evaluated for traffic arriving on this interface before normal forwarding processing sends the traffic toward another interface. Huawei VRP ACLs are commonly used with interface traffic-filter commands to control which packets are permitted or denied at an interface. Huawei’s ACL documentation describes using rules to match protocol and port characteristics and applying an ACL to traffic for packet filtering; see the [Huawei VRP ACL configuration documentation](#) and the [traffic-filter command reference](#).

QUESTION NO: 11

Now there are the following four segments: 10.24.0.0/24, 10.24.1.0/24, 10.24.2.0/24, 10.24.3.0/24, which route can point to these four network segments at the same time? (*)

- A. 10.24.1.0/23
- B. 0.0.0.0/0
- C. 10.24.0.0/22
- D. 10.24.0.0/21

ANSWER: B C D

Explanation:

0.0.0.0/0, 10.24.0.0/22, and 10.24.0.0/21 can each route traffic toward all four listed destination networks. A route prefix matches a destination when the destination address shares the prefix bits of the route. The default route 0.0.0.0/0 has no fixed prefix bits, so it is a matching route for every IPv4 destination, including each of the four 10.24.x.0/24 networks. In forwarding, it is normally used only when no more-specific route is available, according to longest-prefix matching.

The summary 10.24.0.0/22 uses mask 255.255.252.0 and spans 10.24.0.0 through 10.24.3.255. It therefore exactly aggregates the four consecutive /24 networks: 10.24.0.0/24, 10.24.1.0/24, 10.24.2.0/24, and 10.24.3.0/24. The broader 10.24.0.0/21, with mask 255.255.248.0, spans 10.24.0.0 through 10.24.7.255. Because all four specified /24 networks fall within that range, this route also points to all of them. This follows CIDR address aggregation and forwarding-prefix rules described in [RFC 4632](#) and the IPv4 forwarding requirements in [RFC 1812](#).

QUESTION NO: 12

Which of the following messages is used for IPv6 stateless address autoconfiguration? (*)

- A. NS
- B. RA
- C. NA
- D. RS

ANSWER: B D

Explanation:

RA and **RS** are used in IPv6 Stateless Address Autoconfiguration (SLAAC). A host uses an RS (Router Solicitation) message to request configuration information from IPv6 routers on the local link, particularly when it starts up and does not want to wait for the router's periodic advertisements. The router responds by sending an RA (Router Advertisement) message.

The RA is the essential source of SLAAC parameters. It can provide one or more IPv6 prefixes and associated flags, including the Autonomous Address-Configuration flag. When that flag is set for a prefix, the host can combine the advertised prefix with an interface identifier to create its own IPv6 address without requiring a server to allocate the address. Router Advertisements can also supply the default-router information and may carry related network configuration data through options, such as the MTU and recursive DNS server information.

Therefore, RS initiates or accelerates acquisition of router configuration information, while RA delivers the prefix and routing information that enables the host's stateless configuration process. This behavior is defined by IPv6 Neighbor Discovery and the SLAAC specification: [RFC 4861](#) and [RFC 4862](#).

QUESTION NO: 13

How to enter in the vP operating system OView of SPF area 0? (*)

- A. (Huawei-ospf-1) area 0 enable
- B. (Huawei-ospf-1) area 0
- C. (Huawei) ospf area 0
- D. (Huawei-ospf-1) area 0.0.0.0

ANSWER: B D

Explanation:

In Huawei VRP, the `area` command is entered from the OSPF process view, indicated by the `(Huawei-ospf-1)` prompt. The area identifier accepts either a decimal area number or an IPv4-style dotted-decimal area ID. Therefore, `area 0` enters the OSPF area view for the backbone area by using its decimal representation. Likewise, `area 0.0.0.0` enters exactly the same OSPF area view by using the dotted-decimal representation of the backbone area ID. These two forms are equivalent because OSPF represents area IDs as 32-bit values; zero in decimal and 0.0.0.0 in dotted-decimal notation identify the same 32-bit value. After entering this view, area-specific configuration, such as enabling OSPF on interfaces with the `network` command, can be performed. Huawei's OSPF command documentation describes the use of the `area` command in OSPF process view, and the OSPF standard defines area ID 0 as the backbone area. See [Huawei VRP documentation](#) and [RFC 2328](#).

QUESTION NO: 14

The route entry 10.0.0.24/29 may be aggregated from which of the following subnet routes? (*)

- A. 10.0.0.23/30
- B. 10.0.0.28/30
- C. 10.0.0.26/30
- D. 10.0.0.24/30

ANSWER: B C D

Explanation:

The prefix 10.0.0.24/29 uses the mask 255.255.255.248, so its aggregate address block spans 10.0.0.24 through 10.0.0.31. A /30 prefix has mask 255.255.255.252 and represents a four-address block. Therefore, **10.0.0.24/30** represents the range 10.0.0.24–10.0.0.27, and **10.0.0.28/30** represents 10.0.0.28–10.0.0.31. Both ranges are entirely contained within the /29 aggregate. Together, these two contiguous /30 networks exactly form the 10.0.0.24/29 range.

10.0.0.26/30 also belongs to the same aggregatable range. Applying a /30 mask to 10.0.0.26 yields the canonical network prefix 10.0.0.24/30, because the final two host bits are cleared by the mask. Thus, it identifies an address within the 10.0.0.24–10.0.0.27 /30 subnet and can be summarized by 10.0.0.24/29. This follows standard CIDR prefix and route-aggregation behavior, in which an aggregate must cover the component prefixes' address space. See [RFC 4632: Classless Inter-domain Routing](#) and Huawei's [enterprise documentation portal](#).

QUESTION NO: 15

OSThe I reference model from high-level to bottom-level is ()

- A. Application layer, transport layer, network layer, data link layer, physical layer
- B. Application layer Session layer Presentation layer Transport layer Network layer Data link layer Physical layer
- C. Application layer Presentation layer Session layer Network layer Transport layer Data link layer Physical layer
- D. Application layer Presentation layer Session layer Transport layer Network layer Data link layer Physical layer


ANSWER: D

Explanation:

The correct sequence is "Application layer Presentation layer Session layer Transport layer Network layer Data link layer Physical layer." The OSI (Open Systems Interconnection) reference model organizes network communication into seven layers, ordered from the layer closest to end-user applications down to the layer responsible for transmitting raw signals over a physical medium. At the top, the Application layer provides network services to application processes. The Presentation layer handles data representation functions, such as translation, formatting, compression, and encryption. The Session layer establishes, manages, and terminates dialogs between communicating applications. The Transport layer provides end-to-end transport services, including segmentation and reliability mechanisms where applicable. The Network layer handles logical addressing and packet forwarding between networks. The Data Link layer enables node-to-node delivery over a local link, including framing and use of physical addressing. Finally, the Physical layer transmits bits as electrical, optical, or radio signals through the medium. This top-down order is the standard OSI layering model used to describe and troubleshoot network functions. See the ISO overview of the OSI model at [ISO/IEC 7498-1](#) and the Cisco Networking Academy explanation at [Cisco Networking Basics](#).

QUESTION NO: 16

The network shown in the figure below. The host has an ARP cache. Which of the following statements are correct?
(Multiple Choice Questions)



- A. The following entries may appear in the ARP cache of host A 10.0.12.2MAC-C
- B. The router needs to be configured with a static route. Otherwise, host A and host B cannot communicate in both directions
- C. Host A and Host B can communicate in two directions
- D. The following entry 11.0.12.1 MAC-B exists in the ARP cache of host A

ANSWER: A C

Explanation:

The entries “The following entries may appear in the ARP cache of host A 10.0.12.2MAC-C” and “Host A and Host B can communicate in two directions” are correct. ARP operates within a local Layer 2 broadcast domain and resolves an IPv4 address to the MAC address of the next Layer 2 destination. When host A sends traffic to a host on the remote 11.0.12.0 network, it does not ARP for the remote host’s address. Instead, it identifies its configured default gateway as the next hop and uses ARP to resolve the gateway interface address 10.0.12.2. Therefore, host A can create an ARP-cache mapping from 10.0.12.2 to MAC-C. The router then forwards the IP packet between its directly connected interfaces. Because both host subnets are directly connected to the router, the router has routes for both networks automatically through its interface configuration. Return traffic follows the same routing principle: host B sends packets for host A to its local default gateway, and the router forwards them to host A’s subnet. This permits bidirectional communication when the hosts have the appropriate IP addressing, masks, and default gateways. ARP’s local-link behavior is defined in [RFC 826](#); Huawei’s networking documentation and product resources are available through [Huawei Enterprise Support](#).

QUESTION NO: 17

After the network administrator uses the Traceroute function on the router device, in the data packet sent by the router, the value of the Protocol field in the IPv4 header is? ()

- A. 17
- B. 2
- C. 6
- D. 1

ANSWER: D

Explanation:

The correct value is **1**, which identifies ICMP in the IPv4 Protocol field. Huawei devices use ICMP messages for the Traceroute/tracert diagnostic process. The trace sends probes with progressively increased TTL values so that each traversed Layer 3 device can be identified. When a probe’s TTL reaches zero at an intermediate router, that router sends an ICMP Time Exceeded message back to the source. The destination may also respond with an ICMP Echo Reply when it receives the final probe, depending on the specific trace behavior and implementation. Therefore, packets generated by a router as ICMP diagnostic or TTL-expiry responses carry the IPv4 Protocol field value assigned to ICMP: decimal 1. This value is not a port number; it is the IPv4 header’s next-protocol identifier, allowing the receiving host to deliver the IP payload to ICMP for processing. The Internet Assigned Numbers Authority registry and the IPv4 specification define protocol number 1 as ICMP. See the [IANA Protocol Numbers registry](#) and [RFC 792: Internet Control Message Protocol](#).

QUESTION NO: 18

What is the minimum IPv6 link MTU required for an IPv6 link? () (Single-choice question)

- A. 576 bytes
- B. 1000 bytes
- C. 1280 bytes
- D. 1500 bytes

ANSWER: C


Explanation:

The correct answer is **1280 bytes**. Every IPv6 link must support an MTU of at least 1280 bytes. This requirement ensures that IPv6 packets of the minimum supported size can be delivered without relying on IPv4-style router fragmentation.

If a link cannot natively support 1280-byte frames, the link layer must provide a mechanism allowing IPv6 packets of this size to traverse the link. IPv6 nodes can use Path MTU Discovery to identify a smaller path MTU where applicable. The IPv6 minimum link MTU requirement is specified in [RFC 8200](#).

QUESTION NO: 19

As shown in the figure below, are the following statements correct? () (Multiple Choice Questions)



- A. The broadcast address of host A and host B are the same
- B. Host A can ping host B
- C. Host A and Host B cannot obtain each other's MAC address
- D. The following entry 10.0.12.5 MAC-B exists in the ARP cache of host A

ANSWER: B D

Explanation:

Host A can ping host B is correct because the hosts are connected within the same Layer 2 broadcast domain. Although their configured subnet masks result in different local broadcast-address calculations, Host A treats 10.0.12.5 as reachable on its local link and sends an ARP Request as an Ethernet broadcast. Host B receives that request on the shared LAN and returns an ARP Reply containing MAC-B. After this address resolution, Host A can encapsulate ICMP Echo Request packets in Ethernet frames addressed directly to MAC-B, allowing ping connectivity.

The following entry 10.0.12.5 MAC-B exists in the ARP cache of host A is also correct. ARP maps a known IPv4 address on a local Ethernet segment to the destination device's MAC address. Once Host A receives Host B's ARP Reply, it dynamically learns and stores the mapping between 10.0.12.5 and MAC-B in its ARP cache. That cached mapping is then used for subsequent unicast frames sent to Host B until the entry ages out or is refreshed. This behavior follows the ARP operation defined in [RFC 826](#); Huawei documentation also describes ARP as the mechanism used to obtain MAC addresses from IP addresses on a LAN: [Huawei ARP documentation](#).

QUESTION NO: 20

Which of the following statements about IPv6 address configuration are correct? (*)

- A. IPv6 supports DHCPv6 for address configuration
- B. IPv6 address supports automatic configuration in multiple ways
- C. IPv6 addresses can only be configured manually
- D. IPv6 supports stateless automatic configuration

ANSWER: A B D

Explanation:

IPv6 provides several standardized mechanisms for assigning addresses, so “IPv6 address supports automatic configuration in multiple ways” is correct. A host can form an address through Stateless Address Autoconfiguration, using Router Advertisement information such as the advertised prefix, while also deriving an interface identifier according to the applicable addressing rules. This is the basis for the correct statement “IPv6 supports stateless automatic configuration.” IPv6 networks can additionally use DHCPv6, making “IPv6 supports DHCPv6 for address configuration” correct. DHCPv6 can provide stateful address assignment, delegated prefixes, and/or supplementary configuration parameters, depending on the deployment design and Router Advertisement flags. These mechanisms allow administrators to select autonomous addressing, centrally managed addressing, or a combination of both to meet operational requirements. RFC 4862 defines IPv6 Stateless Address Autoconfiguration, including the role of Router Advertisements and autonomous prefix configuration: [RFC 4862](#). DHCPv6 behavior and address/configuration-information assignment are specified in [RFC 8415](#).

QUESTION NO: 21

UDP cannot guarantee the reliability of data transmission, and does not provide functions such as message sequencing and flow control. It is suitable for flows that do not require high transmission reliability, but require high transmission speed and delay ().

- A. True
- B. False

ANSWER: A


Explanation:

True is correct. The intended wording is that UDP is suitable where applications prioritize high transmission speed and low delay, while not requiring transport-layer reliability. UDP is a connectionless transport protocol with a compact header and no connection establishment process. It does not perform acknowledgements, retransmissions, in-order delivery, or receiver-driven flow control at the UDP layer. Avoiding those mechanisms reduces protocol overhead and can help an application deliver time-sensitive traffic without waiting for lost data to be recovered.

This behavior is useful when an application can tolerate occasional packet loss, can handle recovery itself if needed, or considers an old packet less useful than a newly received one. Typical examples include real-time voice and video, interactive gaming, and some telemetry services. The application remains responsible for choosing appropriate safeguards, such as sequence information, buffering, congestion-aware behavior, or error recovery where its service requires them. UDP therefore trades the reliability and ordered byte-stream service of a connection-oriented transport for simple, low-overhead datagram delivery. UDP's basic connectionless datagram model is defined in [RFC 768](#), and deployment guidance for UDP applications is provided by [RFC 8085](#).

QUESTION NO: 22

The network is shown in the figure below. The following statement correct? () (Multiple Choice Questions)



- A. Host A and Host c can ping through
- B. All hosts can ping each other
- C. Host B and Host C can ping through
- D. Host A and Host B cannot ping through

ANSWER: A C D

Explanation:

“Host A and Host c can ping through,” “Host B and Host C can ping through,” and “Host A and Host B cannot ping through” are correct because IP reachability is determined independently by each host’s subnet-mask-based route lookup. In the shown addressing design, Host C regards the addresses of both Host A and Host B as being on its directly connected network and can therefore resolve their MAC addresses using ARP and send frames directly. Host A likewise treats Host C as local, and Host B treats Host C as local, so the two corresponding ping exchanges can complete successfully. For communication between Host A and Host B, however, the respective subnet calculations do not identify the peer as directly connected. Since the topology does not provide a usable default gateway or another route for that traffic, neither host has a valid Layer 3 forwarding path to the other host’s subnet. This illustrates that hosts attached to the same Layer 2 segment can still have selective IP connectivity when their IP addresses and subnet masks create different on-link determinations. Huawei’s IP addressing documentation describes how masks identify network portions of IPv4 addresses, while Huawei routing documentation explains that forwarding requires a matching directly connected or configured route. See [Huawei IP Addressing and Subnetting documentation](#) and [Huawei IP Routing documentation](#).

QUESTION NO: 23

Which of the following fields does the MPLS header include? (*)

- A. EXP
- B. TTL
- C. Label
- D. Tos

ANSWER: A B C

Explanation:

The MPLS label-stack entry, commonly called the MPLS shim header, is a fixed 32-bit structure inserted between the Layer 2 header and the encapsulated packet. It includes a 20-bit **Label** field, which identifies the forwarding equivalence class used by MPLS-capable routers to make label-switching decisions. It also includes a 3-bit **EXP** field. In current MPLS

terminology, these three bits are formally called the Traffic Class (TC) field, but EXP remains widely used in Huawei training material and operational discussions. These bits can carry QoS-related marking and support differentiated forwarding behavior across an MPLS domain. The header further includes an 8-bit **TTL** field. Each label-switching router decrements this value while forwarding the labeled packet, supporting loop prevention and traceroute-style diagnostics in an MPLS network. In addition to these listed fields, every MPLS label-stack entry contains a one-bit Bottom of Stack indicator, which identifies the final label in the stack. The label-stack entry format is specified in [RFC 3032](#); the updated Traffic Class terminology for the historical EXP field is defined by [RFC 5462](#).

QUESTION NO: 24

Which of the following statements about IPv6 address configuration is correct? (*)

- A. IPv6 supports stateless automatic configuration
- B. IPv6 addresses can only be configured manually
- C. IPv6 supports multiple automatic address-configuration methods
- D. IPv6 supports DHCPv6 for address configuration

ANSWER: A C D

Explanation:

IPv6 provides several valid address-configuration mechanisms, allowing hosts to obtain addresses without relying exclusively on manual configuration. Stateless Address Autoconfiguration (SLAAC) enables a host to create an address after receiving Router Advertisement information, including the applicable IPv6 prefix, from an IPv6 router. This mechanism is standardized in [RFC 4862](#) and is a fundamental IPv6 deployment method.

IPv6 also supports more than one automatic configuration method. In addition to SLAAC, a network can use DHCPv6, either to assign IPv6 addresses statefully or to provide supplementary configuration data, such as DNS server information. DHCPv6 behavior and message exchanges are specified in [RFC 8415](#). Therefore, IPv6 supports stateless automatic configuration, supports multiple automatic address-configuration approaches, and supports DHCPv6 for address configuration. These mechanisms allow network administrators to select an approach that fits the addressing, control, and operational requirements of the network.

QUESTION NO: 25

As shown in the broadcast network. OSPF runs on four routers. And in the same area. Same network segment. OSPF will automatically elect one DR and multiple BDRs to achieve a better backup effect. (True or False)



- A. True
- B. False

ANSWER: B

Explanation:

False is correct. On an OSPF broadcast multi-access network, such as an Ethernet LAN, routers on the same OSPF subnet hold a designated-router election to reduce the number of adjacencies that must be formed and maintained. The election produces exactly one Designated Router (DR) and exactly one Backup Designated Router (BDR) for that network segment. All remaining eligible routers are DR Other routers. The BDR is the standby for the DR and takes over the DR role if the current DR becomes unavailable; OSPF does not elect multiple BDRs on the same broadcast segment. The fact that four routers participate in the same OSPF area and IP network does not change this election rule. Router interface priority is considered first, and router ID is used as a tie-breaker among routers with equal priority. A router configured with priority 0 is ineligible to become either DR or BDR. This behavior is defined for broadcast networks in the OSPF specification and is also reflected in Huawei OSPF configuration documentation. See [RFC 2328: OSPF Version 2](#) and [Huawei OSPF documentation](#).

QUESTION NO: 26

The information shown in the figure is the port status information displayed on a switch running STP. According to this information, what is wrong in the following description? () (Single-choice question)

RTID	Port	Role	STP State	Protection
0	GigabitEthernet0/0/1	DESI	FORWARDING	NONE
0	GigabitEthernet0/0/2	DESI	FORWARDING	NONE
0	GigabitEthernet0/0/13	DESI	FORWARDING	NONE
0	GigabitEthernet0/0/21	DESI	FORWARDING	NONE
0	GigabitEthernet0/0/22	DESI	FORWARDING	NONE
0	GigabitEthernet0/0/23	DESI	FORWARDING	NONE

- A. This switch is the root switch in the network
- B. The priority of this switch is 0
- C. This network may only contain this switch
- D. This switch may be connected to 6 other switches

ANSWER: B

Explanation:

"The priority of this switch is 0" is the incorrect statement because the STP status information identifies the switch's bridge priority as 32768, not 0. In STP, a bridge ID is formed from the bridge priority and the bridge MAC address; the bridge with the lowest bridge ID becomes the root bridge. A default STP bridge priority of 32768 is commonly shown in switch STP status output when no lower configured priority has been applied. Therefore, the displayed priority must be read from the bridge information itself rather than inferred from the switch's root-bridge role or its port roles. A value of 0 can be configured as a bridge priority in supported implementations, but it is not the priority shown for this switch in the figure. Huawei's STP documentation describes bridge priority as part of the bridge ID and explains its role in root bridge election. See the [Huawei VRP STP configuration documentation](#) and the [IEEE 802.1D STP standard information](#).

QUESTION NO: 27

The MAC address table of existing switches is shown in the figure. Which of the following statements are correct? () (Single-choice question)

```

Huawei>display mac-address
MAC address table of slot 0:

```

MAC Address	VLAN/ VSI/VSI	PVID/VLAN	CVLAN	Port	Type	LDP/LDP-ID MAC-Tunnel
5489-9811-0b49 1	-	-	-	Eth0/0/3	static	-

Total watching items on slot 0 displayed = 1

```

Huawei>display mac-address
MAC address table of slot 0:

```

MAC Address	VLAN/ VSI/VSI	PVID/VLAN	CVLAN	Port	Type	LDP/LDP-ID MAC-Tunnel
5489-989d-1d30 1	-	-	-	Eth0/0/1	dynamic	0/-
5489-9885-18a8 1	-	-	-	Eth0/0/2	dynamic	0/-

Total watching items on slot 0 displayed = 2

- A. The MAC address 5489-9811-0b49 is manually configured by the administrator
- B. The MAC address 5489-989d-1d30 is manually configured by the administrator
- C. After the switch restarts, all MAC addresses need to be relearned. The MAC address 5489-9885-18a8 is manually configured by the administrator
- D. The MAC address 5489-9885-18a8 is manually configured by the administrator

ANSWER: A

Explanation:

The MAC address 5489-9811-0b49 is manually configured by the administrator. In a Huawei switch MAC address table, the entry type identifies how the device obtained and maintains the mapping between a MAC address, VLAN, and outbound interface. A *static* MAC address entry is created through administrator configuration rather than being learned automatically from received Ethernet frames. It is therefore used when traffic for a known endpoint must always be forwarded through a specified interface, providing a fixed forwarding association. Static MAC entries are configuration-based entries and can be retained after a restart when the configuration is saved, unlike ordinary dynamically learned forwarding entries, which are maintained in the running MAC table and are learned again as traffic arrives. The table shown identifies 5489-9811-0b49 with the static/manual entry characteristic, so it is the address that was configured by the administrator. Huawei documentation describes static MAC address configuration and the distinction between static and dynamically learned MAC forwarding entries in its VRP MAC-address feature documentation: [Huawei Enterprise Support Documentation](#) and [Huawei VRP static MAC address command reference](#).

QUESTION NO: 28

In an IPv4 network, AP supports static and DHCP methods to obtain IP addresses. ()

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. In an IPv4 WLAN deployment, an access point requires an IP address for management and for communication with the AC or other network-management services. Huawei APs can be configured to use a manually assigned, static IPv4 address, which is useful where addressing is planned and fixed IP assignments are required. They can also act as DHCP clients and automatically obtain their IPv4 address, subnet mask, default gateway, and potentially other network parameters from a DHCP server. DHCP-based addressing is commonly used for large-scale AP rollouts because it reduces the manual configuration required for each device. After receiving network connectivity and an address, the AP can use the configured discovery and management mechanisms to establish its connection to the WLAN controller. DHCP is the standard protocol for dynamically supplying IPv4 configuration information to hosts, as specified in [RFC 2131](#). Huawei's enterprise support portal provides WLAN product documentation and configuration guidance for AP network access and address assignment at [Huawei Enterprise Support](#).

QUESTION NO: 29

Which of the following protocols is OSPF based on?

- A. UDP
- B. HTTP
- C. TCP
- D. IP

ANSWER: D

Explanation:

IP is correct. OSPF is a link-state interior gateway protocol that operates directly over IP, rather than using TCP or UDP as a transport-layer protocol. OSPF packets are encapsulated in IP packets with the IP protocol number 89. This direct use of IP lets OSPF exchange Hello packets, link-state advertisements, and other routing-control packets without the connection establishment, port numbers, and reliability mechanisms associated with TCP. Instead, OSPF defines its own neighbor-state procedures, acknowledgements, flooding controls, and retransmission mechanisms to ensure reliable distribution of link-state information within an OSPF routing domain. In IPv4 deployments, OSPF commonly uses the multicast addresses 224.0.0.5 (AllSPFRouters) and 224.0.0.6 (AllDRouters) for relevant control-plane exchanges; these are carried as IP multicast packets using protocol 89. The OSPFv2 specification explicitly states that all OSPF protocol packets are carried directly by IP and identifies 89 as the assigned IP protocol number. Huawei documentation likewise describes OSPF as an IP-based dynamic routing protocol. See [RFC 2328, OSPF Version 2](#) and [Huawei OSPF Configuration Overview](#).

QUESTION NO: 30

When deploying ACL on Huawei equipment, which of the following descriptions is correct is (.)*

- A. ACL can match the TCP/UDP port number of the packet, and can specify the range of port numbers
- B. The same ACL can be called under multiple interfaces
- C. When ACL defines rules, it can only proceed in the order of 10, 20, 30
- D. ACL cannot be used to filter OSPF traffic, because OSPF traffic does not use UDP protocol encapsulation
- E. When calling ACL under the interface, it can only be applied to the outbound direction

ANSWER: A B

Explanation:

ACL can match the TCP/UDP port number of the packet, and can specify the range of port numbers is correct because Huawei advanced ACL rules can match Layer 4 protocol characteristics, including TCP or UDP source and destination port information. Port matching supports operators such as equality, ranges, and comparisons, allowing a policy to identify a service or group of services rather than only an IP address. This makes ACLs useful for traffic classification, packet filtering, and security-policy deployment.

The same ACL can be called under multiple interfaces is also correct. An ACL is a reusable set of numbered rules defined globally on the device. After it is created, the same ACL can be referenced by multiple features and applied to more than one interface when the required matching policy is identical. Reusing an ACL provides consistent policy enforcement and simplifies administration: changes to the ACL rule set are reflected wherever that ACL is referenced. Huawei documentation describes ACLs as rule-based traffic matching tools and documents their application by traffic-policy and interface-related features. See Huawei's [VRP ACL configuration documentation](#) and the [Huawei enterprise documentation portal](#).

QUESTION NO: 31

ip route-static 10.0.2.2 255.255.255.255 10.0.12.2 preference 20, the correct statement about this command is ()

- A. The NextHop of this route is 10.0.12.2
- B. The priority of this route is 20
- C. This route can guide the forwarding of packets with the destination IP address of 10.0.2.2.
- D. This route can guide the forwarding of packets with the destination IP address 10.0.12.2


ANSWER: A B C

Explanation:

Interpreting the command with its intended spaces, it is `ip route-static 10.0.2.2 255.255.255.255 10.0.12.2 preference 20`. The destination address and the all-ones mask define a host static route: only packets whose destination is exactly 10.0.2.2 match this route. The address following the mask, 10.0.12.2, is the next-hop address. After the router selects this route, it performs the required recursive lookup or directly resolves that next hop through an appropriate outgoing interface before forwarding traffic. The `preference 20` parameter sets the route preference to 20. In Huawei VRP routing, route preference is used when multiple routes to the same destination are candidates; the route with the preferred value is selected according to the platform's route-selection rules. Therefore, the statements identifying 10.0.12.2 as the next hop, identifying 20 as the route preference, and stating that the route forwards traffic addressed to 10.0.2.2 are all correct. Huawei's command documentation for static-route configuration is available through [Huawei Enterprise Support](#); the relevant IP-routing concepts are also covered in the [Huawei VRP documentation](#).

QUESTION NO: 32

The network shown in the figure below. From a security perspective, router A refuses to receive OSPF packets and ICMP packets received from interface G0/0/1. Which of the following commands can achieve this requirement? () (Multiple Choice Questions)



- A. acl number 3000
Rule 5 deny 89
Rule 10 deny icmp

Interface GigabitEthernet0/0/1Traffic-filter inbound acl 3000#
- B. acl number 3000
Rule 5 deny 89
Rule 10 deny 1

Interface GigabitEthernet0/0/1Traffic-filter inbound acl 3000#
- C. acl number 2000
Rule 5 deny 89
Rule 10 deny 1

Interface GigabitEthernet0/0/1Traffic-filtet inbound acl 2000#

```
D. acl number 3000
Rule 5 deny ospf
Rule 10 deny icmp
#
Interface GigabitEthernet0/0/1Traffic-filter inbound acl 3000#
```

ANSWER: A B D

Explanation:

The configurations using an advanced ACL in the 3000 range and applying it inbound with `traffic-filter inbound acl 3000` meet the requirement. An inbound traffic filter evaluates packets as they enter GigabitEthernet0/0/1, before they are accepted for further processing by Router A. Therefore, denying the relevant IP protocol types at that interface prevents Router A from receiving OSPF control traffic and ICMP traffic from that direction.

In IPv4 protocol-number notation, OSPF is IP protocol 89 and ICMP is IP protocol 1. Thus, the rules `deny 89` and `deny 1` identify the required packets directly. Huawei VRP advanced ACLs also accept protocol keywords, so `deny ospf` and `deny icmp` express the same filtering intent using named protocol types. These forms can be combined in an ACL and bound to the inbound interface traffic filter as shown. The assigned protocol values are defined by IANA: [IP Protocol Numbers](#). OSPF packet processing is specified in [RFC 2328](#).

QUESTION NO: 33

It is known that there are two entries in the routing table of a router. If the router wants to forward packets with the destination address of 9.1.4.5, the correct one of the following statements is ()

- A. Choose the second item as the best matching item, because the cost of the RIP protocol is small
- B. The second item is selected as the best matching item, because the route is a more accurate match relative to the destination 9.1.4.5
- C. Select the second item as the best match, because the speed of Ethernet0 is faster than the speed of serial0
- D. Chooses the first item as the best matching item, because the priority value of the ospf protocol is higher

ANSWER: B

Explanation:

The second item is selected as the best matching item, because the route is a more accurate match relative to the destination 9.1.4.5. Huawei routers apply longest-prefix matching when forwarding an IP packet: among the active routes whose destination prefixes contain the packet's destination address, the route with the greatest prefix length is selected. A longer prefix represents a smaller, more specific destination network and therefore provides the most precise forwarding instruction for 9.1.4.5. For example, a route covering 9.1.4.0 with a longer mask is preferred over a broader route that also contains that address. Route protocol preference and route cost are used when routes compete for the same destination prefix and mask; they do not replace longest-prefix matching between different prefixes. This behavior ensures that traffic follows the route intended for the most specifically defined network. The underlying IP forwarding rule is also described in [RFC 1812, Section 5.2.4](#) and in [RFC 4632, Section 3.1](#).

QUESTION NO: 34

The configuration information of a certain router is shown in the figure. Which of the following statements are correct? () (Single-choice question)

```
#
aaa
authentication-scheme default
```

```
authentication-scheme huawei
authentication-mode radius
authorization-scheme default
authorization-scheme huawei
accounting-scheme default
domain default
domain default admin
domain huawei
authentication-scheme huawei
authorization-scheme huawei
local-user huawei password cipher 123456
local-user huawei@huawei password cipher 654321
#
```

- A. The domain with the domain name huawei does not use the billing plan
- B. Use the username huawei for authentication. The password needs to be 654321
- C. The authentication method adopted for the domain with the domain name huawei is local authentication
- D. The authorization method for the domain with the domain name huawei is local authorization

ANSWER: D

Explanation:

The authorization method for the domain with the domain name huawei is local authorization is correct. Within the `domain huawei` view, the configuration applies `authorization-scheme huawei`. The displayed `authorization-scheme huawei` has no explicit `authorization-mode` command beneath it. In Huawei VRP AAA, local authorization is the default authorization mode of an authorization scheme when another authorization mode has not been configured. Therefore, after a user is associated with the `huawei` domain, the router uses the `huawei` authorization scheme and performs authorization against locally configured user information and attributes. This is separate from the domain's authentication scheme: the configuration explicitly selects RADIUS authentication under `authentication-scheme huawei`, while the authorization scheme remains at its local default. The local-user entries shown in AAA provide the local user database used by local AAA functions. Huawei's enterprise documentation portal provides VRP AAA configuration references and command descriptions at [Huawei Enterprise Documentation](#). Additional official product and configuration resources are available through [Huawei Enterprise Support](#).

QUESTION NO: 35

The following statement about static routing is wrong ()

- A. Cannot automatically adapt to changes in network topology
- B. Manual configuration by network administrator
- C. Low system performance requirements
- D. Routers need to exchange routing information

ANSWER: D

Explanation:

“Routers need to exchange routing information” is the wrong statement for static routing. A static route is a manually configured forwarding entry: the administrator specifies a destination prefix, mask or prefix length, next hop or outbound interface, and optionally route preference. Once configured, the router installs and uses that route according to normal route-selection and forwarding rules. The route does not result from routers advertising networks to one another, calculating paths, or maintaining neighbor relationships. Those activities belong to dynamic routing protocols, which exchange routing information and can react automatically to learned topology changes. Static routing is therefore especially useful for small, stable networks, stub networks, default-route scenarios, and tightly controlled paths where an administrator wants predictable forwarding behavior. If a link or topology changes, the static configuration generally must be adjusted by an administrator or supplemented with mechanisms such as route tracking; this does not turn the route into an exchanged routing update. Huawei’s routing documentation distinguishes manually configured static routes from protocol-learned routes, while the IP routing requirements in RFC 1812 describe routers’ forwarding and routing responsibilities. See [Huawei Enterprise documentation](#) and [RFC 1812](#).

QUESTION NO: 36

Which of the following refers to network communication? (*)

- A. Use instant messaging software (such as: QQ. WeChat) to chat with friends
- B. Use a computer to access the company's official website
- C. Use computer to watch videos online
- D. Download emails from the company mailbox to your computer

ANSWER: A B C D

Explanation:

Network communication is the exchange of data between endpoints through a network, typically using Internet protocol suites and application-layer services. “Use instant messaging software (such as: QQ. WeChat) to chat with friends” is network communication because chat clients send and receive messages through network-connected servers. “Use a computer to access the company's official website” involves a browser requesting web resources and receiving responses using HTTP or HTTPS over an IP network. “Use computer to watch videos online” is also network communication: the device receives a continuous stream or sequence of video data from an online service. “Download emails from the company mailbox to your computer” likewise transfers messages from a mail server to a mail client through standard email-access protocols such as POP3 or IMAP. Although these activities have different user purposes—messaging, web browsing, video delivery, and email retrieval—they all depend on data being transmitted across a network between a local terminal and remote systems. The Internet protocol architecture that supports these communications is described in [RFC 1122](#), while the web request-and-response model used for website access is defined by [RFC 9110](#).

QUESTION NO: 37

As shown in the figure, there are three data packets captured by the administrator on the network. Which of the following statements is incorrect? (Single choice question)

Source destination protocol info

10.0.12.110.0.12.2TCP50190>telnet [SYN] seq=0win=8192Len=0mss=146010.0.12.2 10.0.12.1 TCP telnet:50190 [SYN,ACK] seq=0 Ack=1 win=8192 Len=0mss=1460

10.0.12.110.0.12.2 TCP50190>telnet(ACK)seq=1Ack=1win=8192len=0

- A. None of these three data packets contain application layer data
- B. These three data packets represent the three-way handshake process of TCP
- C. The telnet client uses port 50190 to establish a connection with the server

D. The IP address of the telnet server is 10.0.12.1. The IP address of the telnet client is 10.0.12.2

ANSWER: D

Explanation:

The statement “*The IP address of the telnet server is 10.0.12.1. The IP address of the telnet client is 10.0.12.2*” is the incorrect statement. The initial TCP segment has source address 10.0.12.1 and shows the source port as 50190, with destination address 10.0.12.2 and destination service *telnet*. Telnet conventionally uses the well-known TCP port 23, while port 50190 is an ephemeral client-side port. Therefore, 10.0.12.1 is the Telnet client initiating the connection, and 10.0.12.2 is the Telnet server receiving the connection on its Telnet service port. The following segment, sent from 10.0.12.2 to 10.0.12.1, contains SYN and ACK flags, confirming that the destination of the first SYN acts as the responding server. The final ACK returns from 10.0.12.1 to 10.0.12.2 and completes the TCP three-way handshake. This directionality follows TCP’s client-server connection establishment model: the client sends SYN to a server’s listening well-known port, and the server replies with SYN-ACK. See the [TCP specification \(RFC 793\)](#) and the [IANA Telnet service-port registration](#).

QUESTION NO: 38

Which of the following statements about IPv6 address configuration are correct? (*)

- A. IPv6 supports stateless automatic configuration
- B. IPv6 address supports automatic configuration in multiple ways
- C. IPv6 supports DHCPv6 for address configuration
- D. IPv6 addresses can only be configured manually

ANSWER: A B C

Explanation:

IPv6 supports several complementary address-configuration mechanisms, allowing hosts to obtain usable addresses with different levels of automation and administrative control. “IPv6 supports stateless automatic configuration” is correct because Stateless Address Autoconfiguration (SLAAC) uses Router Advertisement messages to provide a network prefix and related parameters; the host then creates its own IPv6 address, commonly using a stable or privacy-oriented interface identifier. “IPv6 address supports automatic configuration in multiple ways” is also correct: IPv6 hosts can use SLAAC, DHCPv6, or a combination in which router advertisements supply default-router information while DHCPv6 provides addresses and/or additional configuration data. “IPv6 supports DHCPv6 for address configuration” is correct because stateful DHCPv6 can allocate and manage IPv6 addresses, and stateless DHCPv6 can provide supplementary settings such as DNS recursive-server information without assigning the address itself. These methods enable IPv6 networks to select a deployment model appropriate to operational requirements, ranging from low-touch endpoint onboarding to centrally administered address assignment. The authoritative SLAAC behavior is specified in [RFC 4862](#), while DHCPv6 address-assignment and configuration procedures are specified in [RFC 8415](#).

QUESTION NO: 39

The DHCP request message is sent by broadcast

- A. True
- B. False

ANSWER: B

Explanation:

False is correct because whether a DHCPREQUEST message uses broadcast depends on the client’s DHCP state, so it cannot be stated as an unconditional rule. During initial address selection, after receiving one or more DHCPOFFER

messages, a client broadcasts DHCPREQUEST to identify the selected server and notify any other offering servers that their offers were not selected. A client also broadcasts DHCPREQUEST when attempting to confirm a previously used address after restart or when rebinding because its original lease server has not responded. However, while renewing an active lease, the client normally unicasts DHCPREQUEST directly to the DHCP server that granted the lease. Therefore, the statement is too broad: broadcast is a common and important DHCPREQUEST transmission method, but not the only one. RFC 2131 explicitly defines broadcast DHCPREQUEST behavior for the SELECTING, INIT-REBOOT, and REBINDING states, and unicast behavior for the RENEWING state. See the DHCP client state descriptions in [RFC 2131, section 4.4.1](#) and the DHCPREQUEST message details in [RFC 2131, section 4.3.2](#).

QUESTION NO: 40

On Huawei devices, if you use aaA authentication to authorize. When the remote server does not respond, you can authorize from the network device side ()

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. Huawei AAA supports using local authorization on the network device as a fallback when remote authorization cannot be completed, provided that the relevant AAA scheme and authorization method have been configured to permit local processing. Remote AAA services, such as RADIUS or HWTACACS, are commonly used to centrally manage user privileges and command permissions. However, network operations must be able to continue during a server outage or connectivity failure. A device can therefore use its locally configured user information, privilege levels, and authorization rules when the remote authorization server does not respond. This fallback behavior improves service availability while still allowing administrators to control which local accounts and permissions exist on the device. The exact behavior depends on the configured authorization scheme and the order of authorization methods, so local authorization must be deliberately enabled and populated with appropriate local user privileges before it can act as a reliable backup. Huawei documents AAA as the framework that provides authentication, authorization, and accounting services, including local and remote AAA methods. See the [Huawei AAA configuration documentation](#) and the [Huawei AAA overview](#).

QUESTION NO: 41

As shown in the figure below. Which network device can host A and host B use to achieve communication? () (Single-choice question)



- A. HUB
- B. Hub
- C. Layer 2 switch
- D. router


ANSWER: D

Explanation:

The **router** is the appropriate device because it forwards traffic between separate IP networks. When Host A needs to send data to Host B on a different network segment, Host A sends the packet to its configured default gateway. The router receives that packet on the interface connected to Host A's network, examines the destination IP address, consults its routing information, and forwards the packet through the interface leading toward Host B's network. Each router interface operates in a distinct Layer 3 broadcast domain and can provide the gateway IP address for the hosts attached to that network. This routing function is what enables end-to-end communication when the hosts have different network prefixes or are connected through separate network segments, as illustrated. IPv4 routers are required to make forwarding decisions using the destination address and routing information, a behavior defined in [RFC 1812](#). Huawei enterprise routers are designed to provide this Layer 3 inter-network forwarding capability; see Huawei's [enterprise router product information](#).

QUESTION NO: 42

The network shown in the figure below assumes that all routers can enable OSPF at the same time. Which router is the DR in this network? () (Single-choice question)



- A. Router B
- B. Router D
- C. Router A
- D. Router C

ANSWER: A

Explanation:

Router B is the DR. On an OSPF multi-access broadcast or NBMA segment, routers elect a designated router when they establish adjacency on that shared network. Because the question specifies that all routers enable OSPF at the same time, there is no pre-existing DR or BDR whose non-preemptive status would affect the result. The election therefore compares all eligible routers simultaneously. OSPF first uses the interface priority; assuming the default or equal priority values shown or implied in this topology, the deciding value is the router ID. Router B has the highest router ID among the eligible routers in the figure, so it wins the DR election. The router with the next-highest eligible election value becomes the BDR. Router IDs are compared as unsigned 32-bit values rather than as ordinary dotted-decimal text, which is why the numerically greatest ID is decisive when priorities are equal. This process reduces the number of full adjacencies required on a shared segment and centralizes LSA exchange through the DR. The OSPF DR/BDR election procedure is defined in [RFC 2328](#); Huawei's OSPF configuration documentation is available through the [Huawei Enterprise Support documentation portal](#).

QUESTION NO: 43

The administrator hopes that RTA. uses static routing and preferentially sends data packets to RTB's Loopback0 interface through G0/0/1 port, while G0/0/2 is used as a backup. How to configure on RTA to achieve this requirement ()

- A. ip route-static 10.0.2.2255.255.255.255 10.0.12.2ip
Route-static 10.0.2.2255.n255.255.255 10.0.21.2 preference 40

B. ip route-static 10.0.2.2255.255.255.25 5 10.0.12.2 preference 70 ip route-static/n10.0.2.2255.255.25510.0.21.2

C. ip route-static 10.0.2.2255.255.255.255 10.0.12.2 ip
Route-static 10.0.2.2255.255.255.255 10.0.21.2

D. ip route-static 10.0.2.2255.255.255.255 10.0.12.2 ip
Route-static 10.0.2.2255.255.255.255

E. ip route-static 10.0.2.2 255.255.255.255 10.0.12.2
ip route-static 10.0.2.2 255.255.255.255 10.0.21.2 preference 70

ANSWER: E

Explanation:

ip route-static 10.0.2.2 255.255.255.255 10.0.12.2
ip route-static 10.0.2.2 255.255.255.255 10.0.21.2 preference 70 correctly creates two host routes to RTB's Loopback0 address. The first route uses the next hop on the path connected through G0/0/1, 10.0.12.2. Because no preference is explicitly configured for this route, it uses the default preference of 60 for IPv4 static routes in VRP. The second route uses the next hop on the path connected through G0/0/2, 10.0.21.2, and assigns it a preference of 70. In Huawei VRP, a lower route preference is more preferred, so the route through 10.0.12.2 is installed and selected first. The route through 10.0.21.2 remains a less-preferred candidate and is selected when the preferred route is no longer available. The /32 mask is important because it specifies a route solely to the Loopback0 host address rather than to a broader network. This is the standard floating-static-route design: retain a primary static route at the normal preference and assign a larger preference value to the backup route. See Huawei's [VRP product documentation](#) and [Huawei routing configuration documentation](#) for static-route configuration and route-preference behavior.

QUESTION NO: 44

Telnet is a module that implements the Telnet protocol that comes with Python ()

A. True

B. False

ANSWER: A

Explanation:

True is correct in the Python-version context typically used for HCIA-Datcom network automation material. Python historically included the `telnetlib` standard-library module, which provides a Telnet client implementation. A script can create a `telnetlib.Telnet` object, establish a TCP connection to a device's Telnet service, send commands, and read returned data. This made it a common built-in choice for simple network-device automation scripts, without requiring a separately installed third-party package. The module's official documentation describes it as providing a Telnet client implementation and documents its connection and read/write methods. See the [Python 3.8 telnetlib documentation](#) and the [Python 3.12 library reference](#). For modern development, it is important to recognize the version scope: `telnetlib` was deprecated in Python 3.11 and removed from the standard library in Python 3.13. Nevertheless, the statement accurately reflects the standard-library availability in the Python versions on which this certification question is based.

QUESTION NO: 45

The administrator configures R1 as shown in the figure. But R1 cannot learn the routes of other routers, so it may be.

The reason is? () (Multiple Choice Questions)

```
[R1]ospf
[R1-ospf-1]area 1
[R1-ospf-1-area-0.0.0.1]network 10.0.12.0 0.0.0.255
```

- A. This router did not announce the network connected to neighbors when configuring ospf
- B. This router is not configured with authentication function but the neighbor router is configured with authentication function
- C. When this router is configured, the ospf process number is not configured
- D. The area configured by this routerID and its neighbor router areaID is different

ANSWER: A B D

Explanation:

OSPF can install routes learned from other routers only after the routers have established an OSPF neighbor adjacency and exchanged link-state information. "This router did not announce the network connected to neighbors when configuring ospf" can prevent OSPF from being enabled on the relevant R1 interface. In that case, R1 does not send or receive OSPF Hello packets on the link and cannot form the required adjacency. "This router is not configured with authentication function but the neighbor router is configured with authentication function" is also a valid cause: OSPF neighbors must use compatible authentication settings and credentials for packets on the shared segment. An authentication mismatch prevents successful neighbor establishment. Likewise, "The area configured by this routerID and its neighbor router areaID is different" describes an OSPF area mismatch on the inter-router link. Hello packets carry the area identifier, and routers connected through the same OSPF link must agree on that area before becoming neighbors. These conditions therefore stop the adjacency and subsequent LSDB synchronization that supplies remote OSPF routes. The OSPF neighbor-state and Hello-packet requirements are defined in [RFC 2328](#); Huawei's OSPF configuration documentation is available through the [Huawei Enterprise Support documentation portal](#).

QUESTION NO: 46

When Huawei switches allow STP, the priority of the switch by default is ()

- A. 32768
- B. 8192
- C. 16384
- D. 4096

ANSWER: A

Explanation:

32768 is the default STP bridge priority on Huawei switches. In STP, the root bridge is elected according to the Bridge ID, which combines the bridge priority with the switch's MAC address. A lower Bridge ID is preferred, so configuring a lower priority is a standard way to intentionally make a particular switch the root bridge. Huawei uses 32768 as the default bridge-priority value when STP is enabled and no explicit priority has been configured. Bridge-priority values are represented in increments of 4096, allowing administrators to create a predictable hierarchy when selecting root and backup-root bridges. For example, a distribution-layer switch can be assigned a priority lower than 32768 so that it wins the root-bridge election over switches retaining their default setting. The displayed default is the base bridge-priority value; in VLAN-aware spanning-tree implementations, the effective identifier can also incorporate the system-ID extension associated with the VLAN or spanning-tree instance. Huawei's STP configuration guidance documents the bridge-priority setting and its role in root-bridge election: [Huawei STP bridge priority configuration](#). The general IEEE STP Bridge ID election principle is also summarized by [Cisco's STP root-bridge overview](#).

QUESTION NO: 47

What is the default priority of a VRRP router? () (Single-choice question)

- A. 0
- B. 1

C. 100

D. 255

ANSWER: C

Explanation:

The correct answer is **100**. In VRRP, each router in a VRRP group has a priority value used during Master router election. The router with the highest priority normally becomes the Master router and is responsible for forwarding traffic sent to the virtual IP address.

The default VRRP priority is 100. A router that owns the virtual IP address uses priority 255 and has special election behavior. VRRP priority and Master election are specified in [RFC 5798](#).

QUESTION NO: 48

The output information of the routing table of a certain router is shown in the figure. Which of the following statements is correct? () (Multiple Choice Questions)

<R1> display ip routing-table

Route Flags: R-relay, D-download to fib

Routing Tables: Public						
Destinations 10 Routes:10						
Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
10.0.0.0/8	Static	60	0	RD	10.0.12.2	Ethernet0/0/0
10.0.2.2/32	Static	60	0	RD	10.0.21.2	Ethernet0/0/1

A. The NextHop for this router to reach 10.0.2.2 is 10.0.12.2

B. The NextHop for this router to reach 10.0.0.1 is 10.0.12.2

C. The NextHop for this router to reach 10.0.2.2 is 10.0.21.2

D. The NextHop for this router to reach 10.0.0.1 is 10.0.21.2

ANSWER: B C

Explanation:

The statements “The NextHop for this router to reach 10.0.0.1 is 10.0.12.2” and “The NextHop for this router to reach 10.0.2.2 is 10.0.21.2” correctly apply routing-table lookup. A router forwards a packet by comparing the destination IP address with the destination prefixes in its routing table and selecting the best matching route. In the displayed table, 10.0.0.1 belongs to the 10.0.0.0 network entry, whose listed next hop is 10.0.12.2. Therefore, packets addressed to 10.0.0.1 are sent to 10.0.12.2 for the next forwarding step. Likewise, 10.0.2.2 belongs to the 10.0.2.0 network entry, and that entry specifies 10.0.21.2 as its next hop. The next-hop address is the adjacent router to which R1 sends the packet; it is not necessarily the packet’s final destination. The route flags shown by Huawei identify route status, including whether a route has been downloaded to the forwarding information base, while the destination prefix and NextHop fields provide the information required for this lookup. Huawei documents the routing-table display command and route attributes in its [VRP product documentation](#); the forwarding principle is also defined in [RFC 1812](#).

QUESTION NO: 49

Which of the following statements about SNMP are correct? ()

A. An SNMP NMS can send a Get-Request message to an Agent

- B. An SNMP Agent returns a Response message after processing a request
- C. SNMPv3 can provide authentication and encryption protection
- D. An SNMP Trap message is sent only after the NMS sends a request

ANSWER: A B C

Explanation:

SNMP uses an NMS and Agent model. The NMS sends management requests to an Agent running on a managed device, and the Agent returns a Response message after processing a request. SNMPv3 supports user-based security mechanisms that can provide authentication and privacy protection for management communication.

Community names are used by SNMPv1 and SNMPv2c. SNMP Trap messages are unsolicited notifications sent by an Agent to an NMS; they are not sent only after an NMS request. SNMP protocol operations are specified in [RFC 3416](#), and SNMPv3 security mechanisms are described in [RFC 3414](#).

QUESTION NO: 50

AC. Use CAPWAP control tunnel to transmit management messages to FIT AP. ()

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. In Huawei's centralized WLAN architecture, a Fit AP establishes a CAPWAP tunnel with its AC so that the AC can centrally manage the AP. The CAPWAP control tunnel carries control-plane and management communication, including AP discovery, joining, authentication, keepalive exchanges, configuration delivery, status reporting, and other AP management messages. This channel allows the AC to provision and supervise Fit APs while retaining centralized control over wireless service behavior. CAPWAP defines a control channel specifically for control and management functions between the access controller and wireless termination point; its separate data channel is intended for encapsulated client traffic when centralized forwarding is used. Therefore, management messages sent from an AC to a Fit AP use the CAPWAP control tunnel. See the CAPWAP protocol specification in [RFC 5415](#) and Huawei's [CAPWAP documentation search](#) for Huawei WLAN implementation material.

QUESTION NO: 51

The NAT configuration is as follows, which IP can be converted when accessing the public network (RTA) acl2000 (RTA-acl-basic-2000) rule5 permit source 192.168.1.0.0.0.255 RTA-acl-basic-2000) rule 10 deny source 192.168. 1.30

(RTA-acl-basic-2000) quit RTA) interface serial1/0/0 (RTA-Serial1/0/0) nat outbound. 2000*

- A. 192.168.1.1
- B. 192.168.2.1
- C. 192.168.1.2
- D. 192.168.1.3

ANSWER: A C D

Explanation:

The IP addresses **192.168.1.1**, **192.168.1.2**, and **192.168.1.3** can be translated when traffic leaves Serial1/0/0. The `nat outbound 2000` command associates outbound NAT processing with basic ACL 2000. An address is eligible for translation when its packet matches a permit rule in that ACL. The first rule permits the source range `192.168.1.0 0.0.0.255`, which represents the complete 192.168.1.0/24 network. Consequently, each of these three source addresses is within the permitted range and is selected for NAT processing. ACL rules are evaluated in their configured order, and matching stops at the first applicable rule. Therefore, traffic sourced from 192.168.1.3 matches the preceding subnet permit rule before the later deny rule can be considered. This configuration illustrates why ACL entries used by NAT must be ordered from the most specific matching criteria to more general criteria. Huawei documents outbound NAT as using an ACL to identify internal source addresses that require address translation; see the [Huawei Enterprise documentation portal](#) and the [Huawei Enterprise Support site](#).

QUESTION NO: 52

Which of the following IPv4 addresses the host uses cannot directly access the Internet? ()

- A. 100.1.1.1
- B. 50.1.1.1
- C. 10.1.1.1
- D. 200.1.1.1

ANSWER: C

Explanation:

10.1.1.1 is within the IPv4 private-use block 10.0.0.0/8. RFC 1918 reserves this entire range for use inside private networks, such as enterprise LANs, home networks, and internal cloud environments. These addresses are not globally unique and are not intended to be advertised or routed across the public Internet. Internet service providers and public Internet routers should filter traffic whose source or destination is an RFC 1918 private address, preventing a host using 10.1.1.1 from directly communicating with Internet hosts using that private address as its public identity.

To reach the Internet, a device configured with 10.1.1.1 normally sends traffic to a gateway that performs Network Address Translation. NAT replaces the private source address with a routable public address, maintains a translation entry for the session, and maps returning traffic back to the internal host. This preserves private addressing internally while enabling controlled Internet access. The private-address designation and its intended non-public routing behavior are defined in [RFC 1918](#); the IANA IPv4 special-purpose registry also identifies 10.0.0.0/8 as private-use space: [IANA IPv4 Special-Purpose Address Registry](#).

QUESTION NO: 53

Can the administrator modify the device name for the router under ()?

- A. User-view
- B. Protocol-view
- C. System-view
- D. Interface-view

ANSWER: C

Explanation:

The correct answer is **System-view**. On Huawei VRP-based routers, changing the device name (also called the host name) is a global configuration task. An administrator first enters the system view from the user view by using the `system-view` command, then configures the name with the `sysname` command. For example, `sysname R1` changes the device prompt to reflect the new name, helping administrators identify the router clearly during local-console or remote-management sessions.

The system view is the central configuration context for device-wide settings, including management and foundational system parameters. The hostname is not tied to one physical or logical interface, nor is it configured within the context of a particular routing protocol; it applies to the entire router. Huawei documentation describes entering system view before executing global configuration commands and identifies `sysname` as the command used to configure a device name. See Huawei's [sysname command reference](#) and the [system-view command reference](#).

QUESTION NO: 54

In the STP protocol, which of the following factors will affect the election of the root switch? (*)

- A. MAC address of the switch
- B. IP address of the switch
- C. Switch priority
- D. Switch interface bandwidth
- E. Switch interface ID

ANSWER: A C

Explanation:

STP elects the root switch by comparing each participating switch's Bridge ID (BID). The BID is designed to provide a deterministic, network-wide ordering: the switch with the numerically lowest BID becomes the root bridge. The Bridge ID includes the configurable switch priority and the switch's MAC address. Therefore, changing the switch priority directly affects the root-bridge election and is the normal administrative method for intentionally selecting a root switch. If multiple switches have the same priority, STP uses the MAC-address portion of the Bridge ID as the tie-breaker; the lower MAC address wins. In VLAN-aware variants such as RSTP/MSTP implementations, the extended system ID can also be present in the priority portion of the BID, but the fundamental comparison still depends on priority followed by MAC address. This mechanism lets administrators establish a predictable Layer 2 topology by assigning a lower bridge priority to the switch intended to act as the root. Huawei's STP configuration documentation describes bridge priority as a root-bridge election parameter, and the IEEE 802.1Q standard defines the bridge identifier structure used by spanning-tree protocols. See [Huawei STP configuration documentation](#) and [IEEE 802.1Q](#).

QUESTION NO: 55

On Huawei equipment, if AAA authentication is used for authorization, when the remote server does not respond, authorization can be performed from the network equipment side ()

- A. True
- B. False

ANSWER: A


Explanation:

True is correct. Huawei VRP AAA supports authorization schemes that can use a remote AAA server, such as a RADIUS or HWTACACS server, together with local authorization on the device. When the authorization scheme is configured with a remote method followed by the local method, the device first attempts to obtain authorization information from the remote server. If that server does not respond, the device can use its locally configured authorization data as the fallback processing method. This preserves managed access when the external AAA service is temporarily unavailable, provided that suitable local users, privilege levels, and authorization settings have been configured in advance.

This behavior is controlled by the authorization method sequence in the AAA scheme; it is not an automatic replacement for careful AAA design. Administrators should deliberately configure the local fallback method and ensure that local accounts have appropriately restricted permissions, because fallback authorization is intended to maintain availability while retaining policy control. Huawei's enterprise documentation portal provides the VRP AAA configuration and command references:

QUESTION NO: 56

As shown in the figure below, assuming that the other configurations of the switch are kept in the default state, which port of the following switch will become the designated port? () (Multiple Choice Questions)



- A. G0/0/1 of SWB
- B. SWC's G0/0/2
- C. SWA's G0/0/3
- D. SWC's G0/0/1

ANSWER: A C

Explanation:

In the displayed STP topology, **G0/0/1 of SWB** and **SWA's G0/0/3** are designated ports. STP elects one designated port for each LAN segment: it is the port that advertises the superior BPDU toward that segment. BPDU superiority is evaluated first by root bridge ID, then by the accumulated root-path cost, followed by the sending bridge ID and sending port ID if a tie remains. A port on the root bridge is designated for every segment to which that bridge is connected, so SWA's G0/0/3 is designated on its attached segment. On the inter-switch segment where SWB competes with SWC, SWB's G0/0/1 advertises the better path toward the root bridge under the default STP calculation shown in the diagram. It therefore becomes the designated port for that segment and forwards BPDUs downstream. This designated-port election creates a loop-free forwarding tree while retaining an alternate path that STP can use after a topology change. Huawei VRP STP follows the standard root-bridge, root-port, and designated-port election process described in Huawei's [STP configuration documentation](#) and the IEEE STP specification in [IEEE 802.1D](#).

QUESTION NO: 57

The output information of the sub-interface of a certain router's aggregation port 1 is shown in the figure. Based on this information, which of the following statements is correct? ()

```
[Router A] display Interface Eth-Trunk 1.100Eth-Trunk1.100 current state UP Line protocol current state UP Last line
protocol up time 2019-03-04 10:22:40 UTC. -08:00 Description: HUAWEI, ARSeries, Eth-Trunk1.100 Interface Route Port,
Hash arithmetic. :According to SIP-XOR-DIP, Maximal BW:2G. Current BW:2G, The Maximum Transmit Unit is 1500
Internet Address is 10.0.12.2/24 IP Sending Frames'Format is PKTFIT ETENT 2, Hardwareaddress 1s 00e0-fc3b-2015
Current system time:2019-03-04 10:24:29-08:00 Last 300 seconds input rate0 bits/sec.0 packets/sec Last 300 seconds
output rate 0 bits /seC. 0 packets/sec Realtime 0 seconds inputrate 0 bits/seC. 0 packets/sec Realtime 0 seconds output rate
0 bits/seC. 0 packets/sec Input: 6 packets, 574 bytes,
```

Output: 7 packets, 638 bytes, Input bandwidth utilization 0% Output bandwidth utilization 0%

PortName Status Weight

GigabitEthernet0/0/1 UP 1 GigabitEthernet0/0/2 UP 1

The Number of Ports in Trunk:2 The Number of UP Ports in Trunk:2*

- A. Two links exist on the aggregated port
- B. The IP address of the sub-interface of the aggregate port is 10.0.12.2/24
- C. The sub-interface number of the aggregate port is 100
- D. The sub-interface of the aggregate port carries VLAN TAG100 when forwarding data frames

ANSWER: A B C

Explanation:

The statements “Two links exist on the aggregated port,” “The IP address of the sub-interface of the aggregate port is 10.0.12.2/24,” and “The sub-interface number of the aggregate port is 100” are supported directly by the displayed interface status. The member-port table lists GigabitEthernet0/0/1 and GigabitEthernet0/0/2, while the summary explicitly reports that the trunk contains two ports and that both are up. This also corresponds with the reported current bandwidth of 2 Gbit/s, consistent with two active GigabitEthernet member links. The interface identifier is Eth-Trunk1.100: Eth-Trunk1 identifies aggregation group 1, and the suffix following the period identifies sub-interface 100. In addition, the line “Internet Address is 10.0.12.2/24” explicitly gives the IPv4 address and prefix length configured on this routed sub-interface. Huawei Eth-Trunk interfaces can aggregate multiple physical member interfaces for load sharing and resiliency, and their operational display presents both trunk-member state and Layer 3 interface attributes. See the [Huawei documentation portal](#) for VRP interface and Eth-Trunk configuration references, and Huawei’s [Enterprise Support site](#) for product-specific command documentation.

QUESTION NO: 58

VLAN4095 cannot be created on Huawei switches, and VLAN1 cannot be deleted. ()

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. Huawei VRP-based switches use the standard IEEE 802.1Q VLAN identifier field, which is 12 bits long. Although that field represents values from 0 through 4095, the operational VLAN range available for ordinary VLAN creation is VLAN 1 through VLAN 4094. VLAN ID 4095 is reserved and therefore cannot be created as a user VLAN. This reserved-value treatment ensures compatibility with the VLAN-tagging standard and with switch forwarding behavior.

VLAN 1 is automatically present as the default VLAN when a Huawei switch starts. It provides the initial Layer 2 VLAN context for switch ports unless another VLAN is configured. Because it is the built-in default VLAN required by the system, VRP does not permit its deletion. Administrators can create and manage other VLANs in the supported range and assign interfaces to them, but VLAN 1 remains available as the default system VLAN. These two facts make the complete statement accurate. Huawei’s product documentation is available through the [Huawei Enterprise Documentation portal](#); the relevant tag format is defined by the [IEEE 802.1Q standard](#).

QUESTION NO: 59

As shown in the figure, after the two routers are configured with OSPF, the administrator configures the command on RTA. Is the following description correct? () (Multiple Choice Questions)



- A. RTA will continue to receive, analyze and process OSPF packets sent by RTB
- B. The neighbor relationship between the two routers will be down
- C. The neighbor relationship between the two routers will not be affected
- D. RTA will no longer send OSPF packets

ANSWER: A B D

Explanation:

RTA will continue to receive, analyze and process OSPF packets sent by RTB, the neighbor relationship between the two routers will be down, and RTA will no longer send OSPF packets are correct. In Huawei VRP, configuring an OSPF interface as a silent interface suppresses OSPF packet transmission on that interface while retaining the interface in the OSPF process. The router can therefore still receive and process OSPF packets arriving from the peer, including Hello packets. However, it does not transmit its own OSPF packets, most importantly Hello packets used to maintain bidirectional neighbor communication. RTB consequently stops receiving Hellos from RTA and, after the OSPF dead interval expires, declares the neighbor unavailable. The adjacency is then lost. This behavior is commonly used on interfaces connected to hosts or stub segments: the connected subnet can still be advertised through OSPF, but no OSPF neighbor is formed across that interface. Huawei documents the command in its [VRP product documentation](#); the role of Hello packets and the neighbor dead interval is defined in [RFC 2328](#).

QUESTION NO: 60

Telnetlib is a module that implements the Telnet protocol that comes with Python ()

- A. True
- B. False


ANSWER: A

Explanation:

True is correct in the context of Python versions current for the HCIA-Datacom exam material. `telnetlib` was included in Python's standard library and provides a Telnet client implementation. It supplies the `Telnet` class, which can establish a Telnet session with a network device, send commands, and read returned data. This made it useful in basic network automation scripts, including scripts that log in to a router or switch and collect command output. Because it was part of the standard library, it could traditionally be imported directly with `import telnetlib`, without installing a separate package through pip. Python's historical library documentation describes it as a "Telnet client" module and documents the `telnetlib.Telnet` interface: [Python 3.10 telnetlib documentation](#). A modern caveat is that Python deprecated `telnetlib` in version 3.11 and removed it in Python 3.13; however, that later lifecycle change does not alter the intended answer for this question, which refers to the module as distributed with the Python versions used in the course context. The removal is recorded in [Python's removed modules documentation](#).

QUESTION NO: 61

In the network shown in the figure below, what commands can be entered on the router to make host A ping host B through? ()



- A. RTA:ip route-static 0.0.0.0 0.0.0.0 11.0.12.2
RTB:ip route-static 0.0.0.0 0.0.0.0 11.0.12.1
- B. RTA:ip route-static 10.0.12.5 255.255.255.252 11.0.12.2
RTB:ip route-static 10.0.12.1 255.255.255.252 11.0.12.1
- C. RTA:ip route-static 0.0.0.0 0.0.0.0 11.0.12.1
RTB:ip route-static 0.0.0.0 0.0.0.0 11.0.12.2
- D. RTA:ip route-static 10.0.12.5 255.255.255.252 11.0.12.1
RTB:ip route-static 10.0.12.1 255.255.255.252 11.0.12.2

ANSWER: A B

Explanation:

The two valid configurations provide bidirectional routing between the LANs connected to the two routers. With RTA:ip route-static 0.0.0.0 0.0.0.0 11.0.12.2 and RTB:ip route-static 0.0.0.0 0.0.0.0 11.0.12.1, each router installs a default route whose next hop is the directly connected address of the peer router on the 11.0.12.0 link. Therefore, traffic for an unknown remote destination is forwarded across that link in either direction, including the ICMP request and reply required for a successful ping.

The configuration using routes for 10.0.12.5 and 10.0.12.1 with a 255.255.255.252 mask also works. A /30 mask identifies the respective remote /30 LAN prefixes: the address ending in .5 belongs to the 10.0.12.4/30 subnet, and the address ending in .1 belongs to the 10.0.12.0/30 subnet. Each route points to the peer address on the transit network, so both routers have an explicit path toward the opposite host network. Static routing requires a valid, reachable next hop and return-path routing; both correct configurations meet those requirements. See Huawei's [routing documentation portal](#) and the forwarding requirements described in [RFC 1812](#).

QUESTION NO: 62

Which of the following features does the OSPF protocol have? (*)

- A. Support area division
- B. Prone to routing loops
- C. Calculate the shortest path based on the number of hops
- D. Trigger update

ANSWER: A D

Explanation:

OSPF supports hierarchical area division. An OSPF autonomous system is organized into areas to limit the scope of link-state advertisements and SPF recalculations, improving scalability in larger networks. All non-backbone areas connect logically to the backbone area, and area border routers maintain the required inter-area routing information. This hierarchical design is a defining OSPF capability.

OSPF also uses triggered updates. Rather than relying on periodic full routing-table broadcasts, it originates and floods link-state advertisements when a topology change occurs, such as an interface or neighbor state change. Routers receiving the updated information refresh their link-state databases and run the SPF calculation as needed to converge on the new topology. OSPF additionally refreshes LSAs periodically for database aging and reliability, but topology-change-driven LSA flooding is the relevant triggered-update behavior. RFC 2328 describes OSPF areas and the flooding of changed LSAs; see

[OSPF Version 2 \(RFC 2328\)](#), especially its descriptions of areas and link-state advertisement flooding. Huawei also summarizes OSPF's area-based hierarchical structure in its [OSPF documentation](#).

QUESTION NO: 63

In an OSPF broadcast network, with which routers does a Dothter router exchange link state information? ()

- A. BDR
- B. All OSPF neighbors
- C. DR Other
- D. DR

ANSWER: A D

Explanation:

On an OSPF broadcast multi-access segment, a DR Other (DROther) router maintains full adjacencies with the designated router and backup designated router. Therefore, it exchanges database-description packets, link-state requests, link-state updates, and link-state acknowledgments with the **DR** and **BDR**. These full adjacencies allow the DROther router to receive the synchronized link-state database and reliably flood topology changes through the segment's elected central routers.

This design is fundamental to OSPF scalability on Ethernet-style shared networks. Rather than requiring every router on the segment to form a full adjacency with every other router, the DR and BDR act as the primary exchange points for link-state information. The BDR is fully adjacent as well so that it already has an up-to-date database and can take over quickly if the DR fails. OSPF routers still discover peers and exchange Hello packets on the shared network, but the full link-state-information exchange relevant to this question occurs with the DR and BDR. This behavior is specified for broadcast networks in [RFC 2328, Section 7.5](#) and described in the [OSPF neighbor state machine specification](#).

QUESTION NO: 64

By default, the bridge priority of Huawei switches is 32768 ()

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. In Huawei VRP-based switching devices, the default STP bridge priority is 32768, which is the midpoint of the valid 16-bit bridge-priority range. STP uses the bridge ID to elect the root bridge for a spanning-tree instance; this identifier incorporates the configured bridge-priority value, the system-ID extension where applicable, and the switch MAC address. A lower bridge ID is preferred during root-bridge election. Therefore, an administrator who wants a particular switch to become root commonly configures a priority lower than the default 32768, using the supported priority increments for the relevant STP mode and device software version. The default value provides a neutral starting point so that, without explicit configuration, the eventual election can be determined by the remaining bridge-ID components when priorities are equal. This behavior follows the standard bridge-priority mechanism used by STP and is reflected in Huawei switch STP configuration and operational documentation. See Huawei's [Enterprise Support Documentation portal](#) for platform-specific VRP command references and the IEEE [802.1D standard page](#) for the spanning-tree bridge-ID framework.

QUESTION NO: 65

The output information of a certain router is shown in the figure. Which of the following statements are correct? () (Multiple Choice Questions)

```
<R1>dis ospf peer
OSPF Process 1 with Router ID 10.0.2.2
Neighbors
Area 0.0.0.0 interface 10.0.12.2 ( GigabitEthernet0/0/0 ) 's neighborsRouter
ID:10.0.1.1Address:10.0.12.1
State:Full Mode:Nbr is Slave Priority:1
DR:10.0.12.2BDR:10.0.12.1MTU:0
Dead timer due in 32 sec
Retrans timer interval:5
Neighbor is up for 00:03:30
Authentication Sequence:0]
```

- A. The Router ID of the router is 10.0.1.1
- B. This router is DR
- C. The interface address of this router is 10.0.12.2
- D. The Router ID of the router is 10.0.2.2.

ANSWER: B C D

Explanation:

The correct statements are supported directly by the OSPF neighbor output. The local OSPF process identifies this device with router ID **10.0.2.2**. A router ID is a 32-bit value that uniquely identifies an OSPF router within the routing domain; it is not necessarily the address of the interface on which the neighbor relationship is displayed. The output also identifies the local OSPF-enabled interface using address **10.0.12.2**, so that is the interface address of the router represented by the command output. Finally, the designated-router field shows **10.0.12.2**. Because this matches the local interface address, the local router is the DR for that multi-access OSPF network. OSPF elects a DR on broadcast and NBMA segments to reduce adjacency and LSA-flooding overhead; routers report the elected DR by its interface address on that segment. These fields—local Router ID, local interface address, and DR address—therefore establish that the router ID is 10.0.2.2, the interface address is 10.0.12.2, and the router is DR. See the OSPF neighbor and DR behavior defined in [RFC 2328](#) and Huawei's [Enterprise Support documentation portal](#).

QUESTION NO: 66

On a Huawei router, the administrator configures the following static route. What is the meaning of the parameter 60? () (Single-choice question)

```
[RTA] ip route-static 192.168.10.0 255.255.255.0 10.0.12.2 preference 60
```

- A. The route preference
- B. The cost of the outbound interface
- C. The metric of the next hop
- D. The route aging time

ANSWER: A

Explanation:

The value **60** is the preference of the static route. On Huawei VRP devices, route preference is used to select between routes to the same destination that are learned from different sources. A smaller preference value indicates a higher priority.

The configured route therefore has a preference of 60 and can be selected ahead of another route to the same destination with a larger preference value, provided that other route-selection conditions are comparable.

The value is not an interface cost, a metric carried by the next hop, or an administrative distance configured on the neighboring router. It is a local route-selection attribute of the static route. Huawei documentation for routing configuration is available from the [Huawei Enterprise Support portal](#).

QUESTION NO: 67

Which authorization modes are supported by AAA on Huawei ARG3 series routers? (*)

- A. Not authorized
- B. HWTACACS authorization
- C. Local authorization
- D. Authorization after successful RADIUS authentication


ANSWER: A B C D

Explanation:

Huawei VRP AAA authorization schemes support the full set represented here. **Not authorized** corresponds to the `none` authorization mode, in which no authorization check is performed. **Local authorization** uses locally configured user and authorization information on the router, allowing access rights to be determined without an external authorization server. **HWTACACS authorization** uses an HWTACACS server to return authorization information, such as service permissions and command privileges, after the user is authenticated. **Authorization after successful RADIUS authentication** expresses the `if-authenticated` authorization behavior: once authentication has succeeded, authorization is permitted without requiring a separate HWTACACS authorization exchange. These modes allow an administrator to select authorization behavior according to whether credentials and permissions are maintained locally or through centralized AAA services. Huawei's VRP documentation describes AAA configuration and authorization schemes in its product documentation and command references. See the [Huawei Enterprise documentation portal](#) and [Huawei Enterprise Support](#) for the applicable router software documentation.

QUESTION NO: 68

The output information of the router is shown in the figure. Which of the following statements is wrong? () (Single-choice question)



```
<R1>display ip interface Ethernet 0/0/0
Ethernet0/0/0 current state : UP
Line protocol current state : UP
The Maximum Transait Unit : 1500 bytes
input packets : 3645, bytes : 287468, multicasts : 3590
output packets : 3715, bytes : 294616, multicasts : 3654
Directed-broadcast packets:
  received packets:      0, sent packets:      0
  forwarded packets:    0, dropped packets:    0
ARP packet input number: 29
  Request packet:       13
  Reply packet:         16
  Unknown packet:       0
Internet Address is 10.0.12.1/24
Broadcast address : 10.0.12.255
```

The broadcast address corresponding to the AEthernet0/0/0 interface is 10.0.12.255

- B. The physical link of the Ethernet0/0/0 interface is normal
- C. Ethernet0/0/0 interfaceThe MTU value is 1480

D. Ethernet (The IP address of the /0/0 interface is 10.0.12.1/24)

Answer: C

Explanation:

- A. The physical link of the Ethernet0/0/0 interface is normal
- B. The MTU value of the Ethernet0/0/0 interface is 1480.
- C. The IP address of the Ethernet0/0/0 interface is 10.0.12.1/24.

ANSWER: B

Explanation:

“The MTU value of the Ethernet0/0/0 interface is 1480” is the wrong statement because the interface information shown in the output reports an MTU of **1500 bytes**, not 1480 bytes. MTU is the largest Layer 3 packet size, in bytes, that an interface can transmit without IP fragmentation. On an Ethernet interface using the normal Ethernet frame format, 1500 bytes is the standard IP MTU value and is the value displayed by the router in this case. A value of 1480 would be a distinct configured or negotiated interface characteristic and cannot be inferred merely from a standard Ethernet connection or from the IPv4 address.

When interpreting Huawei interface-display output, the MTU must be read from the explicit MTU field rather than calculated from the address prefix, broadcast address, or physical-link status. This is important for diagnosing packet fragmentation, Path MTU Discovery behavior, and tunnel overhead issues. Huawei VRP interface commands present operational attributes such as interface state, protocol state, addressing, and MTU together; the MTU field is authoritative for the interface’s current packet-size limit. See Huawei’s [VRP product documentation](#) and the IETF’s definition of IP MTU behavior in [RFC 1191](#).

QUESTION NO: 69

Telnet.read in Telnetlib. _The role of all() is to read all data until EOF. If the echo does not return EOF, it will always block ()

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. In Python’s historical `telnetlib` API, the method intended by the question is `Telnet.read_all()`. Its documented behavior is to read data from the Telnet connection until an EOF condition is received. EOF occurs when the remote side closes the connection, so the method cannot complete normally while the connection remains open and no EOF is delivered. Consequently, a call to `read_all()` blocks waiting for additional input or for the remote endpoint to close the Telnet session. This is important in network automation: when communicating with devices that maintain an interactive Telnet CLI session, scripts generally use methods that read until an expected prompt or matching text rather than `read_all()`, unless the device is expected to close the connection. Python’s API documentation explicitly states that `read_all()` “Read[s] all data until EOF; block[s] until connection closed.” The question’s wording uses “echo,” but the relevant technical condition is the arrival of EOF from the Telnet connection, not command echo behavior itself. See the [Python 3.12 Telnet.read_all\(\) documentation](#) and the [telnetlib module reference](#).

QUESTION NO: 70

The controller is the core component of SDN. The controller is connected to the device through the southbound interface. The following is missing from the controller southbound agreement? ()

- A. SNMP

- B. PCEP
- C. OpenFlow
- D. NETCONF

ANSWER: A B C D

Explanation:

Controller southbound communication can use several protocols, depending on whether the controller needs to configure devices, collect device information, install forwarding behavior, or perform path computation and traffic engineering. **SNMP** is used by controllers and network-management systems to retrieve operational and monitoring data from managed devices. **NETCONF** provides a standardized, structured mechanism for installing, modifying, and retrieving device configuration and state data; it is widely used for model-driven network management. **OpenFlow** is a foundational SDN southbound protocol that enables a controller to program forwarding behavior in an OpenFlow-capable data plane. **PCEP** allows a Path Computation Element, including a controller-based PCE, to communicate with Path Computation Clients for path computation, delegation, and traffic-engineering control. Therefore, each listed protocol can be used as part of a controller's southbound communication capability; practical deployments select one or more based on the device type and required control function. The protocol specifications describe these roles in detail: [IETF RFC 6241—NETCONF](#) and [IETF RFC 5440—PCEP](#).

QUESTION NO: 71

Which of the following types does IPv6 address include? (*)

- A. Unicast address
- B. Broadcast address
- C. Multicast address
- D. Anycast address

ANSWER: A C D

Explanation:

IPv6 defines three principal address types: unicast, multicast, and anycast. A unicast address identifies a single network interface, so packets sent to it are delivered to that one interface. IPv6 unicast addressing includes categories such as global unicast, unique local, and link-local addresses. A multicast address identifies a group of interfaces; traffic sent to the multicast address is delivered to all interfaces that have joined the corresponding group. IPv6 relies on multicast for functions such as Neighbor Discovery, including solicitation and router discovery operations. An anycast address can be assigned to multiple interfaces, typically on different devices, while using the same address. A packet sent to an anycast address is routed to one member of that group, normally the topologically nearest one according to routing. These address models support scalable communication and service availability without requiring an IPv6 broadcast address type. Huawei IPv6 configuration and operation therefore use unicast, multicast, and anycast concepts. See the IPv6 addressing architecture in [RFC 4291](#) and Huawei's [IPv6 Addressing Overview](#).

QUESTION NO: 72

The following description of static and dynamic routing is wrong ().

- A. Dynamic routing protocols take up more system resources than static routing protocols.
- B. After the link fails, the static route can automatically complete the network convergence.
- C. After the administrator deploys the dynamic routing protocol in the corporate network, later maintenance and expansion can be more convenient
- D. Static routing is easy to configure and easy to manage when applied in the enterprise

ANSWER: B

Explanation:

“After the link fails, the static route can automatically complete the network convergence.” is the incorrect description. A conventional static route is manually configured with a fixed next hop or outbound interface. It does not exchange routing information, calculate alternate paths, or independently select a replacement path when a failure occurs. Therefore, a link failure will not by itself cause automatic convergence to another route merely because a static route exists. The administrator must intervene, or the network design must include additional mechanisms in advance. On Huawei VRP devices, this can include configuring multiple static routes with different preferences, associating route validity with Network Quality Analysis (NQA) detection, or using Bidirectional Forwarding Detection (BFD) for faster fault detection. Those mechanisms can provide static-route backup or switchover behavior, but they are explicit enhancements rather than inherent automatic convergence of a basic static route. This distinction is important in enterprise design: static routes are deterministic and simple, while resilient failover requires deliberate redundancy and reachability tracking. Huawei’s routing documentation describes static-route configuration and route preference behavior, while the IPv4 router requirements standard explains the forwarding and routing responsibilities of routers. See [Huawei Enterprise Support documentation](#) and [RFC 1812: Requirements for IP Version 4 Routers](#).