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

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

Leaf and spine data centers are used to better accommodate which type of traffic?

A. north-east B. east-west
C. north-west
D. south-east

Answer: B

Explanation:

In modern data centers, the shift toward leaf-spine architectures is driven by the need to handle increased east-west traffic, which is traffic between servers within the same data center. Unlike traditional hierarchical data center designs, where most traffic was "north-south" (between users and servers), modern applications often involve server-to-server communication (east-west) to enable services like distributed databases, microservices, and virtualized workloads.
Leaf-Spine Architecture:
Leaf Layer: This layer consists of switches that connect directly to servers or end-host devices. These switches serve as the access layer.
Spine Layer: The spine layer comprises high-performance switches that provide interconnectivity between leaf switches. Each leaf switch connects to every spine switch, creating a non-blocking fabric that optimizes traffic flow within the data center.
East-West Traffic Accommodation:
In traditional three-tier architectures (core, aggregation, access), traffic had to traverse multiple layers, leading to bottlenecks when servers communicated with each other. Leaf-spine architectures address this by creating multiple equal-cost paths between leaf switches and the spine. Since each leaf switch connects directly to every spine switch, the architecture facilitates quick, low-latency communication between servers, which is essential for east-west traffic flows.
Juniper's Role:
Juniper Networks provides a range of solutions that optimize for east-west traffic in a leaf-spine architecture, notably through:
QFX Series Switches: Juniper’s QFX series switches are designed for the leaf and spine architecture, delivering high throughput, low latency, and scalability to accommodate the traffic demands of modern data centers.
EVPN-VXLAN: Juniper uses EVPN-VXLAN to create a scalable Layer 2 and Layer 3 overlay network across the data center. This overlay helps enhance east-west traffic performance by enabling network segmentation and workload mobility across the entire fabric. Key Features That Support East-West Traffic:
Equal-Cost Multipath (ECMP): ECMP enables the use of multiple paths between leaf and spine switches, balancing the traffic and preventing any one path from becoming a bottleneck. This is crucial in handling the high volume of east-west traffic.

Low Latency: Spine switches are typically high-performance devices that minimize the delay between leaf switches, which improves the efficiency of server-to-server communications.
Scalability: As the demand for east-west traffic grows, adding more leaf and spine switches is straightforward, maintaining consistent performance without redesigning the entire network.
In summary, the leaf-spine architecture is primarily designed to handle the increase in east-west traffic within data centers, and Juniper provides robust solutions to enable this architecture through its switch platforms and software solutions like EVPN-VXLAN.

Question: 2

When troubleshooting an OSPF neighborship, you notice that the router stopped at the ExStart state. What is the cause of this result?

A. The priority is set to 255.
B. There is an interval timing mismatch.
C. There is an area ID mismatch.
D. There is an MTU mismatch.

Answer: D

Explanation:

When an OSPF (Open Shortest Path First) neighborship is stuck in the ExStart state, it usually points to a mismatch in Maximum Transmission Unit (MTU) settings between two routers trying to establish the adjacency. The ExStart state is where OSPF routers negotiate the master-slave relationship and exchange DBD (Database Description) packets.

Step-by-Step Breakdown:

OSPF Neighbor States: OSPF goes through several states to establish an adjacency with a neighbor: Down: No hello packets have been received.

Init: Hello packets are received, but bidirectional communication isn't confirmed. 2-Way: Bidirectional communication is established.

ExStart: The routers are negotiating who will be the master and who will be the slave, and begin to exchange DBD packets.

Exchange: The routers start exchanging the database information. Loading: The routers process the Link-State Advertisements (LSAs). Full: The adjacency is fully established.

MTU Mismatch Issue:

During the ExStart state, both OSPF routers must agree on their MTU values. If there is an MTU mismatch between the two routers, OSPF neighbors will fail to move from the ExStart to the Exchange state. The router with the larger MTU setting will not accept DBD packets from the router with a smaller MTU because the packets may exceed the smaller MTU size.

In Juniper devices, this behavior can be identified by examining the MTU settings using the show interfaces command and ensuring both routers have matching MTU configurations. To resolve this issue, either match the MTU settings on both routers or configure OSPF to ignore MTU mismatches using the command set protocols ospf ignore-mtu.

Juniper Reference:

Junos Command: show ospf neighbor helps diagnose neighbor states.

MTU Adjustment: set interfaces <interface-name> mtu <size> can be used to set the MTU values correctly.

Question: 3

Which statement is correct about aggregate routes?

A. Thedefaultnext hopis discard.
B. Thedefaultnext hopis readvertise.
C. Thedefaultnext hopis resolve.
D. Thedefaultnext hopis reject.

Answer: D

Explanation:

An aggregate route is a summarized route that is created by combining multiple specific routes into a single, broader route. In Junos OS, when an aggregate route is configured, its default next hop is set to reject.

Step-by-Step

Aggregate Route:

Aggregate routes are used to reduce the size of routing tables by representing a collection of more specific routes with a single summary route. They help improve routing efficiency and scalability, especially in large networks.

Default Next Hop Behavior:

When you configure an aggregate route in Junos OS, it has a reject next hop by default.

The reject next hop means that if a packet matches the aggregate route but there is no more specific route in the routing table for that destination, the packet will be discarded, and an ICMP "destination unreachable" message is sent to the source.

This behavior helps to prevent routing loops and ensures that traffic isn't forwarded to destinations for which there is no valid route.

Modifying Next Hop:

If needed, the next hop behavior of an aggregate route can be changed to discard (which silently drops the packet) or to another specific next hop. However, by default, the next hop is set to reject.

Juniper Reference:

Junos Command: set routing-options aggregate route <route> reject to configure an aggregate route with a reject next hop.

Verification: Use show route to verify the presence and behavior of aggregate routes.

Question: 4

Which Junos OS routing table stores IPv6 addresses?

A. inet.0
B. inet0.6
C. inet.6
D. inet6.0

Answer: D

Explanation:

In Junos OS, routing information is stored in different routing tables depending on the protocol and address family. For IPv6 addresses, the routing table used is inet6.0.

Step-by-Step

Routing Tables in Junos:

inet.0: This is the primary routing table for IPv4 unicast routes. inet6.0: This is the primary routing table for IPv6 unicast routes. inet.3: This routing table is used for MPLS-related routing.

Other routing tables, like inet.1, inet.2, are used for multicast and other specific purposes. inet6.0 Routing Table:

When IPv6 is enabled on a Juniper router, all the IPv6 routes are stored in the inet6.0 table. This includes both direct routes (connected networks) and learned routes (from dynamic routing protocols like OSPFv3, BGP, etc.).

Verification:

To view IPv6 routes, the command show route table inet6.0 is used. This will display the contents of the IPv6 routing table, showing the network prefixes, next-hop addresses, and protocol information for each route.
Juniper Reference:
Junos Command: Use show route table inet6.0 to check IPv6 routing entries.
IPv6 Routing: Ensure that the IPv6 protocol is enabled on interfaces and that routing protocols like OSPFv3 or BGP are properly configured for IPv6 traffic handling.

Question: 5

What is the primary purpose of an IRB Layer 3 interface?

A. to provide load balancing
B. to provide a default VLAN ID
C. to provide inter-VLAN routing
D. to provide port security

Answer: C
Explanation:

The primary purpose of an IRB (Integrated Routing and Bridging) interface is to enable inter-VLAN routing in a Layer 3 environment. An IRB interface in Junos combines the functionality of both Layer 2 bridging (switching) and Layer 3 routing, allowing devices in different VLANs to communicate with each other.
Step-by-Step Breakdown:
VLANs and Layer 2 Switching:
Devices within the same VLAN can communicate directly through Layer 2 switching. However, communication between devices in different VLANs requires Layer 3 routing.

IRB Interface for Inter-VLAN Routing:
The IRB interface provides a Layer 3 gateway for each VLAN, enabling routing between VLANs. Without an IRB interface, devices in different VLANs would not be able to communicate.
Configuration:
In Juniper devices, the IRB interface is configured by assigning Layer 3 IP addresses to it. These IP addresses serve as the default gateway for devices in different VLANs.
Example configuration:

set interfaces irb unit 0 family inet address 192.168.1.1/24 set vlans vlan-10 l3-interface irb.0
This allows VLAN 10 to use the IRB interface for routing.
Juniper Reference:
IRB Use Case: Inter-VLAN routing is essential in data centers where multiple VLANs are deployed, and Juniper's EX and QFX series switches support IRB configurations for this purpose.

Question: 6

Which two statements describe an IP fabric? (Choose two.)

A. An IP fabric allows devices to always be one hop away.
B. An IP fabric depends on Layer 2 switching.

C. An IP fabric uses spine and leaf devices.
D. An IP fabric provides traffic load sharing.

Answer: C, D

Explanation:

An IP fabric is a network topology designed to provide a scalable, low-latency architecture that is typically implemented in modern data centers. It uses spine and leaf switches and enables efficient traffic load sharing across the network.

Step-by-Step Breakdown:

Spine-Leaf Architecture:

Leaf Devices: These switches connect to servers and edge devices within the data center. Each leaf switch connects to every spine switch.

Spine Devices: These high-performance switches interconnect all the leaf switches. There are no direct connections between leaf switches or spine switches. This architecture ensures that any two endpoints within the fabric are only one hop away from each other, minimizing latency.

Traffic Load Sharing:

An IP fabric leverages Equal-Cost Multipath (ECMP) to distribute traffic evenly across all available paths between leaf and spine switches, providing effective load balancing. This ensures that no single link becomes a bottleneck and that traffic is distributed efficiently across the network.

Juniper Reference:

Juniper provides QFX Series switches optimized for IP fabric topologies, allowing for scalable deployments in modern data centers.

EVPN-VXLAN: Often used in IP fabrics to extend Layer 2 services across the fabric with Layer 3 underlay, enabling both efficient routing and bridging.

Question: 7

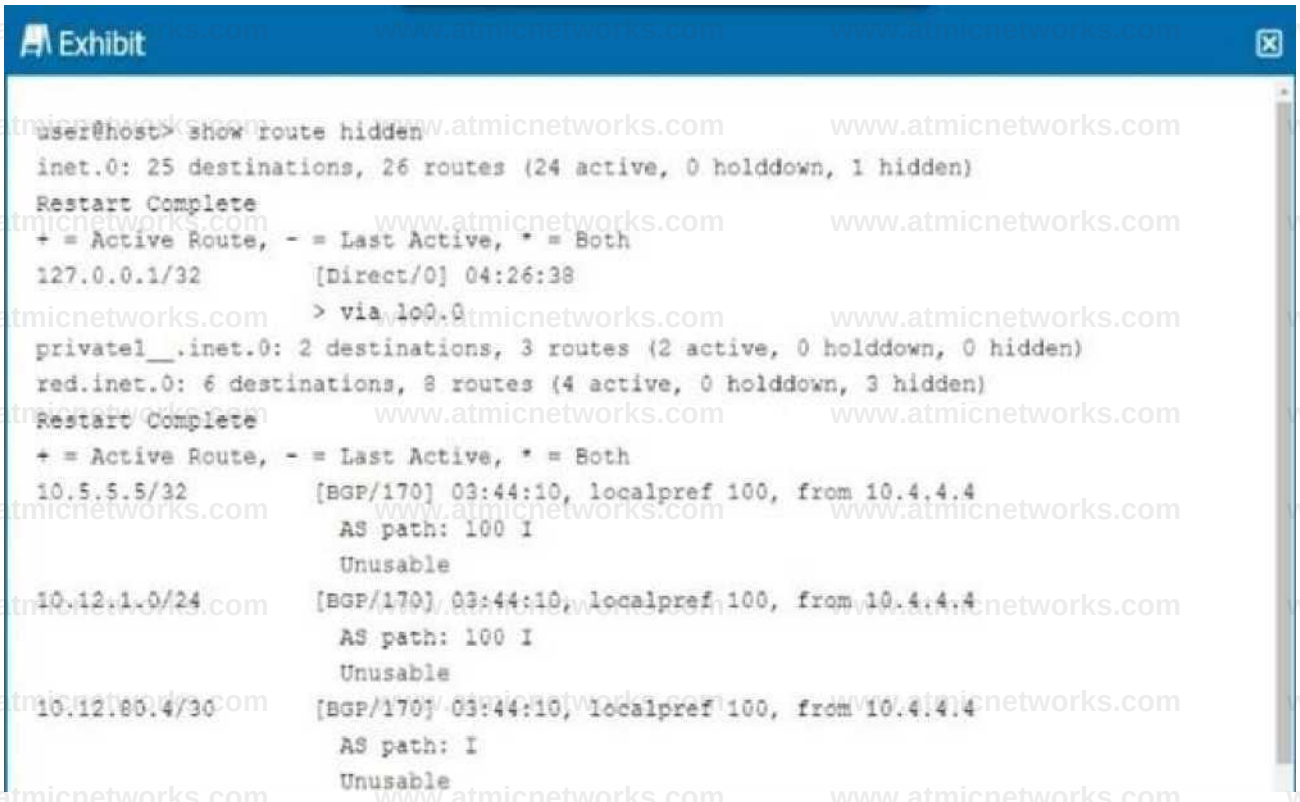

Referring to the exhibit, why are the BGP routes hidden?

A. Load balancing is not enabled.
B. There are too many hops to the destination.
C. The BGP next hop is unreachable.
D. Other routes are selected because of better metrics.

Answer: C

Explanation:

In the exhibit, the BGP routes are marked as hidden. This typically happens when the routes are not considered valid for use, but they remain in the routing table for reference. One common reason for BGP routes being hidden is that the next hop for these routes is unreachable.

Step-by-Step Breakdown:

BGP Next Hop:

In BGP, when a route is received from a neighbor, the next hop is the IP address that must be reachable for the route to be used. If the next hop is unreachable (i.e., the router cannot find a path to the next-hop IP), the route is marked as hidden.

Analyzing the Exhibit:

The exhibit shows that the BGP next hop for all hidden routes is 10.4.4.4. If this IP is unreachable, the BGP routes from that neighbor will not be considered valid, even though they appear in the routing table.

Verification:

Use the command show route 10.4.4.4 to check if the next-hop IP is reachable.

If the next-hop is not reachable, the BGP routes will be hidden. Resolving the next-hop reachability issue (e.g., fixing an IGP route or an interface) will allow the BGP routes to become active.

Juniper Reference:

Junos Command: show route hidden displays routes that are not considered for forwarding.

Troubleshooting: Check the next hop reachability for hidden BGP routes using show route <next- hop>.

Question: 8

Which statement is correct about the BGP AS path when advertising routes?

A. The order of the AS path is not significant.
B. The local AS number is added to the end of the AS path.
C. The order of the AS path is only significant in IBGP.
D. The local AS number is added to the beginning of the AS path.

Answer: D

Explanation:

The BGP AS (Autonomous System) path attribute is crucial in path selection and loop prevention. Each BGP router appends its local AS number to the beginning of the AS path when it advertises a route to an external BGP (eBGP) peer.

Step-by-Step Breakdown:

AS Path Attribute:

The AS path is a sequence of AS numbers that a route has traversed to reach a destination. Each AS adds its number to the front of the path, allowing BGP to track the route's history.

Why the Local AS is Added at the Beginning:

When advertising a route to an eBGP neighbor, a BGP router adds its own AS number to the beginning of the AS path. This ensures that the AS path reflects the route's journey accurately from the origin to the destination, and prevents loops in BGP. If the route returns to the same AS, the router will detect its AS number in the path and reject the route, preventing routing loops.
Order of the AS Path:
The order is significant because BGP uses it to select the best path. A shorter AS path is preferred, as it indicates fewer hops between the source and destination.
Juniper Reference:
AS Path Attribute: Junos devices append the local AS at the start of the AS path before advertising the route to an external peer.

Question: 9

Which statement is correct about a three-stage IP fabric underlay?

A. Every ingress interface into the fabric is only two hops away from the egress interface.
B. Every spine device can communicate directly with other spine devices.
C. Every leaf device can communicate directly with other leaf devices.
D. Every server that connects to a three-stage IP fabric must be multihomed.

Answer: A
Explanation:

In a three-stage IP fabric (also known as a Clos fabric), traffic between any two points (ingress to egress) in the fabric is only two hops away.
Step-by-Step Breakdown:

Three-Stage IP Fabric:
Leaf Layer: Leaf switches connect directly to servers and edge devices.
Spine Layer: Spine switches provide connectivity between leaf switches but do not connect to each other directly.
Two-Hop Communication:
In this architecture, every leaf switch is connected to every spine switch. Therefore, when a packet enters the fabric via an ingress leaf switch, it is forwarded to a spine switch, which then directs the packet to the correct egress leaf switch. This path always involves exactly two hops:
Ingress leaf → Spine → Egress leaf. Benefits:
This consistent two-hop path ensures predictable latency and makes the network highly scalable while maintaining low complexity.
Juniper Reference:
IP Fabric Architecture: This two-hop property of Clos fabrics is a hallmark of spine-leaf designs, as supported by Juniper's QFX and EX switches in data centers.

Question: 10

A routing policy has been created to advertise OSPF routes in BGP. Which statement is correct in this scenario?

A. Apply thepolicyas anexport policywithin BGP.
B. Apply thepolicyas anexport policywithin OSPF.
C. Apply thepolicyas animport policywithin BGP.
D. Apply thepolicyas animport policywithin OSPF.

Answer: A
Explanation:

When advertising OSPF routes into BGP, the appropriate routing policy should be applied as an export policy in BGP.
Step-by-Step Breakdown:
OSPF to BGP Route Advertisement:
Routes learned via OSPF (a dynamic IGP) need to be exported into BGP to be advertised to external BGP peers. In Junos OS, this is done using export policies.
Export Policies in BGP:
An export policy controls which routes are advertised out of a BGP session. In this scenario, the routing policy must be applied to BGP as an export policy to export the OSPF-learned routes to external BGP peers.
Policy Configuration:

Example configuration:

set policy-options policy-statement EXPORT_OSPF term 1 from protocol ospf set policy-options policy-statement EXPORT_OSPF term 1 then accept
set protocols bgp group <group-name> export EXPORT_OSPF This policy ensures that only OSPF routes are exported into BGP.
Juniper Reference:
Routing Policy: Export policies are used in BGP to control route advertisements to peers, including those learned via OSPF.

Question: 11

Which statement is correct about member interfaces when creating a LAG?

A. The interface's MTU settings must match on all member interfaces.
B. The interface's duplex settings and link speed must be the same on all member interfaces.
C. Member interfaces must all be allocated on the same chassis when using a Virtual Chassis.
D. Member interfaces must all be allocated on the same PFE.

Answer: B
Explanation:

When creating a LAG (Link Aggregation Group) in Junos, the duplex settings and link speed must be the same across all member interfaces.
Step-by-Step Breakdown:
LAG Overview:
A LAG combines multiple physical interfaces into a single logical interface to increase bandwidth and provide redundancy. All member links must act as a single cohesive unit.
Interface Requirements:
Duplex: All member interfaces must operate in the same duplex mode (either full-duplex or half- duplex). Mismatched duplex settings can cause performance issues, packet drops, or interface errors. Link Speed: All interfaces in the LAG must have the same link speed (e.g., all interfaces must be 1 Gbps or 10 Gbps). Mismatched speeds would prevent the interfaces from functioning correctly within the LAG.
Configuration and Validation: Ensure that all member interfaces have identical settings before adding them to the LAG. These settings can be checked using the show interfaces command, and the LAG can be configured using:

set interfaces ae0 aggregated-ether-options link-speed 10g set interfaces ge-0/0/1 ether-options 802.3ad ae0
Juniper Reference:
LAG Configuration: Duplex and link speed must be consistent across member interfaces to ensure proper LAG operation in Juniper devices.

Question: 12

Which three actions are required to implement filter-based forwarding? (Choose three.)

A. You must create an instance-type forwarding routing instance.
B. You must create an instance-type vrf routing instance.
C. You must create a match filter.
D. You must create a security policy.
E. You must create a RIB group.

Answer: A, C, E
Explanation:

Filter-Based Forwarding (FBF) in Junos OS allows traffic to be routed based on specific criteria such as source address, rather than just the destination address. This is useful in scenarios like policy routing or providing multiple paths for different types of traffic.
Step-by-Step Breakdown:

Instance-Type Forwarding:
You must create an instance-type forwarding routing instance. This routing instance allows for different routing tables based on the incoming packet filter.
Command:

set routing-instances FBF-instance instance-type forwarding Match Filter:
You need to create a filter to match the traffic that will be forwarded according to your custom routing policy. This filter is applied to an interface to determine which traffic will use the custom forwarding instance.
Command Example:

set firewall family inet filter FBF-filter term 1 from source-address <address> set firewall family inet filter FBF-filter term 1 then routing-instance FBF-instance RIB Group:
A RIB (Routing Information Base) group is necessary to share routes between the primary routing table and the custom routing instance. This allows FBF traffic to use the routing information from other routing tables.
Command Example:

set routing-options rib-groups FBF-group import-rib inet.0
set routing-instances FBF-instance routing-options rib-group FBF-group
Juniper Reference:
FBF Configuration: Filter-based forwarding requires these specific steps to redirect traffic to a custom routing table based on filter criteria.

Question: 13

Which signaling protocol is used for EVPN?

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

Answer: D
Explanation:

EVPN (Ethernet Virtual Private Network) is a standard protocol used for building Layer 2 and Layer 3 VPNs over an IP or MPLS network. The signaling protocol used for EVPN is BGP (Border Gateway Protocol).
Step-by-Step Breakdown:
BGP as the EVPN Signaling Protocol:
EVPN uses BGP to exchange MAC address reachability information between routers (PE devices). This enables devices to learn which MAC addresses are reachable through which PE devices, facilitating Layer 2 forwarding across an IP or MPLS core.

BGP Extensions for EVPN:
BGP is extended with new address families (e.g., EVPN NLRI) to carry both MAC and IP address information, allowing for scalable and efficient multi-tenant network solutions.
Juniper Reference:
Junos EVPN Configuration: Juniper uses BGP as the control plane for EVPN to exchange MAC and IP

route information between different data center devices.

Question: 14

Which operation mode command will display the mapping between the VLAN ID and ports on a switch?

A. show route
B. show ethernet-switching table
C. show interfaces terse
D. show vlans

Answer: D

Explanation:

To display the mapping between VLAN IDs and ports on a Juniper switch, the show vlans command is used.
Step-by-Step Breakdown:
VLAN Information:
The show vlans command displays detailed information about VLAN configurations, including the VLAN ID, associated interfaces (ports), and VLAN membership.
Command Example:
show vlans

This command will provide an output listing each VLAN, its ID, and the interfaces associated with the VLAN, enabling network engineers to quickly verify VLAN to port mappings.

Juniper Reference:
VLAN Verification: Use the show vlans command to verify which VLANs are configured on the switch and the ports that are members of those VLANs.

Question: 15

Within your router, you want to verify that you are learning routes from a remote BGP peer at IP address 10.10.100.1. Which command would satisfy the requirement?

A. show route receive-protocol bgp 10.10.100.1
B. show route protocol bgp table inet.0 10.10.100.1
C. show route advertise-protocol bgp 10.10.100.1
D. show route protocol bgp source-gateway 10.10.100.1

Answer: A

Explanation:

To verify that your router is learning routes from a remote BGP peer at a specific IP address (e.g., 10.10.100.1), the correct command to use is show route receive-protocol bgp.
Step-by-Step Breakdown:
BGP Route Learning:
The show route receive-protocol bgp command displays the routes that have been received from a specified BGP peer. This helps in confirming that the remote peer is sending routes correctly and that

your router is receiving them.
Command Example:
show route receive-protocol bgp 10.10.100.1

This will show all routes that have been received from the BGP peer with IP address 10.10.100.1.
Juniper Reference:
BGP Route Verification: Use this command to troubleshoot and verify that routes from a specific BGP peer are being received.

Question: 16

When a MAC limiting violation occurs, the switch performs which two actions by default? (Choose two.)

A. No logging takes place.
B. It causes Layer 2 loops.
C. The port is disabled.
D. It drops the packet.

Answer: C, D
Explanation:

When a MAC limiting violation occurs on a Juniper switch, the switch will perform the following actions by default:

Step-by-Step Breakdown:
Port Disabled:
When the number of MAC addresses on an interface exceeds the configured limit, the port is automatically disabled to prevent further violations. This is a protective mechanism to prevent MAC address flooding.
Packet Dropped:
Additionally, packets from the violating MAC address are dropped to prevent any further communication from that address. This ensures that only valid MAC addresses are allowed to communicate through the interface.
Example Configuration:

set ethernet-switching-options secure-access-port interface <interface-name> mac-limit 5

If more than five MAC addresses are learned, the port is disabled, and excess packets are dropped.
Juniper Reference:
MAC Limiting: When the switch detects a MAC limiting violation, it disables the port and drops further packets from the violating MAC addresses to maintain network security.

Question: 17

What information in the Ethernet header is used to populate the bridging table?

A. destination address
B. source address
C. type

D. protocol

Answer: B

Explanation:

The source MAC address in the Ethernet header is used to populate the bridging table (also called the MAC address table) on a switch. When a frame arrives at a switch, the switch examines the source MAC address and records it along with the ingress port in its MAC address table.
Step-by-Step Breakdown:
Learning Process:
When an Ethernet frame arrives on a switch port, the switch looks at the source MAC address and adds this MAC address to the MAC table along with the port it was received on. This process is called MAC learning.
Purpose:
The switch uses this information to determine the correct port to send frames destined for that MAC address in future transmissions, thus ensuring efficient Layer 2 forwarding.
Juniper Reference:
Ethernet Switching: Juniper switches use source MAC addresses to build and maintain the MAC address table, which is essential for Layer 2 switching.

Question: 18

You are configuring an aggregate route. In this scenario, which two statements are correct? (Choose two.)

A. Reject will silently drop the traffic.
B. Discard will silently drop the traffic.
C. Reject will send an ICMP Destination Unreachable message back to the sender.
D. Discard will send an ICMP Destination Unreachable message back to the sender.

Answer: B, C

Explanation:

When configuring an aggregate route, you have options for how to handle traffic that matches the route but does not match any more specific route in the routing table. Two actions can be taken: discard and reject.
Step-by-Step Breakdown:
Discard:
The discard option will silently drop packets that match the aggregate route. No notification is sent to the sender, and the packet is simply dropped.
Reject:

The reject option will drop the packet and also send an ICMP Destination Unreachable message back to the sender. This informs the sender that the packet could not be delivered because there is no specific route available.
Juniper Reference:
Aggregate Routes: The reject and discard next-hop options provide different levels of feedback when packets cannot be routed, and they can be used to control how unreachable destinations are handled.

Question: 19

What are two requirements for an IP fabric? (Choose two.)

A. a Layer 3 routing protocol
B. a single connection between each spine and leaf
C. a single connection between each leaf
D. a Layer 2 switching protocol

Answer: A, B

Explanation:

An IP fabric is a network architecture commonly used in data centers to provide scalable, high-throughput connectivity using a spine-leaf topology.
Step-by-Step Breakdown:
Layer 3 Routing Protocol:
An IP fabric relies on a Layer 3 routing protocol, typically BGP or OSPF, to provide routing between the leaf and spine switches. This ensures efficient traffic forwarding across the network.
Single Connection Between Spine and Leaf:
In an IP fabric, each leaf switch connects to every spine switch with a single connection. This ensures that traffic between any two leaf switches can travel through the spine layer in just two hops.
Juniper Reference:
Spine-Leaf Design: Juniper's IP fabric implementations are designed for scalability and low-latency routing, often using protocols like BGP for Layer 3 control.

Question: 20

What is the main purpose of Bidirectional Forwarding Detection (BFD)?

A. to detect network path failures
B. to determine if the forwarding routes are correct
C. to detect the forwarding protocol
D. to determine packet round-trip latency

Answer: A

Explanation:

Bidirectional Forwarding Detection (BFD) is a network protocol used to detect failures in the network path between two devices quickly.
Step-by-Step Breakdown:
Path Failure Detection:

BFD provides a low-overhead mechanism for detecting failures in forwarding paths across Layer 3 networks. It is much faster than traditional routing protocol timers and can detect failures within milliseconds.
BFD in Routing:
BFD can be integrated with routing protocols like OSPF, BGP, or IS-IS to trigger a faster convergence when a network path goes down.
Juniper Reference:

BFD Configuration: Juniper devices use BFD to monitor network paths and ensure fast failure detection, enhancing network resilience.

Question: 21

Which statement is correct about per-flow load balancing?

A. Packets associated with the same flow are sent through different egress ports.
B. The packets are guaranteed to arrive at their destination in a different order in which they were sent.
C. Packets associated with the same flow are sent through the same egress port.
D. The packets are guaranteed to arrive at their destination in the same order in which they were sent.

Answer: C

Explanation:

Per-flow load balancing ensures that packets within the same flow are always forwarded over the same path, ensuring that packet order is preserved.

Step-by-Step Breakdown:
Flow Definition:
A flow is typically defined by a combination of packet attributes like source/destination IP, source/destination port, and protocol type. Packets that belong to the same flow are routed over the same path to avoid reordering.
Per-Flow Behavior:
In per-flow load balancing, the hashing algorithm ensures that all packets in a particular flow use the same egress port, maintaining order across the network.
Juniper Reference:
Load Balancing in Juniper: This method ensures that flows are balanced across multiple paths while preventing packet reordering within a single flow.

Question: 22

You want to minimize topology disruptions in your network when the rpd process restarts on a device. Which service would accomplish this task?

A. Bidirectional Forwarding Detection (BFD)
B. link aggregation groups
C. graceful restart (GR)
D. Virtual Chassis

Answer: C

Explanation:

Graceful Restart (GR) is a feature that allows a router to maintain forwarding even when the routing process (e.g., the rpd process in Junos) is restarting, minimizing disruption to the network.
Step-by-Step Breakdown:
Graceful Restart Function:
During a GR event, the forwarding plane continues to forward packets based on existing routes, while the

control plane (rpd process) is restarting. This prevents traffic loss and maintains routing stability.
Minimizing Disruptions:
GR is particularly useful in ensuring continuous packet forwarding during software upgrades or routing protocol process restarts.
Juniper Reference:
Graceful Restart in Junos: GR ensures high availability by maintaining forwarding continuity during control plane restarts, enhancing network reliability.

Question: 23

Which two statements are true about how switches handle Layer 2 traffic? (Choose two.)

A. The MAC address is learned based on the destination MAC address.
B. The MAC address is learned based on the source MAC address.
C. Traffic is forwarded based on the source MAC address.
D. Traffic is forwarded based on the destination MAC address.

Answer: B, D
Explanation:

In Layer 2 switching, switches learn MAC addresses based on the source MAC address of incoming frames and forward frames based on the destination MAC address.
Step-by-Step Breakdown:
MAC Learning:
When a switch receives a frame, it records the source MAC address and the port on which it arrived. This allows the switch to know where to send traffic destined for that MAC address.
Forwarding Based on Destination:
The switch then looks at the destination MAC address and forwards the frame out of the port associated with that MAC address. If the MAC is unknown, the switch floods the frame to all ports.
Juniper Reference:
Layer 2 Switching: Juniper switches use source MAC addresses to build MAC tables and forward traffic based on the destination MAC address.

Question: 24

What are two consequences of having all network devices in a single collision domain? (Choose two.)

A. The amountof network resourceconsumption does not change.
B. The chanceof packet collision is decreased.
C. The chanceof packet collision is increased.
D. The amountof network resourceconsumption is increased.

Answer: C, D
Explanation:

A collision domain is a network segment where data packets can "collide" with one another when being sent on the same network medium.
Step-by-Step Breakdown:
Increased Collision Probability:
If all devices are in a single collision domain, the likelihood of packet collisions increases as more devices

attempt to send packets simultaneously, leading to network inefficiencies.
Increased Resource Consumption:
More collisions result in increased network resource consumption as devices need to retransmit packets, causing higher utilization of bandwidth and slowing down network performance.
Juniper Reference:
Collision Domains: Proper network segmentation using switches reduces collision domains, thereby improving network performance and reducing packet collisions.

Question: 25

Which statement is correct about IBGP?

A. It requires a physical full mesh.
B. It requires a logical full mesh.
C. It ensures that the local and remote peers use different AS numbers.
D. It ensures that duplicate AS numbers are not present in the AS path.

Answer: B

Explanation:

In IBGP (Internal Border Gateway Protocol), all routers within the same AS (Autonomous System) must have a logical full-mesh topology. This means that every IBGP router must be able to communicate with every other IBGP router directly or indirectly to ensure proper route propagation. Step-by-Step Breakdown:
Logical Full Mesh:
In an IBGP setup, routers do not re-advertise routes learned from one IBGP peer to another IBGP peer. This rule is in place to prevent routing loops within the AS.
To ensure full route propagation, a logical full mesh is required, meaning every IBGP router must peer with every other IBGP router in the AS. This can be done either directly or via route reflection or confederation.
Physical Full Mesh Not Required:
The physical topology does not need to be a full mesh, but the BGP peering relationships must form a logical full mesh. Techniques like route reflectors or BGP confederations can reduce the need for manual full-mesh peering.
Juniper Reference:
IBGP Configuration: IBGP logical full mesh requirements can be simplified using route reflectors to avoid the complexity of manually configuring many IBGP peers.

Question: 26

Which three technologies improve high availability and convergence in a data center network? (Choose three.)

A. graceful restart (GR)
B. Bidirectional Forwarding Detection (BFD)
C. link loss adjacency
D. Failover Group (FG)
E. link aggregation group (LAG)

Answer: A, B, E

Explanation:

High availability and fast convergence are critical in data center networks to minimize downtime and maintain optimal performance. The following technologies contribute to achieving these goals: Graceful Restart (GR):
GR allows routers to maintain forwarding state during control plane restarts, ensuring continuous packet forwarding while minimizing network disruptions.
Bidirectional Forwarding Detection (BFD):
BFD provides fast detection of path failures, allowing routing protocols to converge quickly by detecting link failures much faster than traditional timers.
Link Aggregation Group (LAG):
LAG increases both redundancy and bandwidth by combining multiple physical links into one logical link, providing load balancing and fault tolerance.
Juniper Reference:

High Availability Techniques: These technologies are fundamental in ensuring rapid recovery and failover within Juniper-based data center environments.

Question: 27

Which two statements are correct about rules for EBGP and IBGP? (Choose two.)

A. EBGP peers have a TTL of 1, while IBGP peers have a TTL of 255.
B. EBGP peers have a TTL of 255, while IBGP peers have a TTL of 1.
C. EBGP routes are more preferred than IBGP routes.
D. IBGP routes are more preferred than EBGP routes.

Answer: A, C

Explanation:

EBGP (External BGP) and IBGP (Internal BGP) operate with different rules due to the nature of their relationships.
Step-by-Step Breakdown:
TTL Differences:
EBGP: By default, EBGP peers have a TTL of 1, meaning they must be directly connected, or the TTL needs to be manually increased for multihop EBGP.
IBGP: IBGP peers within the same AS have a TTL of 255, as they are expected to communicate over multiple hops within the AS.

Preference for EBGP Routes:
Routes learned via EBGP are typically preferred over IBGP routes. This is because EBGP routes are considered more reliable since they originate outside the AS, while IBGP routes are internal.
Juniper Reference:
BGP Configuration: The different handling of TTL and route preferences between EBGP and IBGP ensures proper route selection and security within Junos-based networks.

Question: 28

Which statement is correct about an IRB interface?

A. An IRBinterface switches traffic within the same VLAN.
B. An IRB interface trunks together VLANs on different switches.
C. An IRB interface is a physical Layer 3 interface that connects VLANs together.
D. An IRB interface is a Layer 3 interface that can be used to route between VLANs.

Answer: D

Explanation:

An IRB (Integrated Routing and Bridging) interface provides routing functionality between VLANs at Layer 3, allowing devices in different VLANs to communicate with each other.
Step-by-Step Breakdown:
IRB Functionality:
The IRB interface enables routing between different VLANs by acting as a Layer 3 gateway. Traffic within the same VLAN is handled by Layer 2 switching, while traffic between VLANs is routed through the IRB interface.
Layer 3 Routing Between VLANs:
Each VLAN can be assigned an IP address on the IRB interface, which allows traffic to flow between VLANs based on Layer 3 IP routing.
Juniper Reference:
IRB Interface Configuration: Juniper supports IRB for inter-VLAN routing on devices like the EX and QFX series switches, facilitating Layer 3 communication in data centers.

Question: 29

You want to enable a Junos device to support aggregated Ethernet interfaces. In this scenario, which configuration hierarchy would you use?

A. [edit switch-options]
B. [edit system]
C. [edit interfaces]
D. [edit chassis]

Answer: D

Explanation:

To configure aggregated Ethernet (AE) interfaces on a Junos device, the configuration is done under the [edit chassis] hierarchy.

Step-by-Step Breakdown:
Chassis Configuration:
The chassis configuration is responsible for enabling the hardware to support Link Aggregation Groups (LAGs), allowing multiple physical interfaces to be bundled into a single logical interface for load

balancing and redundancy.
Command Example:

set chassis aggregated-devices ethernet device-count <number>

This command enables a specific number of aggregated Ethernet interfaces on the device.
Juniper Reference:
LAG Configuration in Junos: The chassis hierarchy is used to allocate and manage hardware resources for aggregated Ethernet interfaces in Juniper devices.

Question: 30

A switch receives a frame with a MAC address of FF-FF-FF-FF-FF-FF. Which action will the switch take on this frame?

A. It will flood it out of all interfaces, except for the ingress interface.
B. It will flood it out of all interfaces, except for the directly connected VLAN.
C. It will flood it out of all interfaces, except for the next-hop interface.
D. It will flood it out of all interfaces.

Answer: A
Explanation:

A MAC address of FF-FF-FF-FF-FF-FF is the Ethernet broadcast address. When a switch receives a frame with this destination MAC address, it is required to forward the frame to all interfaces except the one it was received on.
Step-by-Step Breakdown:
Broadcast Frame Handling:
When a frame with the broadcast MAC address is received, the switch will flood it out of all active ports that belong to the same VLAN as the incoming frame. The broadcast frame is not sent back out of the ingress interface (the interface where the frame was originally received).
Purpose of Flooding:
Broadcasting is used to ensure that the frame reaches all devices within the broadcast domain (all devices within the same VLAN), which may not have a specific entry for the MAC address in their MAC address table.
Juniper Reference:
Layer 2 Frame Forwarding: Juniper switches flood broadcast frames to all ports in the same VLAN, except the port the frame was received on.

Question: 31

Referring to the exhibit, you notice that after committing the configuration, the ae0 and ae1 interfaces appear in a link down state.

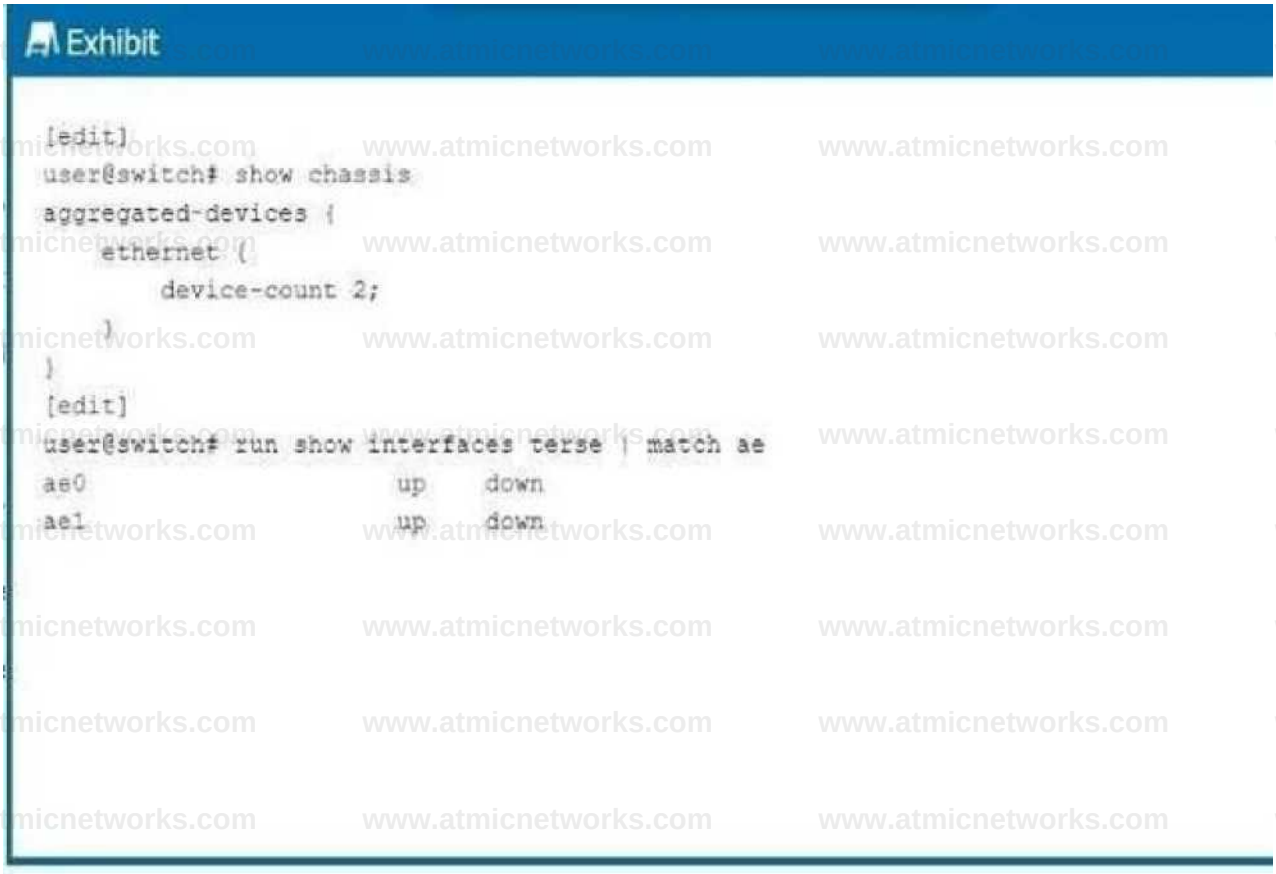

Which statement is correct in this scenario?

A. No operational interfaces have been added to the LAG interfaces.
B. No traffic is traversing the LAG interfaces.
C. The LAG interfaces are in a passive state.
D. The LAG interfaces are in aggressive mode.

Answer: A

Explanation:

In the exhibit, the ae0 and ae1 interfaces are in a link down state. This occurs when no physical interfaces (member interfaces) have been added to the LAG (Link Aggregation Group) interfaces, or the member interfaces are not operational.
Step-by-Step Breakdown:
LAG Configuration:
A LAG interface (aggregated Ethernet interface) is a logical interface that combines multiple physical interfaces for redundancy and increased bandwidth. The LAG will only be operational if at least one member interface is active and configured correctly.
No Operational Member Interfaces:
If no member interfaces are added or if the member interfaces are down, the LAG will remain in a down state, as shown in the exhibit for ae0 and ae1.
Resolution:
Verify that physical interfaces have been added to the LAG using commands like:
LAG Interface Status: In Juniper, the link status of the LAG depends on its member interfaces, which must be operational for the LAG to function.

Question: 32

Which two statements are correct about VLAN tags? (Choose two.)

A. VLAN tags carry a VLAN ID and priority.
B. VLAN tags are required on access ports.
C. VLAN tags require multiple forwarding tables.
D. VLAN tags can be inserted or removed by trunk interfaces.

Answer: A, D

Explanation:

VLAN tags are used in Ethernet frames to identify and differentiate traffic between multiple VLANs. They are especially important for devices like switches that handle multiple VLANs on the same physical link.
Step-by-Step Breakdown:
VLAN Tag Contents:
VLAN ID: The tag contains a 12-bit VLAN ID field that identifies the VLAN to which the frame belongs.

Priority: The tag also includes a 3-bit priority field (also known as 802.1p priority) used for QoS (Quality of Service) to prioritize traffic.
Trunk Ports and VLAN Tagging:
Trunk Ports are used to carry traffic for multiple VLANs across a single link. These interfaces insert (tag) VLAN identifiers into frames when they leave the switch and remove (untag) them when frames enter the switch.
Access Ports:
VLAN tags are typically not used on access ports (ports that connect to end devices) since those ports are configured to be part of a single VLAN, and the traffic doesn't need VLAN tags.
Juniper Reference:
VLAN Tagging: Juniper switches support VLAN tagging and ensure that frames are tagged or untagged as they traverse trunk or access ports, respectively.

Question: 33

Exhibit:

B Exhibit

■edit protocols cspf‘
□sergrouter# show
area 0.9.0.0 (
interface xe~0/0/4.0 {
Md-liveness-detection (mr.iir.nt-interval 400; multiplier 5;

Referring to the exhibit, at which interval will the interface be considered down if no hello packets are received?

A. 2000 seconds
B. 400 milliseconds
C. 400 seconds
D. 2000 milliseconds

Answer: D

Explanation:

The exhibit shows the configuration of Bidirectional Forwarding Detection (BFD) for OSPF on interface xe-0/0/4.0, with the following parameters: minimum-interval: 400 milliseconds multiplier: 5
Step-by-Step Breakdown:
BFD Liveness Detection:
BFD is used to detect link failures at sub-second intervals, providing faster convergence times for routing protocols like OSPF. The minimum-interval is the time between BFD control packets (in milliseconds), and the multiplier indicates how many missed BFD packets trigger a failure.
Calculating Failure Detection Time:
The failure detection interval is calculated as:
Failure Interval=minimum-interval×multiplier ext{Failure Interval} = ext{minimum-interval} imes ext{multiplier}Failure Interval=minimum-interval×multiplier In this case:
400?milliseconds×5=2000?milliseconds(2seconds)400 \, ext{milliseconds} imes 5 = 2000 \, ext{milliseconds} (2 seconds)400milliseconds×5=2000milliseconds(2seconds) Conclusion: If no BFD control packets are received within 2000 milliseconds (2 seconds), the interface will be

considered down, triggering OSPF to recalculate routes.
Juniper Reference:
BFD Configuration: BFD parameters such as minimum-interval and multiplier are used to fine-tune the failure detection time for faster convergence.

Question: 34

In the Junos OS, which feature is used to create an alternate next hop with a unique preference for a static route?

A. Preference
B. Resolve
C. Next-hop
D. Qualified-next-hop

Answer: D

Explanation:

In Junos OS, the qualified-next-hop feature is used to specify an alternate next hop for a static route, along with a unique preference value.
Step-by-Step Breakdown:
Qualified-Next-Hop:
A qualified-next-hop allows you to define multiple next hops for a static route, each with its own preference. This provides flexibility by allowing the router to choose the best available next hop based on

reachability and preference.
Use Case:

If the primary next hop becomes unreachable, the router can automatically switch to the alternate next hop defined by the qualified-next-hop with a higher preference value.
Command Example:
set routing-options static route 10.10.10.0/24 qualified-next-hop 192.168.1.1 preference 5 set routing-options static route 10.10.10.0/24 qualified-next-hop 192.168.1.2 preference 10

Preference:
The next hop with the lowest preference is chosen first. If it becomes unavailable, the router will use the higher preference next hop.
Juniper Reference:
Qualified-Next-Hop: This feature is used to configure backup or alternate next hops for static routes in Juniper devices.

Question: 35

Exhibit:

B Exhibit

[edit] user^routerl show protocols bgp (
group 1 I local-as 65101;

neighbor 172.1€.1.1 { peer-as €5261;

1

Iedit] userQroutert show routing-opt ions router-id 1$2.let.1JO.1; autonomous-systern €50C0;

Referring to the exhibit, which statement is correct?

A. The configuration will commit successfully and BGP group1will operate as IBGP.
B. The configuration will commit successfully and BGP group1will operate asEBGP.
C. BGP group 1 requires a type external parameter.
D. BGP group 1 requires a type internal parameter.

Answer: B

Explanation:

In the exhibit, BGP is configured with local AS 65101 and a neighbor at 172.16.1.1 in peer AS 65201. This setup involves two different Autonomous Systems (AS), indicating an External BGP (EBGP) configuration.
Step-by-Step Breakdown:

EBGP vs. IBGP:
EBGP is used between routers in different ASes. In this case, the local AS is 65101 and the peer AS is 65201, meaning the BGP session is EBGP.
IBGP is used between routers within the same AS, which is not applicable here as the AS numbers are different.
BGP Group Configuration:
The configuration does not require a type external parameter because Junos OS automatically recognizes the session as EBGP when the local and peer AS numbers are different.
The BGP session will operate as EBGP, and the configuration will commit successfully.
Juniper Reference:
BGP Configuration: In Juniper, EBGP is automatically recognized when the local and peer AS numbers differ, without needing to specify type external.

Question: 36

Which statement is correct about areas in OSPF?

A. An OSPF area is used to segment Layer 2 broadcast domains.
B. OSPF areas are used to isolate the effects of a broadcast storm.
C. OSPF areas are used to reduce the size of the link-state database.
D. An OSPF area is used to signify the autonomous system to which each device belongs.

Answer: C
Explanation:

In OSPF (Open Shortest Path First), areas are used to segment a network into smaller, more manageable pieces to improve scalability. By dividing a network into areas, OSPF can reduce the size of the link-state database (LSDB), which helps routers process updates more efficiently.
Step-by-Step Breakdown:
Purpose of OSPF Areas:
OSPF areas allow for hierarchical routing within the OSPF domain. Routers in the same area have identical LSDBs, but routers in different areas do not exchange full link-state information. Instead, they exchange summarized routes, which reduces the LSDB size and CPU/memory usage.
Benefits:
Reducing the LSDB size improves scalability and ensures faster convergence in larger networks. Area 0 is the backbone area, and all other areas must connect to it, forming a hierarchical structure.
Juniper Reference:
OSPF Configuration: Areas in OSPF are configured to optimize network performance by limiting the scope of link-state advertisements (LSAs) to within an area.

Question: 37

What are two reasons why you would deploy an IP fabric instead of a traditional Layer 2 network in a data center? (Choose two.)

A. Layer 2 networks only support a single broadcast domain.
B. IP fabrics are better suited to smaller networks where scale is less important.
C. Layer 3 networks support load balancing.
D. Layer 2 networks are susceptible to loops.

Answer: C, D

Explanation:

IP fabrics are Layer 3-centric network designs often used in data centers due to their scalability, efficient routing, and loop-free architecture.

Step-by-Step Breakdown:

Layer 3 Load Balancing:

IP fabrics use Equal-Cost Multipath (ECMP) to distribute traffic across multiple paths, providing effective load balancing and improving bandwidth utilization. This capability is absent in traditional Layer 2 networks, which do not support ECMP for routing decisions.

Layer 2 Loops:

Layer 2 networks are prone to loops because of the lack of TTL (Time-to-Live) mechanisms. Spanning Tree Protocol (STP) is required to prevent loops, but it can introduce inefficiencies by blocking links. In contrast, IP fabrics based on Layer 3 protocols are loop-free and do not need STP.

Juniper Reference:

IP Fabric: Juniper's IP fabric solutions offer efficient Layer 3 routing with built-in load balancing and loop prevention, making them ideal for modern data center architectures.

Question: 38

Which two statements are correct about EVPN-VXLAN overlay networking? (Choose two.)

A. It is the only option to provide reachability between servers that reside in the same network segment in a data center.
B. BGP provides the control plane within the overlay network.
C. An encapsulation of the original packet is required to transport the packet across the network.
D. OSPF provides the control plane within the overlay network.

Answer: B, C

Explanation:

EVPN-VXLAN is an overlay technology used in data center networks to extend Layer 2 services over a Layer 3 network.

Step-by-Step Breakdown:

BGP Control Plane:

BGP (Border Gateway Protocol) is used as the control plane for EVPN-VXLAN. BGP advertises MAC addresses and IP address reachability information across the VXLAN network, enabling efficient multi-tenant Layer 2 connectivity over a Layer 3 infrastructure.

Encapsulation:

VXLAN (Virtual Extensible LAN) encapsulates Layer 2 frames into Layer 3 packets. This encapsulation allows Layer 2 traffic to be transported across a Layer 3 network, effectively creating a tunnel for Ethernet frames.

Juniper Reference:

EVPN-VXLAN Configuration: Juniper supports EVPN-VXLAN with BGP as the control plane, allowing scalable Layer 2 connectivity over a routed infrastructure in modern data centers.

Question: 39

A generated route is configured under which hierarchy?

A. [editpolicy-options]
B. [edit routing-instance]
C. [edit routing-options]
D. [edit protocols]

Answer: C
Explanation:

A generated route in Junos OS is configured under the [edit routing-options] hierarchy. Step-by-Step Breakdown:
Generated Routes:
A generated route is created based on the presence of more specific routes in the routing table. It acts as a summary route and is generated when any of its contributing routes are active. This is commonly used to create aggregate routes in OSPF, BGP, or other protocols.
Configuration Hierarchy:
The configuration for generated routes is placed under [edit routing-options], where other static and routing policies are also defined.
Command Example:
set routing-options generate route 10.10.0.0/16
Juniper Reference:
Routing Options: Juniper routers use the routing-options hierarchy to configure generated routes and other static routing behaviors.

Question: 40

MACsec provides protection against which two types of threats? (Choose two.)

A. Data decryption
B. Playback attacks
C. Hashing attacks
D. Man-in-the-middle attack

Answer: B, D
Explanation:

MACsec (Media Access Control Security) provides data confidentiality, integrity, and origin authenticity at Layer 2, protecting against several types of threats.
Step-by-Step Breakdown:
Man-in-the-Middle Attack Protection:
MACsec encrypts traffic at Layer 2, preventing man-in-the-middle attacks where an attacker intercepts and manipulates traffic between two communicating devices. Since the data is encrypted, any intercepted packets are unreadable.
Protection Against Playback Attacks:
MACsec also protects against playback attacks by using sequence numbers and timestamps to ensure that old, replayed packets are not accepted by the receiver.

Juniper Reference:
MACsec Configuration: Juniper devices support MACsec for securing Layer 2 communications, ensuring protection against replay and man-in-the-middle attacks in sensitive environments.

Question: 41

Exhibit:

AT Exhibit

```
[edit routing-options! userGrouter* show
static ( defaults I
        preference 7;

route 0.0.0.0/0 (
    next-hop 172.25.20.254 qualified-next-hop *72 preference 4;
```

Referring to the exhibit, which next hop will be preferred in the routing table?

A. Next hop IP address 172.25.20.254 will be preferred.
B. Neither next hop will be preferred.
C. Next hop IP address 172.25.20.200 will be preferred.
D. Both next hops will be preferred.

Answer: C

Explanation:

In the exhibit, we see a static route configuration with two possible next hops for the default route (0.0.0.0/0):
next-hop 172.25.20.254 with the default preference of 7. qualified-next-hop 172.25.20.200 with a preference of 6. Step-by-Step Breakdown:
Preference Value:
In Junos OS, the preference value is used to determine which route should be preferred in the routing table. The lower the preference value, the higher the priority for the route.
Comparison:
In this case:
The next hop 172.25.20.254 has a preference of 7.
The qualified-next-hop 172.25.20.200 has a preference of 6. Preferred Next Hop:
Since 172.25.20.200 has a lower preference (6) compared to 172.25.20.254 (7), it will be the preferred next hop in the routing table, assuming both next hops are reachable.
Juniper Reference:
Qualified Next Hop: In Junos, static routes with multiple next-hop options are selected based on the preference value, with the lower value being preferred.

Question: 42

Layer 2 interfaces operate in which two modes? (Choose two.)

A. Access
B. Modular
C. Trunk
D. Tagged

Answer: A, C

Explanation:

Comprehensive Detailed Step by Step Explanation with all Juniper Data Center References Layer 2 interfaces on a switch operate in two key modes: Access and Trunk.

Step-by-Step Breakdown:

Access Mode:

Access ports are used to connect end devices, like PCs or servers, and they are assigned to a single VLAN. These interfaces handle untagged traffic and do not pass VLAN tags.

Example: A port assigned to VLAN 10 will only handle traffic for that VLAN. Trunk Mode:

Trunk ports are used to connect switches or other networking devices that need to handle traffic from multiple VLANs. Trunk interfaces carry tagged traffic, allowing multiple VLANs to traverse the same physical link.

Trunk ports typically use 802.1Q VLAN tagging to differentiate between VLANs.

Juniper Reference:

Access and Trunk Ports: Juniper switches use these modes to manage VLAN traffic at Layer 2, with access ports handling untagged traffic and trunk ports handling tagged traffic from multiple VLANs.

Question: 43

Exhibit:

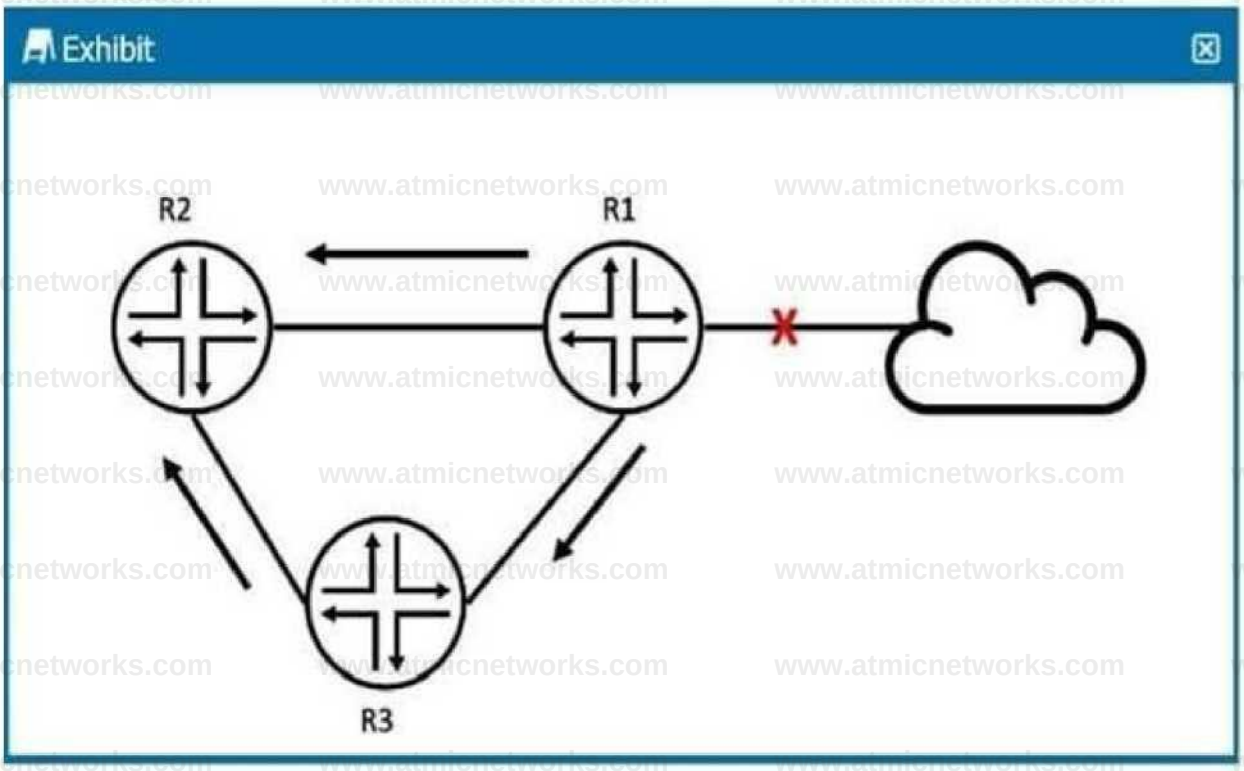

R2 received an OSPF update from R1, and it received the same update from R3. Referring to the exhibit, what will R2 do?

A. R2 ignores the update from R1.
B. R2 does nothing with R3's update.
C. R2 ignores the update from R3.
D. R2 acknowledges R3 and discards it.

Answer: C

Explanation:

In the exhibit, R2 receives the same OSPF update from both R1 and R3. OSPF has mechanisms to prevent unnecessary processing of duplicate LSAs (Link-State Advertisements).

Step-by-Step Breakdown:

OSPF LSA Processing:

OSPF uses LSAs to exchange link-state information between routers. When a router receives an LSA, it checks if it already has a copy of the LSA in its Link-State Database (LSDB).

Duplicate LSAs:

If R2 has already received and processed the update from R1, it will ignore the update from R3 because it already has the same LSA in its database. OSPF uses the concept of flooding, but it does not reprocess LSAs that it already knows about.

R2 Behavior:

R2 will keep the update from R1 (the first one it received) and will ignore the same LSA from R3, as it is already in the LSDB.

Juniper Reference:

OSPF LSA Processing: Junos adheres to OSPF standards, ensuring that duplicate LSAs are not processed multiple times to avoid unnecessary recalculations.

Question: 44

What is the definition of a trunk interface on a switch?

A. An interface that carries multiple VLANs.
B. An interface that carries high bandwidth.
C. An interface that connects directly to powerful servers.
D. An interface that carries excess traffic.

Answer: A

Explanation:

A trunk interface on a switch is used to carry traffic for multiple VLANs between switches or between a switch and another network device, like a router. Trunk interfaces use 802.1Q tagging to identify which VLAN the traffic belongs to.

Step-by-Step Breakdown:

Trunk Ports:

Trunk ports are typically used for inter-switch links or switch-to-router links where multiple VLANs need to be carried over the same physical connection.

VLAN traffic is tagged with a VLAN ID to ensure that it is properly identified as it crosses the trunk link.

802.1 Q VLAN Tagging:

Trunk ports use 802.1Q to tag Ethernet frames with the VLAN ID. This ensures that frames are correctly forwarded to the appropriate VLANs on the other side of the trunk.

Juniper Reference:

Trunk Interface Configuration: In Juniper switches, trunk ports are configured to carry tagged traffic for multiple VLANs, which is essential for interconnecting multiple network segments.

Question: 45

Which two statements about IBGP are correct? (Choose two.)

A. By default, IBGP has a TTLof 1.
B. IBGP uses AS path for loopprevention.
C. By default, IBGP has a TTLof 255.
D. IBGP uses full mesh for loop prevention.

Answer: C, D

Explanation:

IBGP (Internal Border Gateway Protocol) is used to exchange routing information between routers within the same AS (Autonomous System).
Step-by-Step Breakdown:
TTL of 255:
By default, IBGP sessions are established with a TTL (Time to Live) value of 255. This allows IBGP neighbors to communicate over multiple hops within the AS without requiring any additional configuration.
Full Mesh Requirement:
IBGP requires a logical full mesh between all IBGP routers to ensure that routing information is fully distributed within the AS. Since IBGP does not propagate routes learned from one IBGP peer to another by default, a full mesh topology is needed unless route reflectors or BGP confederations are used.
Juniper Reference:
IBGP Full Mesh: Juniper recommends using route reflectors in large networks to simplify IBGP full- mesh requirements.

Question: 46

Exhibit:

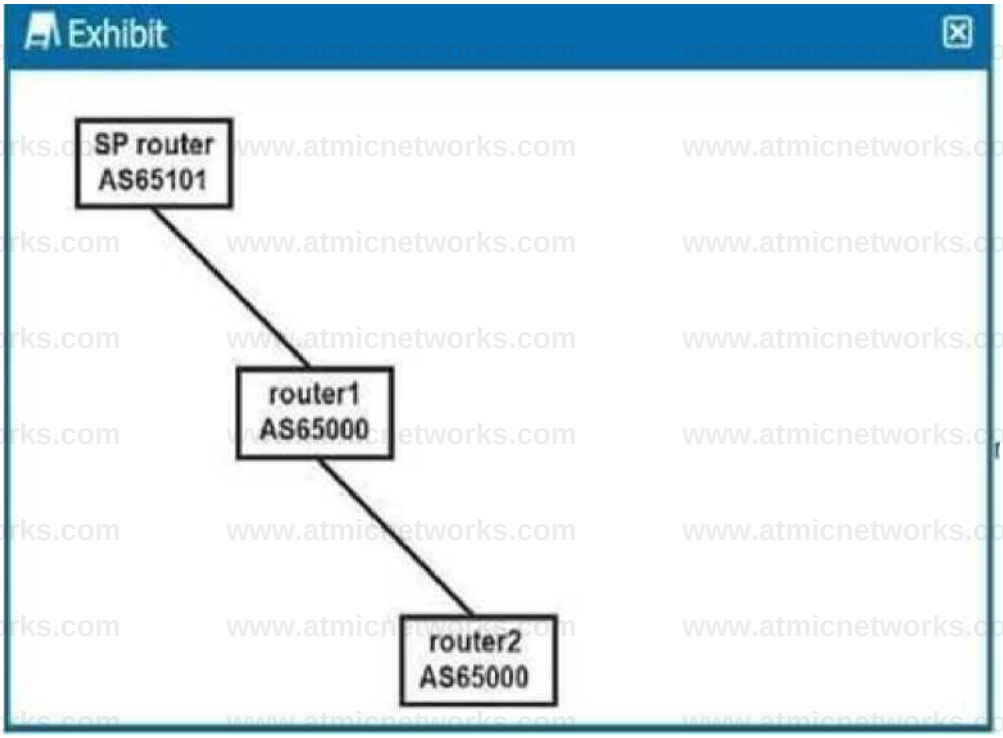

Referring to the exhibit, which two statements are correct about default BGP advertisements? (Choose two.)

A. When routes advertised by router2 are received by the SP router, they will contain the next-hop address of router2.
B. When routes advertised by router2 are received by the SP router, they will contain the next-hop address of router1.
C. When routes advertised by the SP router are received by router2, they will contain the next-hop address of the SP router.
D. When routes advertised by the SP router are received by router2, they will contain the next-hop address of router1.

Answer: B, D

Explanation:

The exhibit shows a BGP peering scenario between three routers: router1 and router2 are part of the same AS (AS65000), while the SP router is in a different AS (AS65101). This indicates an EBGP (External BGP) peering between the SP router and router1, and IBGP between router1 and router2.
Step-by-Step Breakdown:
Next-Hop Behavior in BGP:
IBGP: In IBGP, the next-hop address is not modified when advertising routes within the same AS. Thus, when router1 advertises routes learned from router2 to the SP router, it will keep the next-hop address of router1, not router2.
EBGP: In EBGP, the next-hop address is modified. When router1 receives routes from the SP router, it will advertise them to router2 with the next-hop address of router1.
Route Propagation:
Routes received by router1 from router2 will be advertised to the SP router with router1 as the next hop. Similarly, routes advertised by the SP router will be passed on to router2, with router1 remaining as the next hop.

Juniper Reference:
BGP Next-Hop: Juniper's BGP implementations follow standard BGP next-hop behavior, where the next-hop is modified in EBGP but not in IBGP, ensuring proper route advertisement across autonomous systems.

Question: 47

When considering bidirectional forwarding detection, which two statements are correct? (Choose two.)

A. The BFD default minimum interval is 3.
B. You can configure BFD per interface within the protocol stanza.
C. The BFD operation always consists of minimum intervals and multipliers.
D. The BFD default multiplier is 5.

Answer: B, C
Explanation:

Bidirectional Forwarding Detection (BFD) is a protocol used to detect faults in the forwarding path between two routers. It provides rapid failure detection, enhancing the performance of routing protocols like OSPF, BGP, and IS-IS.
Step-by-Step Breakdown:

Per Interface Configuration:
BFD can be configured on a per-interface basis within the protocol stanza (e.g., OSPF, BGP). This allows granular control over where BFD is enabled and the failure detection intervals for specific interfaces.
Minimum Interval and Multiplier:
BFD uses a minimum interval (the time between BFD control packets) and a multiplier (the number of missed packets before the path is declared down). The combination of these two defines the detection time for failures.
Juniper Reference:
BFD Configuration: In Juniper, BFD is configurable within routing protocol stanzas, with the failure detection mechanism always based on minimum intervals and multipliers.

Question: 48

How does OSPF calculate the best path to a particular prefix?

A. It finds the path with the numerically lowest cost.
B. It finds the path with the shortest autonomous system path.
C. It finds the path with the least number of hops.
D. It finds the path with the numerically lowest route preference.

Answer: A
Explanation:

OSPF (Open Shortest Path First) calculates the best path based on the cost of the route, which is derived from the bandwidth of the interfaces along the path.

Step-by-Step Breakdown:
OSPF Path Selection:

OSPF assigns a cost to each link, typically based on the link's bandwidth (higher bandwidth equals lower cost).
The OSPF algorithm computes the shortest path to a destination by adding the costs of all links in the path. The path with the numerically lowest total cost is chosen as the best path.
Cost Calculation:
The OSPF cost can be manually adjusted or automatically calculated using the default formula:
Cost=Reference BandwidthLink Bandwidth ext{Cost} = rac{ ext{Reference Bandwidth}}{ ext{Link Bandwidth}}Cost=Link BandwidthReference Bandwidth
Juniper Reference:
OSPF Best Path Selection: OSPF selects the path with the lowest cumulative cost, ensuring efficient use of higher-bandwidth links in Junos networks.

Question: 49

Which statement about switches is correct?

A. Each port is a member of VLAN 2 by default.
B. Every port is in a unique collision domain.
C. Each port is in a unique broadcast domain by default.
D. All ports reside in the same collision domain.

Answer: B
Explanation:

Each port on a modern switch creates a separate collision domain. This allows multiple devices to communicate simultaneously without collisions on different ports.
Step-by-Step Breakdown:
Collision Domain:
A collision domain is a network segment where data packets can collide if two devices send packets simultaneously.
On a switch, each port creates a separate collision domain, so collisions only occur if two devices connected to the same port (through a hub, for instance) try to send data at the same time.
Switches vs Hubs:
Unlike hubs, which have one large collision domain, switches isolate collisions to individual ports, improving performance.
Juniper Reference:
Switch Port Behavior: In Juniper switches, each port operates in its own collision domain, enhancing network efficiency by reducing the chances of packet collisions.

Question: 50

By default, which two statements are correct about BGP advertisements? (Choose two.)

A. BGP peers advertise routes received from EBGP peers to other IBGP peers.
B. BGP peers advertise routes received from IBGP peers to other IBGP peers.
C. BGP peers advertise routes from EBGP peers to other IBGP peers using its own address as the next hop.
D. BGP peers advertise routes from IBGP peers to EBGP peers using its own address as the next hop.

Answer: A, D

Explanation:

BGP (Border Gateway Protocol) has specific rules for route advertisement between peers. Step-by-Step Breakdown:
EBGP to IBGP Route Propagation:
BGP peers advertise routes learned from EBGP peers to IBGP peers within the same AS. This ensures that routes learned from external networks are propagated internally within the AS.
IBGP to EBGP Route Propagation:
Routes learned from IBGP peers can be advertised to EBGP peers, but when advertising these routes, the router uses its own IP address as the next hop.
IBGP Split Horizon:
By default, IBGP peers do not advertise routes learned from one IBGP peer to another IBGP peer. This rule (IBGP split horizon) prevents routing loops within an AS.
Juniper Reference:
BGP Advertisement Rules: Junos adheres to BGP standards, where IBGP peers do not propagate routes to other IBGP peers, but EBGP peers receive IBGP routes with the advertising router as the next hop.

Question: 51

Which static routing parameter will silently drop the packet if it is set as the next hop?

A. Reject
B. Resolve
C. Readvertise
D. Discard

Answer: D

Explanation:

When the discard option is configured as the next hop for a static route, it silently drops any packets that match the route without sending any notification to the sender.
Step-by-Step Breakdown:
Discard Behavior:
If a route uses the discard next hop, the router drops the packet without generating any ICMP message or error back to the sender. This is useful for creating null routes to prevent routing loops or blackhole traffic intentionally.
Reject vs. Discard:
The reject next hop, in contrast, drops the packet but sends an ICMP Destination Unreachable message back to the source.
Juniper Reference:
Static Route Behavior: In Junos, the discard option ensures packets matching a static route are dropped silently, providing a way to discard traffic without alerting the source.

Question: 52

What are two device roles in a five-member Virtual Chassis? (Choose two.)

A. PFE

B. Control-board
C. Line card
D. Routing-engine

Answer: C, D Explanation:

In a Virtual Chassis (VC) configuration, multiple Juniper switches are interconnected to form a single logical device. Each member switch in the Virtual Chassis plays a specific role.
Step-by-Step Breakdown:
Line Card Role:
Member switches acting as line cards provide additional ports for traffic forwarding but do not perform control or routing functions. These switches depend on the routing engine to handle control-plane tasks.

Routing Engine Role:
A switch in the routing-engine role is responsible for control-plane operations such as routing protocol management and control of the Virtual Chassis.
Virtual Chassis Roles:
Master Routing Engine: Handles control-plane functions and manages the entire Virtual Chassis. Backup Routing Engine: Takes over if the master fails.
Line Card: Provides additional ports and handles data-plane operations.
Juniper Reference:
Virtual Chassis: In a five-member Virtual Chassis, multiple switches act as line cards, while one or more switches are designated as the routing engines (master and backup).

Question: 53

Which state in the adjacency process do OSPF routers check the MTU size?

A. Init
B. Exchange
C. Done
D. ExStart

Answer: B
Explanation:

In OSPF, routers exchange link-state information in different stages to establish full adjacency. The MTU size is checked during the Exchange state.
Step-by-Step Breakdown:
OSPF Adjacency Process:
OSPF routers go through multiple stages when forming an adjacency: Down, Init, 2-Way, ExStart, Exchange, Loading, and Full.

Exchange State:
During the Exchange state, OSPF routers exchange Database Description (DBD) packets to describe their link-state databases. The MTU size is checked at this stage to ensure both routers can successfully exchange these packets without fragmentation.
If there is an MTU mismatch, the routers may fail to proceed past the Exchange state.
Juniper Reference:
MTU Checking in OSPF: Junos uses the Exchange state to check for MTU mismatches, ensuring that

routers can properly exchange database information without packet fragmentation issues.

Question: 54

Exhibit:

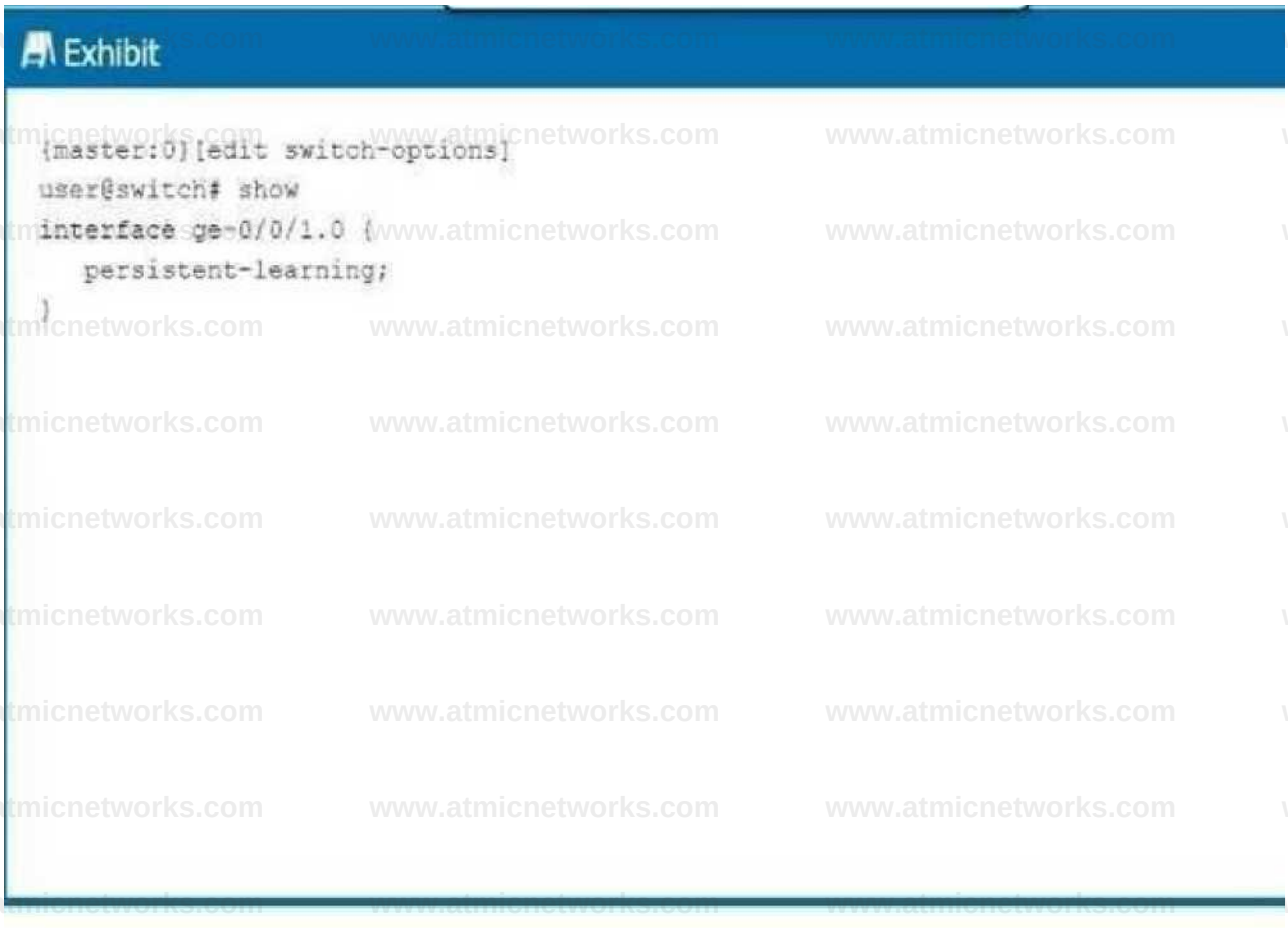

Referring to the exhibit, which behavior does this configuration enable on the ge-0/0/1.0 interface?

A. This configuration enables a MAC address learned on the interface to be persistently retained in the Ethernet-switching table, even after a reboot.
B. This configuration enables the device to place a MAC address that persistently causes network errors into a special protected VLAN.
C. This configuration enables the device to shut down the interface when a particular MAC address persistently sends broadcast traffic.
D. This configuration enables the interface to learn and remember MAC addresses, until the device is rebooted.

Answer: A

Explanation:

The configuration in the exhibit shows the persistent-learning feature enabled on interface ge- 0/0/1.0.
Step-by-Step Breakdown:
Persistent Learning:
Persistent-learning ensures that the MAC addresses learned on the interface are retained in the Ethernet-switching table, even after a device reboot. This prevents the need to re-learn MAC addresses after the device restarts, improving stability and reducing downtime.

Use Case:
This feature is particularly useful in environments where the re-learning of MAC addresses could cause temporary disruptions or delays in communication, such as in critical Layer 2 network segments.
Command Example:

set switch-options interface ge-0/0/1.0 persistent-learning
Juniper Reference:
Persistent MAC Learning: In Junos, enabling persistent-learning ensures that learned MAC addresses are not lost during reboots, contributing to smoother network operations in environments where stability is crucial.

Question: 55

When evaluating BGP routes, what will be evaluated first?

A. Thelocal preference value
B. The AS path
C. The MED value
D. The origin value

Answer: A
Explanation:

In BGP (Border Gateway Protocol), when evaluating multiple routes to the same destination, the first attribute that is considered is the local preference value. The local preference is a BGP attribute used to influence outbound routing decisions within an Autonomous System (AS).
Step-by-Step Breakdown:
Local Preference:
The local preference attribute is used to determine which path is preferred for traffic leaving the AS. The higher the local preference value, the more preferred the route.
BGP Path Selection:
The BGP path selection process evaluates the following attributes in this order: Local Preference (higher is preferred)
AS Path (shorter is preferred) Origin (IGP > EGP > incomplete)
MED (Multi-Exit Discriminator) (lower is preferred)
Juniper Reference:
BGP Path Selection: In Junos, the local preference attribute is the first to be evaluated when determining the best path for outbound traffic.

Question: 56

What is the default route preference of a static route in the Junos OS?

A. 0
B. 10
C. 1
D. 5

Answer: D
Explanation:

In Junos OS, the default route preference for a static route is 5. Route preference values are used to determine which route should be installed in the routing table when multiple routes to the same destination are available.

Step-by-Step Breakdown:
Static Route Preference:
A static route, by default, has a preference of 5, making it a highly preferred route. Lower preference values are more preferred in Junos, meaning static routes take precedence over most dynamic routing protocol routes, such as OSPF (preference 10) or BGP (preference 170).

Route Preference:
Route preference is a key factor in the Junos routing decision process. Routes with lower preference values are preferred and installed in the forwarding table.

Juniper Reference:
Static Routes: In Junos, the default preference for static routes is 5, making them more preferred than most dynamic routes.

Question: 57

You are creating an IP fabric underlay and want to use OSPF as your routing protocol. In this scenario, which statement is correct?

A. All leaf devices must be configured in separate OSPF areas.
B. All leaf and spine devices must be the same model to ensure the proper load-balancing behavior.
C. Interface speeds should be the same throughout the fabric to ensure that all links are utilized.
D. All spine devices must use the same router ID.

Answer: C
Explanation:

When creating an IP fabric underlay using OSPF as the routing protocol, consistent interface speeds are important to ensure optimal traffic distribution and utilization of all links.
Step-by-Step Breakdown:
OSPF and Interface Speeds:
OSPF calculates the cost of a link based on its bandwidth. The default cost calculation in OSPF is:

, Reference Bandwidth Cost = 7 — ;—. , ,
Interface Bandw idth

If interface speeds vary significantly, OSPF may choose paths with lower cost (higher bandwidth), resulting in some links being underutilized. Equal Utilization:

To ensure that all links are equally utilized in an IP fabric, it is recommended to maintain uniform interface speeds across the fabric. This ensures balanced load sharing across all available paths. Juniper Reference:
IP Fabric with OSPF: Juniper recommends consistent interface speeds to maintain even traffic distribution and optimal link utilization in IP fabric underlay designs.

Question: 58

What are three correct layer names used in legacy hierarchical network design? (Choose three.)

A. Access layer
B. Modular layer
C. Aggregation layer
D. Core layer
E. Function layer

Answer: A, C, D
Explanation:

In legacy hierarchical network design, three key layers are used to create a scalable and structured network:
Step-by-Step Breakdown:
Access Layer:
The access layer is where end devices, such as computers and IP phones, connect to the network. It typically involves switches that provide connectivity for devices at the edge of the network.
Aggregation Layer (Distribution Layer):
The aggregation layer (also called the distribution layer) aggregates traffic from multiple access layer devices and applies policies such as filtering and QoS. It also provides redundancy and load balancing.
Core Layer:
The core layer provides high-speed connectivity between aggregation layer devices and facilitates traffic within the data center or between different network segments.
Juniper Reference:
Legacy Hierarchical Design: Juniper networks often follow the traditional three-layer design (Access, Aggregation, and Core) to ensure scalability and high performance.

Question: 59

Which two statements are correct about aggregate routes and generated routes? (Choose two.)

A. An aggregate route does not have a forwarding next hop.
B. An aggregate route has a forwarding next hop.
C. A generated route has a forwarding next hop.

D. A generated route does not have a forwarding next hop.

Answer: A, C
Explanation:

Aggregate routes and generated routes are used to create summarized routes in Junos, but they behave differently in terms of forwarding.
Step-by-Step Breakdown:
Aggregate Routes:
An aggregate route summarizes a set of more specific routes, but it does not have a direct forwarding next hop. Instead, it points to the more specific routes for actual packet forwarding.
Generated Routes:

A generated route also summarizes specific routes, but it has a forwarding next hop that is determined based on the availability of contributing routes. The generated route can be used to directly forward traffic.
Juniper Reference:
Aggregate and Generated Routes: In Junos, aggregate routes rely on more specific routes for forwarding, while generated routes can forward traffic directly based on their next-hop information.

Question: 60

Which route is preferred by the Junos OS software routing tables?

A. Static
B. Aggregate
C. Direct
D. BGP

Answer: C
Explanation:

In Junos OS, direct routes are the most preferred routes in the routing table, having the highest priority.
Step-by-Step Breakdown:
Direct Routes:
Direct routes represent networks that are directly connected to the router's interfaces. Since these routes are directly accessible, they are assigned the highest priority and always take precedence over other types of routes.
Preference Values:
Direct routes have a preference of 0, which is the most preferred in Junos. Static routes, OSPF routes, and BGP routes have higher preference values and will only be used if there are no direct routes to the destination.
Juniper Reference:
Direct Route Preference: In Junos, direct routes are always preferred over other routes, ensuring that the router forwards traffic through locally connected networks.

Question: 61

Exhibit:

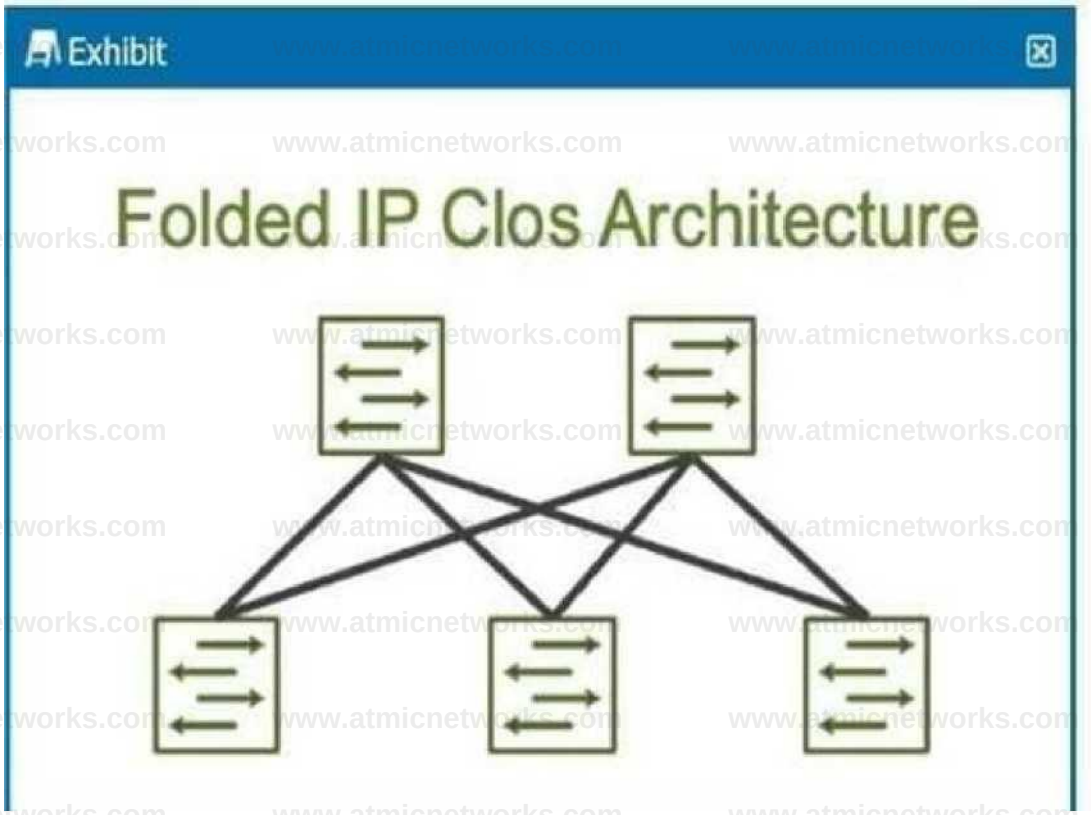

How many stages are shown in the exhibit?

A. 2
B. 5
C. 6
D. 3

Answer: D

Explanation:

The exhibit shows a Folded IP Clos Architecture, which is also referred to as a 3-stage Clos network design. This architecture typically consists of two layers of switches:

Spine Layer: The top row of switches. Leaf Layer: The bottom row of switches. Step-by-Step Breakdown:

Clos Architecture:

A 3-stage Clos network has two types of devices: spine and leaf. In this design, each leaf switch connects to every spine switch, providing a high level of redundancy and load balancing.

Stage

Stage 1: The first set of leaf switches. Stage 2: The spine switches.

Stage 3: The second set of leaf switches.

The Folded Clos architecture shown here effectively "folds" the 3-stage design by combining the ingress and egress leaf layers into one, reducing it to two visible layers, but still maintaining the overall 3-stage architecture.

Juniper Reference:

IP Clos Architecture: The 3-stage Clos design is commonly used in modern data centers for high availability, redundancy, and scalability.

Question: 62

Exhibit:

W Exhibit

```
]edit routing-options]
□ser@RouterI show static I
routs 3.0.0.0/0 (
    next-hop 172.25.11.254;
    qualified-next-hop 172.25.11.2*': ( preference 140;
```

Referring to the exhibit, what is the route preference of the 172.25.11.254 next hop?

A. 5
B. 10
C. 130
D. 140

Answer: A

Explanation:

In the exhibit, we see two next-hop addresses for the default static route (0.0.0.0/0): The first next hop is 172.25.11.254, with no specified preference.
The second next hop is 172.25.11.200, with a specified preference of 140. Step-by-Step Breakdown:
Default Static Route Preference:
If no preference is explicitly set for a next hop in Junos, it defaults to 5 for static routes.

Determining Preference:
In this case, the next hop 172.25.11.254 does not have an explicit preference defined, so it will use the default value of 5. The second next hop has a preference of 140, which is higher, meaning it will only be used if the primary next hop is unavailable.
Juniper Reference:
Static Route Preference: In Junos, the default preference for static routes is 5, and this value is applied unless overridden by the preference parameter.

Question: 63

When using spine and leaf fabric architectures, what is the role of each device? (Choose two.)

A. Spine nodes are used for host connectivity.
B. Spine nodes are used for transit to other leaf nodes.
C. Leaf nodes are used for traffic to other leafs.
D. Leaf nodes are used for host connectivity.

Answer: B, D Explanation:

In a spine-leaf fabric architecture, which is commonly used in data center designs, each device has a distinct role to ensure efficient and scalable network traffic flow.
Step-by-Step Breakdown:
Spine Nodes:
The spine nodes form the backbone of the fabric and are responsible for transit traffic between leaf nodes. They connect to every leaf switch and provide multiple paths for traffic between leaf nodes, ensuring redundancy and load balancing.
Leaf Nodes:
The leaf nodes are used for host connectivity. These switches connect to servers, storage, or edge routers. They also connect to the spine switches to reach other leaf switches.
Juniper Reference:
Spine-Leaf Architecture: In Juniper's IP fabric designs, spine switches handle inter-leaf communication, while leaf switches manage host and endpoint connectivity.

You are troubleshooting a downed BGP session.

Exhibit

```
jse:6ro'j:er> show tap summary match "AS IConnect I Active" Feer AS
InPict outPkt OutQ state|*Active/Received/Accepted/DaKped...        Flaps Last up
10.10.53.2 1111 0 00                                                10 l*0d 3:51:17 Connect
```

Question: 64

Referring to the exhibit, what is the cause of the problem?

A. The UDP session between the peers has not been established.
B. The local peer has sent an Open message but not received one from the remotepeer.

C. The TCP session between the peers has not been established.
D. The local peer has sent an Update message but not received one from the remote peer.

Answer: C

Explanation:

The BGP session in the exhibit shows the state as Connect, which indicates that the TCP session between the BGP peers has not been fully established.

Step-by-Step Breakdown:

BGP State "Connect":

The Connect state is the second stage in the BGP finite state machine (FSM). At this stage, BGP is trying to establish a TCP session with the peer, but the session has not yet been successfully established.

A successful TCP three-way handshake (SYN, SYN-ACK, ACK) is required before BGP can progress to the OpenSent state, where the peers exchange BGP Open messages.

Possible Causes:

A firewall blocking TCP port 179.

Incorrect IP addresses or network connectivity issues between the BGP peers.

Juniper Reference:

BGP Troubleshooting: In Junos, if a BGP session is stuck in the Connect state, the issue is likely due to a failure in establishing the underlying TCP connection.

Question: 65

What is the behavior of the default export policy for OSPF?

A. Accept all routes.
B. Reject all routes.
C. Redistribute all routes.
D. Forward all routes.

Answer: B

Explanation:

In Junos, the default export policy for OSPF is to reject all routes from being exported. Step-by-Step Breakdown:

Default Export Policy:

By default, OSPF in Junos does not export any routes to other routing protocols or neighbors. This is a safety mechanism to prevent unintended route advertisements.

Custom Export Policies:

If you need to export routes, you must create a custom export policy that explicitly defines which routes to advertise.

Example: You can create an export policy to redistribute static or connected routes into OSPF.

Juniper Reference:

OSPF Export Behavior: In Juniper devices, the default policy for OSPF is to reject route advertisements unless explicitly configured otherwise through custom policies.

Question: 66

Which of the following is a key characteristic of traditional (multitier) data center architectures?

A. Simple scalability and flexible deployment
B. Single-tier with direct server access
C. A hierarchical model with distinct layers (access, aggregation, and core)
D. Seamless integration with cloud environments

Answer: C
Explanation:

Traditional multitier data center architectures follow a hierarchical model, consisting of access, aggregation, and core layers. This structure provides scalability and segmentation, but it can be less flexible compared to modern designs.

Question: 67

In an IP fabric architecture with spine and leaf switches, which layer handles the forwarding of traffic within the fabric?

A. Leaf switches
B. Spine switches
C. Core switches
D. Access switches

Answer: A
Explanation:

In a spine/leaf architecture, leaf switches handle traffic forwarding within the fabric. The spine switches are responsible for connecting all leaf switches but do not perform direct traffic forwarding between end devices.

Question: 68

Which of the following is a major advantage of using an IP fabric architecture (spine/leaf) over traditional multitier designs?

A. Simplicity in design and configuration
B. Increased redundancy and lower latency
C. Better cost-efficiency for smaller environments
D. Less complexity in scaling

Answer: B
Explanation:

The spine/leaf architecture offers better redundancy and reduced latency due to its fully meshed topology. This design ensures that any path failure does not significantly impact traffic, providing more robust performance.

Question: 69

When considering Layer 2 strategies in a data center, which of the following is a common challenge?

A. Complex IP addressing and routing
B. Increased multicast traffic
C. Lack of redundancy
D. Difficult scaling and network segmentation

Answer: D

Explanation:

Scaling and network segmentation are significant challenges in Layer 2 strategies. As the Layer 2 domain grows, the complexity increases, particularly when dealing with broadcast storms and the need for larger broadcast domains.

Question: 70

Which of the following describes the key difference between overlay and underlay networks in a data center?

A. Overlay networks operate at the physical layer, underlay at the logical layer
B. Overlay networks use virtual encapsulation to abstract the physical network
C. Underlay networks are used for application-specific protocols
D. Overlay networks are entirely dependent on physical links

Answer: B

Explanation:

Overlay networks abstract the physical network by using virtual encapsulation technologies (such as VXLAN), while the underlay network refers to the physical infrastructure (routers, switches, etc.) that supports the overlay.

Question: 71

Which of the following technologies is used in data center architectures to extend Layer 2 connectivity across Layer 3 boundaries?

A. Ethernet VPN (EVPN)
B. Virtual Routing and Forwarding (VRF)
C. Spanning Tree Protocol (STP)
D. OpenFlow

Answer: A

Explanation:

Ethernet VPN (EVPN) is used in data centers to extend Layer 2 connectivity across Layer 3 boundaries. EVPN provides a more efficient and scalable solution for bridging Layer 2 segments in a multipath, distributed network.

Question: 72

What is the primary purpose of Virtual Extensible LAN (VXLAN) in modern data center networks?

A. To reduce the number of VLANs
B. To enable IP address aggregation
C. To simplify IP address management
D. To create Layer 2 virtual networks over a Layer 3 infrastructure

Answer: D
Explanation:

VXLAN is primarily used to create Layer 2 virtual networks over a Layer 3 infrastructure. This enables data center operators to scale their networks while maintaining Layer 2 isolation across Layer 3 boundaries.

Question: 73

Which of the following best describes the role of a spine switch in an IP fabric architecture?

A. Acts as a gateway between internal and external networks
B. Connects directly to end devices and servers
C. Routes traffic between leaf switches in the fabric
D. Provides high availability for access switches

Answer: C
Explanation:

Spine switches in an IP fabric architecture are responsible for routing traffic between leaf switches. They do not connect to end devices but facilitate communication across the fabric.

Question: 74

In a traditional multitier architecture, what is the main function of the core layer?

A. Handling access control and security policies
B. Connecting the data center to external networks
C. Managing routing between access and aggregation layers
D. Providing storage area network (SAN) connectivity

Answer: B
Explanation:

The core layer in a traditional multitier architecture primarily connects the data center to external networks, including internet and private WAN connections, ensuring high-speed routing between internal and external environments.

Question: 75

Which of the following best describes an advantage of using an underlay network in a data center?

A. Simplified configuration and management of network overlays
B. Faster IP address allocation
C. Support for the use of virtual network functions (VNFs)
D. Simplified troubleshooting due to a well-defined physical layer

Answer: D
Explanation:

An underlay network simplifies troubleshooting because it involves the physical network infrastructure, which is well-defined and less abstract compared to the overlay network that uses virtualized resources.

Question: 76

Which of the following is a primary function of the leaf switch in a spine/leaf architecture?

A. Connecting to the external internet gateway and managing external traffic flow
B. Forwarding traffic between different leaf switches within the same data center
C. Handling traffic from end devices and connecting to the spine switches
D. Managing security policies and enforcing access controls across the entire network

Answer: C
Explanation:

Leaf switches connect directly to end devices and facilitate communication with spine switches. They manage the traffic within the fabric and rely on spine switches to forward it across the network.

Question: 77

What role does Ethernet VPN (EVPN) play in data center architectures?

A. Enables the transport of Layer 2 frames over a Layer 3 network
B. Provides an alternative to VXLAN for network tunneling
C. Simplifies IP routing and network management in large-scale environments
D. Acts as a security mechanism for virtualized data center environments

Answer: A
Explanation:

EVPN is designed to transport Layer 2 Ethernet frames over a Layer 3 network, extending Layer 2 connectivity across distributed networks and providing an efficient way to bridge network segments.

Question: 78

Which of the following is the main benefit of using a spine/leaf architecture in data centers?

A. Simple integration with WAN networks

B. Increased scalability and lower latency
C. Easier configuration and management of VLANs
D. Enhanced security due to isolation of network segments

Answer: B
Explanation:

Spine/leaf architecture provides increased scalability and lower latency due to its fully meshed topology. The design ensures that data is routed efficiently with minimal delay, even as the network grows.

Question: 79

How does VXLAN differ from traditional VLANs in data center architectures?

A. VXLAN requires no network switches
B. VXLAN only works in Layer 1 networks
C. VXLAN is limited to single data center environments
D. VXLAN supports more than 4096 unique networks

Answer: D
Explanation:

VXLAN supports over 16 million unique identifiers (VXLAN IDs), which is a significant increase from the 4096 limit imposed by traditional VLANs. This allows for greater flexibility and scalability in large data centers.

Question: 80

In a data center that implements EVPN, what is the primary advantage of using this technology?

A. It eliminates the need for network switches
B. It provides secure encryption of all traffic
C. It simplifies the implementation of multi-tenant environments
D. It improves routing efficiency by using software-defined networking (SDN)

Answer: C
Explanation:

EVPN simplifies the implementation of multi-tenant environments by providing efficient Layer 2 and Layer 3 connectivity between tenants, offering improved scalability and redundancy.

Question: 81

What is the primary purpose of an underlay network in a spine/leaf architecture?

A. To define and manage the VLANs for each network segment
B. To provide reliable physical connectivity and transport for overlay networks
C. To manage and enforce security policies across the entire network infrastructure
D. To enable and support the implementation of software-defined networking (SDN) solutions

Answer: B
Explanation:

The underlay network provides the physical infrastructure that supports the overlay networks, which are responsible for virtual network configurations and traffic forwarding across the data center.

Question: 82

What is one of the challenges of using traditional multitier architectures in modern data centers?

A. Increased network latency due to multiple tiers of switches
B. Lack of scalability for modern workloads
C. Complex management and configuration of virtualized network environments
D. Inefficient use of VLANs

Answer: A
Explanation:

Traditional multitier architectures can introduce increased network latency due to the multiple tiers (access, aggregation, and core), which can slow down communication as traffic must traverse through each layer.

Question: 83

What is the function of a spine switch in a data center's IP fabric architecture?

A. To act as a point of failure for high availability
B. To connect all leaf switches and provide inter-switch connectivity
C. To provide security by enforcing access control policies
D. To manage traffic from external sources into the data center

Answer: B
Explanation:

Spine switches connect all leaf switches in the fabric, enabling efficient communication between them. They ensure that data can flow seamlessly across the network without bottlenecks.

Question: 84

Which Layer 2 protocol is typically used in data centers to prevent loops in the network?

A. Virtual Router Redundancy Protocol (VRRP)
B. Open Shortest Path First (OSPF)
C. Border Gateway Protocol (BGP)
D. Spanning Tree Protocol (STP)

Answer: D
Explanation:

Spanning Tree Protocol (STP) is widely used in data centers to prevent network loops by creating a loop-free logical topology. It ensures that only one path exists between switches, even in redundant configurations.

Question: 85

In an Ethernet VPN (EVPN) implementation, what is the role of the Type 5 EVPN route?

A. To advertise the IP address of an endpoint
B. To advertise the MAC address of an endpoint
C. To provide a method for VXLAN-to-VXLAN tunneling
D. To carry Layer 3 IP prefixes for routing across the fabric

Answer: D
Explanation:

The Type 5 EVPN route is used to advertise Layer 3 IP prefixes across the EVPN fabric. This allows for both Layer 2 and Layer 3 connectivity to be extended across the data center network.

Question: 86

Which Layer 2 protocol is typically used in data centers to prevent loops in the network?

A. Spanning Tree Protocol (STP)
B. Open Shortest Path First (OSPF)
C. Border Gateway Protocol (BGP)
D. Virtual Router Redundancy Protocol (VRRP)

Answer: A
Explanation:

Spanning Tree Protocol (STP) is widely used in data centers to prevent network loops by creating a loop-free logical topology. It ensures that only one path exists between switches, even in redundant configurations.

Question: 87

In an Ethernet VPN (EVPN) implementation, what is the role of the Type 5 EVPN route?

A. To advertise the IP address of an endpoint in the network

B. To advertise the MAC address of an endpoint for proper traffic forwarding
C. To carry Layer 3 IP prefixes for routing across the fabric
D. To provide a method for VXLAN-to-VXLAN tunneling

Answer: C
Explanation:

The Type 5 EVPN route is used to advertise Layer 3 IP prefixes across the EVPN fabric. This allows for both Layer 2 and Layer 3 connectivity to be extended across the data center network.

Question: 88

Which architecture is more suitable for environments where rapid scaling and low latency are critical?

A. Traditional multitier architecture
B. IP fabric (spine/leaf) architecture
C. Hybrid architecture with both Layer 2 and Layer 3
D. Single-tier architecture

Answer: B
Explanation:

IP fabric (spine/leaf) architecture is designed to provide rapid scaling and low latency by ensuring that all leaf switches are directly connected to spine switches, allowing for efficient and resilient data flow.

Question: 89

In a Layer 3 strategy, what is the primary function of routing?

A. Maintaining ARP tables to map IP addresses to MAC addresses
B. Segmenting broadcast domains and limiting broadcast traffic
C. Controlling traffic flow within a VLAN
D. Directing traffic between different IP subnets

Answer: D
Explanation:

In a Layer 3 strategy, routing's primary function is to direct traffic between different IP subnets by using routing protocols and making decisions based on IP addresses.

Question: 90

What is the primary benefit of using Overlay networks in modern data center designs?

A. Provides separation between virtual networks and physical infrastructure
B. Simplifies the management and scalability of VLAN configurations
C. Reduces the need for complex high availability configurations in the network

D. Increases the complexity of routing decisions across the network architecture

Answer: A
Explanation:

Overlay networks provide separation between virtual networks and physical infrastructure, which allows for greater flexibility, scalability, and isolation of virtualized workloads across the data center.

Question: 91

Which of the following best describes a key advantage of using EVPN-VXLAN in data center

interconnects?

A. It simplifies IP routing between data centers
B. It extends Layer 2 connectivity over Layer 3 networks while supporting multi-tenancy
C. It eliminates the need for VLANs in the data center
D. It requires no configuration of spine and leaf switches

Answer: B
Explanation:

EVPN-VXLAN allows for the extension of Layer 2 connectivity over Layer 3 networks while supporting multi-tenancy, offering efficient and scalable data center interconnect solutions.

Question: 92

Which statement describes the primary function of a spine switch in a data center IP fabric?

A. Provides Layer 3 routing between leaf switches for efficient traffic management
B. Connects to end-user devices, enabling access to the network
C. Aggregates data from all devices across the network to ensure high availability
D. Facilitates traffic forwarding between leaf switches without being a bottleneck

Answer: D
Explanation:

Spine switches facilitate traffic forwarding between leaf switches without being a bottleneck in the system. They provide high-speed connectivity across the fabric, ensuring low-latency and fault-tolerant communication.

Question: 93

Which is a key characteristic of IP fabric (spine/leaf) architectures in terms of scalability?

A. It supports vertical scaling with additional bandwidth
B. It scales horizontally by adding additional leaf switches
C. It requires a hierarchical approach with core, distribution, and access layers
D. It is not suitable for large data center environments

Answer: B
Explanation:

IP fabric (spine/leaf) architectures scale horizontally by adding more leaf switches. This design enables flexible scalability while maintaining consistent performance across the network.

Question: 94

How does VXLAN support large-scale data center networks?

A. It allows Layer 3 routing over Layer 2 networks
B. It limits broadcast traffic to a single data center

C. It provides Layer 2 isolation across a Layer 3 infrastructure
D. It integrates directly with legacy network protocols

Answer: C
Explanation:

VXLAN provides Layer 2 isolation over a Layer 3 network, allowing for better segmentation and scalability. It extends Layer 2 domains across different data centers while leveraging the underlay's Layer 3 infrastructure.

Question: 95

What is the primary benefit of using EVPN to extend Layer 2 connectivity in data centers?

A. It increases the overall bandwidth of the network for better performance
B. It enables virtual routing and switching to support advanced network functions
C. It offers efficient Layer 2 extensions over Layer 3, reducing broadcast traffic
D. It simplifies Layer 3 routing across multi-site deployments

Answer: C
Explanation:

EVPN provides efficient Layer 2 extensions over a Layer 3 network, reducing the need for excessive broadcast traffic and enhancing scalability and redundancy in large data centers.

Question: 96

Which of the following is a key advantage of using an underlay network in modern data center designs?

A. Simplifies the deployment of Layer 3 VPNs
B. Provides the physical network infrastructure for traffic forwarding
C. Eliminates the need for overlay networks by handling all traffic
D. Ensures full visibility and control over Layer 2 broadcast domains

Answer: B
Explanation:

The underlay network provides the physical infrastructure necessary to forward traffic between network devices. It is the foundation for overlay networks, which rely on the underlay for efficient data transport.

Question: 97

What does the term "spine/leaf" refer to in an IP fabric network?

A. A network design where spine switches provide connectivity between leaf switches
B. The connection between Layer 2 devices and Layer 3 routers within the network
C. The layer in the network responsible for IP routing and inter-switch communication
D. A topology that isolates network traffic between servers, storage systems, and other devices

Answer: A

Explanation:

"Spine/leaf" refers to the architecture where spine switches provide connectivity between leaf switches. This design is used in data centers to ensure redundancy, low latency, and easy scalability.

Question: 98

In the context of EVPN, what does the Type 2 route advertise?

A. MAC address to IP address mapping
B. The BGP next-hop IP address for route propagation
C. Layer 3 IP prefixes for routing across the network fabric
D. Layer 2 MAC addresses for virtual network segments

Answer: D
Explanation:

EVPN Type 2 routes advertise Layer 2 MAC addresses for virtual network segments, allowing for efficient Layer 2 forwarding across the network and providing end-to-end connectivity in large-scale virtualized environments.

Question: 99

What is one of the primary benefits of using a spine/leaf network architecture over traditional Layer 2 designs?

A. Simpler configuration
B. Higher performance and scalability
C. Better integration with external networks
D. Reduced complexity in handling broadcast traffic

Answer: B
Explanation:

Spine/leaf architectures provide higher performance and scalability by creating a fully meshed network where leaf switches are connected to spine switches, allowing for efficient traffic flow without bottlenecks.

Question: 100

What does the Layer 3 strategy in a data center typically involve?

A. The use of VLANs to segment traffic
B. Relying on broadcast domains for traffic isolation
C. Direct routing between subnets to optimize traffic flow
D. Ensuring secure communication through firewalls

Answer: C
Explanation:

A Layer 3 strategy focuses on routing between subnets, optimizing traffic flow by managing IP addresses and enabling efficient data transmission between different network segments.

Question: 101

Which type of network encapsulation is used in VXLAN to extend Layer 2 networks across Layer 3 boundaries?

A. IP encapsulation
B. MPLS encapsulation
C. Ethernet encapsulation
D. UDP encapsulation

Answer: D
Explanation:

VXLAN uses UDP encapsulation to extend Layer 2 networks over Layer 3 boundaries. This allows for the isolation of virtual networks over a shared physical infrastructure while enabling better scalability.

Question: 102

Which of the following describes the role of a leaf switch in a spine/leaf architecture?

A. Connects end devices and routes traffic to other leaf switches
B. Routes external traffic to and from the data center
C. Acts as the point of failure in the fabric
D. Only forwards traffic between spine switches

Answer: A
Explanation:

Leaf switches connect to end devices and facilitate communication to and from other leaf switches. They also interact with spine switches for forwarding traffic across the fabric.

Question: 103

In a data center network utilizing EVPN, which function is provided by the Type 1 EVPN route?

A. A) Advertising Layer 2 MAC addresses for endpoint identification
B. Mapping Layer 3 IP addresses to MAC addresses
C. Enabling inter-VLAN routing for communication across VLANs
D. Advertising IP prefixes for routing to external networks

Answer: B
Explanation:

The Type 1 EVPN route is used for mapping Layer 3 IP addresses to MAC addresses, which enables the network to efficiently direct traffic to the correct endpoint.

Question: 104

In a Layer 2 design, which technology is typically used to prevent broadcast storms in a network?

A. OSPF
B. Spanning Tree Protocol (STP)
C. Virtual Routing and Forwarding (VRF)
D. Border Gateway Protocol (BGP)

Answer: B
Explanation:

Spanning Tree Protocol (STP) is used in Layer 2 designs to prevent broadcast storms by blocking redundant paths and ensuring that only one active path exists between devices. This prevents loops and excessive broadcast traffic.

Question: 105

Which Layer 2 technology is used to extend the Ethernet segment over an IP network in a data center?

A. EVPN
B. MPLS
C. VXLAN
D. OSPF

Answer: C

Explanation:

VXLAN (Virtual Extensible LAN) is used to extend Ethernet segments over an IP network, allowing Layer 2 networks to span Layer 3 boundaries. This enables better scalability and flexibility in data center environments.

Question: 106

In a spine/leaf architecture, what is the primary role of the spine switches?

A. To handle traffic from end devices
B. To connect leaf switches and facilitate inter-switch communication
C. To provide access to external networks
D. To control routing between subnets

Answer: B
Explanation:

In a spine/leaf architecture, spine switches connect all leaf switches and facilitate inter-switch communication. They do not directly interact with end devices but are critical for efficient data forwarding across the fabric.

Question: 107

Which of the following is a key benefit of using Overlay networks in data center architectures?

A. Isolation of virtualized networks from the underlying physical infrastructure
B. Simplification of IP address allocation and management across virtual networks
C. Improved management of broadcast traffic and reduced network congestion
D. Easier configuration and management of physical network devices in the data center

Answer: A
Explanation:

Overlay networks provide isolation of virtualized networks from the physical infrastructure, allowing for greater flexibility in network design and scalability. This abstraction helps to manage traffic and resources independently of the underlying physical hardware.

Question: 108

Which of the following are advantages of using a spine/leaf architecture in data centers? (Choose two)

A. Increased scalability and reduced latency
B. Easier integration with external WAN environments
C. Better fault tolerance and redundancy
D. Simplified Layer 2 management across large networks

Answer: A, C
Explanation:

Spine/leaf architecture offers increased scalability and reduced latency by ensuring direct paths between leaf switches via the spine. Additionally, it provides better fault tolerance and redundancy because the fully meshed design allows multiple paths between devices.

Question: 109

Which of the following are key features of Ethernet VPN (EVPN)? (Choose two)

A. It integrates with MPLS to create a hybrid Layer 2/Layer 3 network
B. It eliminates the need for using VXLAN for data center interconnects
C. It allows for the extension of Layer 2 Ethernet segments across a Layer 3 network
D. It provides an alternative to MPLS for Layer 2 and Layer 3 services

Answer: C, D
Explanation:

EVPN provides an alternative to MPLS by offering both Layer 2 and Layer 3 connectivity over a Layer 3 network. It enables the extension of Layer 2 Ethernet segments across a Layer 3 infrastructure, which is useful in data center interconnects.

Question: 110

Which of the following are characteristics of Layer 3 strategies in data center architectures? (Choose two)

A. Traffic is routed between different subnets
B. It involves the use of routing protocols for communication between subnets
C. Broadcast traffic is handled within VLANs
D. It simplifies IP address management for large-scale networks

Answer: A, B
Explanation:

In Layer 3 strategies, traffic is routed between different subnets using routing protocols. This approach allows for efficient traffic management across multiple subnets and is essential for large-scale data center environments.

Question: 111

Which of the following are common use cases for VXLAN in data centers? (Choose two)

A. Reducing the complexity of IP routing
B. Extending Layer 2 networks across Layer 3 boundaries
C. Segmenting traffic within a single Layer 3 network
D. Enabling multi-tenant data center environments

Answer: B, D
Explanation:

VXLAN is primarily used to extend Layer 2 networks over Layer 3 boundaries, making it suitable for multi-tenant data center environments. It allows different tenants to maintain isolated networks while using the same physical infrastructure.

Question: 112

Which of the following are benefits of using Overlay networks in modern data centers? (Choose two)

A. They simplify physical network configurations
B. They provide virtual network isolation from physical infrastructure
C. They reduce the complexity of the underlay network
D. They increase the capacity for Layer 2 traffic

Answer: B, C
Explanation:

Overlay networks provide virtual network isolation, allowing for greater flexibility in network design. Additionally, they reduce the complexity of the underlay network by abstracting the physical network infrastructure from the virtualized network layer.

Question: 113

Which of the following are key differences between traditional multitier architectures and IP fabric (spine/leaf) architectures? (Choose two)

A. IP fabric architectures support faster scaling by adding more leaf switches
B. Traditional multitier architectures are typically more scalable than IP fabric architectures
C. IP fabric architectures offer better redundancy and lower latency due to their meshed design
D. Traditional multitier architectures have a flat topology that simplifies routing

Answer: A, C
Explanation:

IP fabric architectures offer faster scaling by adding more leaf switches without significant design changes. They also provide better redundancy and lower latency because of the fully meshed spine/leaf topology, ensuring efficient data forwarding.

Question: 114

Which of the following are true regarding the role of spine switches in an IP fabric architecture? (Choose two)

A. Spine switches handle traffic forwarding between leaf switches
B. Spine switches connect directly to end-user devices
C. Spine switches provide high availability and fault tolerance across the network
D. Spine switches perform Layer 3 routing for external connections

Answer: A, C
Explanation:

Spine switches in an IP fabric architecture are responsible for forwarding traffic between leaf switches. They provide high availability and fault tolerance by offering multiple paths for data to flow, ensuring continuous connectivity even in case of failures.

Question: 115

What is the purpose of Ethernet switching in a Layer 2 network?

A. To route traffic between different IP subnets
B. To forward traffic based on MAC addresses
C. To manage IP address assignments
D. To segment the network into multiple broadcast domains

Answer: B
Explanation:

Ethernet switching forwards traffic based on MAC addresses at Layer 2. It ensures that frames are delivered to the correct destination device within a local network, eliminating the need for IP routing.

Question: 116

Which of the following statements about VLAN tagging is true?

A. VLAN tags are added to frames at Layer 3
B. VLAN tags are only applied to multicast traffic
C. VLAN tagging is used only for routing purposes
D. VLAN tags are used to distinguish between different Layer 2 broadcast domains

Answer: D
Explanation:

VLAN tags are added to Ethernet frames to distinguish traffic between different VLANs at Layer 2, creating isolated broadcast domains within a Layer 2 network.

Question: 117

What is the primary benefit of using VLANs in a network?

A. To extend the range of IP addresses within the network

B. To enhance the performance of Layer 3 routing across the network
C. To logically segment the network and reduce broadcast traffic
D. To enable faster Ethernet switching

Answer: C
Explanation:

VLANs allow logical segmentation of the network, reducing broadcast traffic and improving network efficiency by isolating traffic within each VLAN. This enhances security and scalability.

Question: 118

Which port mode in VLAN configuration is used when a port is intended to carry traffic for multiple VLANs?

A. Trunk
B. Hybrid
C. Access
D. Management

Answer: A
Explanation:

A trunk port is used to carry traffic for multiple VLANs, allowing the transmission of VLAN-tagged frames between switches. Trunk ports support multiple VLANs through tagging.

Question: 119

What is the function of Integrated Routing and Bridging (IRB) in a Layer 2 network?

A. To isolate broadcast traffic from the VLANs
B. To segregate different types of traffic in the same VLAN
C. To prevent loops in the Layer 2 network
D. To enable Layer 3 routing for VLAN traffic within the same switch

Answer: D
Explanation:

IRB allows routing between VLANs while maintaining Layer 2 bridging functionality. It enables communication between VLANs using Layer 3 routing without needing external routers.

Question: 120

Which of the following is an example of a Layer 2 broadcast frame?

A. IP packet
B. ARP request
C. MAC address lookup frame
D. ICMP Echo request

Answer: B
Explanation:

An ARP request is a Layer 2 broadcast frame, as it is sent to all devices on a network segment to resolve an IP address to a MAC address. This allows the sender to discover the MAC address of a device.

Question: 121

Which protocol is used by switches to prevent loops in a Layer 2 network?

A. Routing Information Protocol (RIP)
B. Border Gateway Protocol (BGP)
C. Open Shortest Path First (OSPF)
D. Spanning Tree Protocol (STP)

Answer: D
Explanation:

Spanning Tree Protocol (STP) is used in Layer 2 networks to prevent loops by blocking redundant paths and ensuring there is only one active path for data to flow between switches.

Question: 122

What happens when a switch receives a frame with an unknown destination MAC address?

A. The switch drops the frame
B. The switch floods the frame to all ports in the VLAN
C. The switch sends a request to the router for the MAC address
D. The switch forwards the frame to the default gateway

Answer: B
Explanation:

When a switch receives a frame with an unknown destination MAC address, it floods the frame to all ports within the VLAN (except the incoming port) to ensure it reaches the correct destination.

Question: 123

Which command is used to configure a switch port as a trunk port on a Junos OS switch?

A. set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members all
B. set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan trunk
C. set interfaces ge-0/0/1 unit 0 family ethernet-switching port-mode trunk

D. set interfaces ge-0/0/1 unit 0 family ethernet-switching mode trunk

Answer: C
Explanation:

The command set interfaces ge-0/0/1 unit 0 family ethernet-switching port-mode trunk is used to configure a port as a trunk, which allows it to carry traffic for multiple VLANs.

Question: 124

Which of the following is a key function of VLANs in network security?

A. VLANs allow for the creation of secure VPN tunnels
B. VLANs provide encryption for sensitive traffic
C. VLANs isolate traffic between different network segments, improving security
D. VLANs allow for the creation of external firewall rules

Answer: C
Explanation:

VLANs improve network security by isolating traffic between different segments, ensuring that only devices within the same VLAN can communicate directly, preventing unauthorized access between different segments.

Question: 125

Which type of VLAN configuration is used when you want to assign a port to a single VLAN?

A. Access
B. Trunk
C. Hybrid
D. Management

Answer: A
Explanation:

Access VLANs are configured for ports assigned to a single VLAN. These ports handle traffic for one VLAN and do not carry traffic for multiple VLANs like trunk ports.

Question: 126

Which of the following is the purpose of VLAN tagging in a network?

A. To reduce the number of broadcast domains
B. To segregate traffic into different logical networks over the same physical link
C. To define the layer 3 routing for a VLAN
D. To manage IP address assignments for different VLANs

Answer: B
Explanation:

VLAN tagging allows traffic to be segregated into different logical networks (VLANs) over the same physical link. The tag in the frame ensures that the receiving device knows which VLAN the frame belongs to.

Question: 127

Which of the following statements is true about the operation of the Spanning Tree Protocol (STP)?

A. STP blocks all redundant links to prevent network congestion
B. STP disables the use of VLANs in a network
C. STP sends broadcasts to find MAC addresses
D. STP assigns a unique path cost for each port in the network

Answer: D
Explanation:

STP assigns a unique path cost to each port to determine the best path for forwarding traffic. It uses this cost to block redundant links that could cause loops in the network.

Question: 128

In a Layer 2 network, what is the primary purpose of a MAC address table?

A. To store routing information for Layer 3 networks
B. To store the IP address of each device in the network
C. To map MAC addresses to ports to forward frames to the correct destination
D. To manage VLAN configurations

Answer: C
Explanation:

A MAC address table maps each MAC address to a specific port on the switch. This allows the switch to forward frames to the correct destination port based on the MAC address in the frame header.

Question: 129

How does the "native VLAN" feature work in a trunk port configuration?

A. It is the VLAN that is used for untagged traffic on the trunk port
B. It is the VLAN used for management traffic
C. It is the VLAN that is used to route traffic between VLANs
D. It is used to isolate traffic from external networks

Answer: A
Explanation:

The native VLAN is the VLAN used for untagged traffic on a trunk port. When a frame does not have a VLAN tag, it is assumed to belong to the native VLAN.

Question: 130

Which Layer 2 protocol is used to prevent loops in a network by dynamically blocking redundant paths?

A. ARP (Address Resolution Protocol)
B. OSPF (Open Shortest Path First)
C. STP (Spanning Tree Protocol)
D. VRRP (Virtual Router Redundancy Protocol)

Answer: C
Explanation:

Spanning Tree Protocol (STP) is used to prevent loops in a network by dynamically blocking redundant paths and ensuring there is only one active path for data to flow through the network.

Question: 131

What is the main difference between an "access" port and a "trunk" port in VLAN configuration?

A. Access ports carry traffic for only one VLAN, while trunk ports carry traffic for multiple VLANs
B. Access ports require manual configuration, while trunk ports are automatically configured
C. Access ports only transmit untagged traffic, while trunk ports transmit both tagged and untagged traffic
D. There is no difference; both port types function the same

Answer: A
Explanation:

Access ports carry traffic for only one VLAN, while trunk ports are designed to carry traffic for multiple VLANs, supporting VLAN tagging to distinguish the traffic.

Question: 132

In a Layer 2 network, what is the function of a MAC address table?

A. To map IP addresses to physical MAC addresses
B. To map MAC addresses to specific ports on a switch
C. To forward traffic to the appropriate router
D. To prevent network loops by blocking redundant paths

Answer: B
Explanation:

The MAC address table in a switch maps MAC addresses to specific ports, allowing the switch to forward frames to the correct destination port based on the MAC address of the frame's destination.

Question: 133

What happens if a switch receives an Ethernet frame with an unrecognized or unknown destination MAC address?

A. The switch drops the frame
B. The switch forwards the frame to the default gateway
C. The switch floods the frame to all ports in the VLAN
D. The switch sends an ARP request to find the MAC address

Answer: C
Explanation:

If a switch receives a frame with an unrecognized or unknown destination MAC address, it floods the frame to all ports within the VLAN (except the port on which the frame was received), to ensure the frame reaches its destination.

Question: 134

Which of the following describes the purpose of VLAN tagging in Ethernet frames?

A. To mark a frame as part of a multicast group
B. To assign an IP address to a frame
C. To encrypt frames to prevent unauthorized access
D. To identify the VLAN to which a frame belongs for forwarding

Answer: D
Explanation:

VLAN tagging allows frames to be marked with a VLAN ID so that switches can forward them to the correct VLAN. This ensures that traffic from different VLANs is properly segregated even when using the same physical link.

Question: 135

In Junos OS, which command would you use to configure a port as an access port for VLAN 100?

A. set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members 100
B. set interfaces ge-0/0/1 unit 0 family ethernet-switching port-mode trunk
C. set vlans vlan100 interface ge-0/0/1
D. set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan trunk members 100

Answer: A
Explanation:

The correct command to configure a port as an access port for VLAN 100 is set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members 100, which assigns the port to the specified VLAN.

Question: 136

What is the primary role of the Spanning Tree Protocol (STP) in a Layer 2 network?

A. To encrypt traffic between switches
B. To filter traffic based on MAC addresses
C. To enable inter-VLAN routing between different VLANs
D. To prevent loops by blocking redundant paths

Answer: D
Explanation:

The primary role of Spanning Tree Protocol (STP) is to prevent loops in Layer 2 networks by blocking redundant paths. It ensures that only one active path exists between any two devices in the network.

Question: 137

Which of the following is an advantage of using VLANs in a network?

A. They allow Layer 2 segmentation without needing multiple physical switches
B. They enable multiple IP addresses to be assigned to a single device
C. They increase the complexity of routing protocols
D. They prevent all forms of network traffic

Answer: A
Explanation:

VLANs allow for logical segmentation of a network into multiple broadcast domains without needing multiple physical switches. This reduces broadcast traffic and improves network efficiency.

Question: 138

When configuring VLANs on a switch, what type of port is typically used to connect to end devices, such as computers or printers?

A. Trunk port
B. Access port
C. Hybrid port
D. Management port

Answer: B
Explanation:

Access ports are used to connect end devices to the network. These ports carry traffic for only one

VLAN, unlike trunk ports, which carry traffic for multiple VLANs.

Question: 139

Which command is used in Junos OS to create a VLAN?

A. set vlans vlan100 vlan-id 100
B. set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members 100
C. set vlan 100 create
D. set vlans vlan100 interface ge-0/0/1

Answer: A
Explanation:

To create a VLAN in Junos OS, you use the set vlans vlan100 vlan-id 100 command, where vlan100 is the name of the VLAN, and 100 is the VLAN ID.

Question: 140

What is the function of a trunk port in a network?

A. It carries traffic for a single VLAN
B. It connects two devices in the same VLAN
C. It carries traffic for multiple VLANs using VLAN tags
D. It allows untagged traffic to pass through

Answer: C
Explanation:

Trunk ports are used to carry traffic for multiple VLANs between switches. They use VLAN tags to differentiate the traffic of each VLAN and ensure proper forwarding.

Question: 141

Which of the following is an appropriate use of the Integrated Routing and Bridging (IRB) feature?

A. To map MAC addresses to IP addresses within the network
B. To allow trunk ports to carry VLAN traffic over Layer 3 networks
C. To isolate VLAN traffic from external networks for security
D. To enable Layer 3 routing between different VLANs while maintaining Layer 2 bridging

Answer: D
Explanation:

IRB allows Layer 3 routing between different VLANs while still maintaining Layer 2 bridging within each VLAN. This feature is commonly used in networks that need to route between VLANs without using an external router.

Question: 142

In a network utilizing VLANs, what happens to traffic that is not tagged with a VLAN ID?

A. It is dropped by all switches
B. It is treated as multicast traffic
C. It is automatically forwarded to the default gateway
D. It is automatically assigned to the native VLAN

Answer: D
Explanation:

Traffic that is not tagged with a VLAN ID is assigned to the native VLAN by default, which is typically the VLAN that carries untagged traffic on trunk ports.

Question: 143

Which of the following are functions of VLANs in a Layer 2 network? (Choose two)

A. Isolating broadcast domains
B. Routing traffic between different VLANs
C. Reducing network congestion by creating multiple logical networks
D. Increasing the number of available IP addresses in the network

Answer: A, C **Explanation:**

VLANs help isolate broadcast domains, which reduces broadcast traffic within each VLAN, and they segment the network into multiple logical networks, improving network performance and scalability.

Question: 144

Which of the following are advantages of using VLAN tagging in a trunk port? (Choose two)

A. It enables the transfer of untagged traffic across VLANs
B. It increases the bandwidth of the physical link
C. It isolates traffic between VLANs on a trunk link
D. It allows a trunk port to carry traffic for multiple VLANs

Answer: C, D Explanation:

VLAN tagging on trunk ports enables the transmission of traffic from multiple VLANs across a single physical link and ensures that each frame is correctly identified and forwarded to the appropriate VLAN, maintaining isolation between VLANs.

Question: 145

Which of the following are benefits of using Spanning Tree Protocol (STP) in a Layer 2 network? (Choose two)

A. It prevents broadcast storms by blocking redundant paths
B. It creates multiple forwarding paths to increase bandwidth
C. It eliminates the need for IP routing between VLANs
D. It prevents network loops by dynamically blocking redundant paths

Answer: A, D
Explanation:

STP prevents network loops by dynamically blocking redundant paths and ensures there is only one active path between switches in a network, effectively preventing broadcast storms and ensuring stable network operation.

Question: 146

Which of the following are correct methods for configuring VLANs on a Junos OS switch? (Choose two)

A. Using the set vlans command to assign an IP address to the VLAN
B. Using the set interfaces command to assign a port to a specific VLAN
C. Using the show interfaces command to configure VLANs
D. Using the set interfaces command to configure the VLAN ID for the interface

Answer: B, D
Explanation:

The set interfaces command is used to assign ports to specific VLANs and configure the VLAN ID on the interface. The set vlans command is used to define and configure VLAN settings but not to assign IP addresses directly.

Question: 147

Which of the following are features of Integrated Routing and Bridging (IRB) in Junos OS? (Choose two)

A. IRB enables Layer 3 routing between VLANs without external routers B. IRB allows routing between different VLANs within the same network C. IRB bridges VLAN traffic on the same physical interface D. IRB disables VLAN tagging for inter-VLAN traffic

Answer: A, B Explanation:

IRB allows Layer 3 routing between VLANs within the same network while maintaining Layer 2 bridging within each VLAN. It eliminates the need for external routers for inter-VLAN communication.

Question: 148

Which of the following are valid port modes in VLAN configuration? (Choose two)

A. Access
B. Trunk
C. Hybrid
D. Isolation

Answer: A, B Explanation:

Access ports are used to carry traffic for a single VLAN, while trunk ports carry traffic for multiple VLANs. Hybrid is not a common standard port mode, and Isolation is not typically used to define port modes in VLAN configurations.

Question: 149

Which of the following are typical troubleshooting steps when VLANs are not communicating across a trunk link? (Choose two)

A. Check if the switches are running different versions of Spanning Tree Protocol (STP)
B. Verify that both switches have the same VLAN database
C. Ensure that both ends of the trunk link are configured for the same native VLAN
D. Confirm that the router is configured to route traffic between VLANs

Answer: B, C
Explanation:

When troubleshooting VLAN communication issues across trunk links, you should ensure that both ends of the trunk link are configured with the same native VLAN and verify that the VLAN database is synchronized across the switches.

Question: 150

Which of the following are valid reasons to configure a port in access mode for a specific VLAN? (Choose two)

A. To prevent devices from sending VLAN-tagged frames
B. To enable multiple VLANs on a single port
C. To connect a device that should only communicate within a single VLAN

D. To isolate a device from other VLANs for security purposes

Answer: C, D

Explanation:

Access mode is used to connect devices that belong to a single VLAN, preventing them from sending VLAN-tagged frames. It is also useful for isolating devices for security purposes by placing them in their dedicated VLAN.

Question: 151

Which of the following statements is true about static routes in a protocol-independent routing configuration?

A. Static routes are dynamically learned and updated
B. Static routes require no manual configuration
C. Static routes are manually configured and do not change unless modified by the administrator
D. Static routes can only be used in Layer 2 networks

Answer: C

Explanation:

Static routes are manually configured and remain unchanged unless explicitly modified by the network administrator. They do not adapt automatically to network changes, unlike dynamic routing protocols.

Question: 152

What is the primary purpose of an aggregate route in protocol-independent routing?

A. To exclude specific routes from the routing table
B. To divide a large network into smaller subnets
C. To automatically assign routes based on network size
D. To summarize multiple routes into a single route

Answer: D

Explanation:

Aggregate routes are used to summarize multiple smaller routes into a single larger route, reducing the size of the routing table and improving routing efficiency.

Question: 153

Which of the following best defines a Martian address?

A. A private IP address reserved for use within local networks and not routable on the public Internet
B. An address that is used exclusively for broadcast communication within a local network
C. An address that is not routable on the public Internet, typically falling within reserved IP ranges
D. A special address used to route traffic between multiple routing tables

Answer: C

Explanation:

Martian addresses are IP addresses that fall within reserved address ranges and are not routable on the public Internet. These are usually detected as invalid or unreachable when routed.

Question: 154

What is the function of Routing Information Base (RIB) groups in protocol-independent routing?

A. They allow for the separation of routing information for different routing protocols
B. They store static route configurations
C. They are used to filter traffic based on routing protocol
D. They store aggregated routes and determine the forwarding of traffic based on routing decisions

Answer: A

Explanation:

RIB groups allow the separation of routing information for different routing protocols, enabling better organization and management of routing data. This helps in scenarios with multiple routing protocols running simultaneously.

Question: 155

Which of the following is true regarding load balancing in protocol-independent routing?

A. Load balancing can only be applied to multicast traffic
B. Load balancing is not supported in protocol-independent routing configurations
C. Load balancing is only effective for static routes in a routing table
D. Load balancing allows traffic to be distributed across multiple routes to the same destination

Answer: D

Explanation:

Load balancing in protocol-independent routing enables traffic to be distributed across multiple routes to the same destination, improving bandwidth utilization and network performance.

Question: 156

How does filter-based forwarding (FBF) work in a protocol-independent routing configuration?

A. FBF allows traffic to be forwarded based on filtering criteria, such as source or destination address
B. FBF uses routing protocols to dynamically route traffic based on network load
C. FBF only allows traffic to be routed to the nearest network destination
D. FBF applies only to Layer 3 traffic

Answer: A

Explanation:

Filter-based forwarding (FBF) enables traffic to be forwarded based on specific filtering criteria, such as

source IP address, destination IP address, or protocol type, providing more granular control over routing decisions.

Question: 157

Which of the following is an example of a generated route in a protocol-independent routing configuration?

A. A route that is manually configured by an administrator
B. A route that is automatically created for local network interfaces
C. A route that summarizes multiple destination addresses into one
D. A route created for each network destination based on its subnet mask

Answer: B

Explanation:

Generated routes are automatically created for local network interfaces. These routes are added to the routing table to ensure proper routing of traffic to directly connected networks.

Question: 158

Which of the following components does not directly participate in protocol-independent routing?

A. Routing Information Base (RIB)
B. Forwarding Information Base (FIB)
C. Static route configurations
D. Dynamic routing protocols

Answer: D

Explanation:

Dynamic routing protocols are not part of protocol-independent routing, which focuses on static, aggregate, and generated routes. Dynamic protocols, such as OSPF or BGP, dynamically adjust routing information, while protocol-independent routing relies on manual configurations and predefined routes.

Question: 159

When configuring static routes in a protocol-independent routing setup, which metric is typically associated with the route?

A. Bandwidt
B. Hop count
C. Administrative distance
D. Priority

Answer: C

Explanation:

Static routes are assigned an administrative distance, which is a value used to compare routes. A lower

administrative distance indicates a preferred route. Static routes generally have a fixed administrative distance of 1.

Question: 160

Which of the following statements is true regarding filter-based forwarding (FBF)?

A. FBF can be used to forward traffic based on policies such as source address or application type
B. FBF can only filter traffic based on IP address in the network
C. FBF requires the use of a dynamic routing protocol to function properly
D. FBF is specifically designed for forwarding multicast traffic based on defined filters

Answer: A
Explanation:

Filter-based forwarding (FBF) allows for traffic forwarding decisions based on various policies, including source address, destination address, application type, and other criteria, providing more flexible routing.

Question: 161

In which of the following scenarios would you use a generated route in protocol-independent routing?

A. To route traffic based on a specific subnet mask
B. To automatically add routes for directly connected interfaces
C. To aggregate multiple routes into one
D. To define backup routes that provide failover capabilities in case of primary route failure

Answer: B
Explanation:

Generated routes are automatically created for directly connected interfaces, ensuring that traffic destined for local networks is correctly routed without manual configuration.

Question: 162

What is the role of an aggregate route in a network?

A. To summarize multiple smaller networks into a larger, more efficient route
B. To define backup routes that can be used for network failover in case of primary route failure
C. To define static routes to remote networks
D. To route traffic based on the destination IP address and forwarding decisions

Answer: A
Explanation:

An aggregate route summarizes multiple smaller networks into one larger, more efficient route. This reduces the size of the routing table, helping improve routing efficiency and reduce complexity.

Question: 163

How does load balancing work in protocol-independent routing?

A. It uses the least-cost route for all traffic
B. It sends traffic through the fastest route available, based on network metrics
C. It prioritizes traffic based on its source IP address
D. It divides traffic evenly across multiple paths to the same destination

Answer: D
Explanation:

Load balancing divides traffic across multiple available paths to the same destination, optimizing the network's throughput and ensuring better utilization of available routes.

Question: 164

Which of the following is true about Martian addresses?

A. They are used to route traffic between different routing tables
B. They are private addresses used for internal network communications
C. They belong to reserved address ranges and cannot be routed on the Internet
D. They are used exclusively in multicast communications

Answer: C
Explanation:

Martian addresses fall within reserved address ranges and are not routable on the public Internet. These addresses are typically detected as invalid when encountered in a network.

Question: 165

What would be the result of applying a filter-based forwarding policy that routes traffic to a specific interface based on its source address?

A. Traffic would be forwarded to the designated interface if the source address matches the policy
B. Traffic would be dropped if the source address does not match the policy
C. The routing table would be updated to include the new source address
D. The interface would automatically be assigned the source address

Answer: A
Explanation:

In filter-based forwarding (FBF), traffic is forwarded to the designated interface if the source address matches the configured policy. This allows for more granular control over routing decisions based on address filtering.

Question: 166

Which type of route is automatically created for locally connected networks in protocol-independent

routing?

A. Static route
B. Aggregate route
C. Dynamic route
D. Generated route

Answer: D

Explanation:

Generated routes are automatically created for locally connected networks. These routes ensure that traffic destined for directly connected subnets is correctly routed without needing to be manually added to the routing table.

Question: 167

In protocol-independent routing, which of the following routing components is used to summarize multiple routes into a single route for efficiency?

A. Static route
B. Generated route
C. Aggregate route
D. Load balancing route

Answer: C

Explanation:

Aggregate routes summarize multiple smaller routes into a single route, reducing the size of the routing table and simplifying route management while maintaining efficient traffic forwarding.

Question: 168

How would you configure load balancing to distribute traffic across multiple equal-cost paths in a protocol-independent routing setup?

A. By assigning different administrative distances to each route
B. By configuring multiple static routes to the same destination
C. By defining a routing policy that prioritizes certain routes
D. By enabling dynamic routing protocols to select the optimal path

Answer: B

Explanation:

To configure load balancing, multiple static routes can be defined to the same destination with equal cost, allowing traffic to be distributed across those paths for better utilization of network resources.

Question: 169

Which of the following is the primary function of a Routing Information Base (RIB) group in protocol-independent routing?

A. To separate and organize routing information for different routing instances
B. To aggregate multiple routes from different sources into a single summarized route
C. To store only dynamically learned routes obtained from routing protocols
D. To route traffic based on specific load balancing criteria across multiple paths

Answer: A

Explanation:

RIB groups are used to separate and organize routing information for different routing instances. They allow the network to manage multiple routing tables, enabling better segregation and control of routing data for various applications or services.

Question: 170

What is the primary benefit of using filter-based forwarding (FBF) in a network?

A. It allows routing decisions based solely on Layer 3 addressing
B. It enhances network performance by automatically filtering out unwanted traffic
C. It enables advanced routing based on specific policies like source address or traffic type
D. It optimizes routing by learning traffic patterns dynamically

Answer: C

Explanation:

Filter-based forwarding (FBF) enables advanced routing decisions based on policies such as source address, destination address, or application type. This allows for more granular control of traffic forwarding based on predefined filtering rules.

Question: 171

Which of the following components can be used to manually route traffic between two different VLANs in protocol-independent routing?

A. Dynamic route
B. Aggregate route
C. Generated route
D. Static route

Answer: D

Explanation:

Static routes are manually configured and can be used to route traffic between two different VLANs by specifying the destination network and the next-hop IP address, ensuring that traffic flows between

VLANs.

Question: 172

What is the primary purpose of using Martian addresses in a network?

A. To route traffic between different subnets
B. To designate routes for multicast traffic
C. To configure private addressing for internal networks
D. To identify and filter out invalid or unreachable IP addresses

Answer: D

Explanation:

Martian addresses are used to identify and filter out invalid or unreachable IP addresses in the routing table. These addresses typically fall within reserved address ranges and cannot be routed on the public Internet.

Question: 173

In protocol-independent routing, what is the key characteristic of aggregate routes?

A. They are dynamically learned through routing protocols
B. They combine multiple routes into a single, summarized route
C. They are only used for load balancing purposes
D. They automatically detect Martian addresses

Answer: B

Explanation:

Aggregate routes combine multiple smaller routes into a single summarized route, helping reduce the size of the routing table and improving routing efficiency by summarizing network destinations.

Question: 174

How does filter-based forwarding (FBF) handle traffic that does not match a defined policy?

A. It forwards the traffic using the default routing table B. It drops the traffic to ensure network security C. It queues the traffic until it matches a policy D. It forwards the traffic to the nearest router

Answer: A

Explanation:

If traffic does not match a defined policy in filter-based forwarding (FBF), it is forwarded using the default routing table, ensuring that all traffic has a path to its destination, even if it doesn't meet specific filtering criteria.

Question: 175

Which of the following is a feature of load balancing in protocol-independent routing?

A. It automatically increases the bandwidth of availablepaths
B. It distributes traffic across multiple paths with the samecost
C. It prioritizes traffic based on its type or application
D. It merges multiple routing tables into a single path

Answer: B
Explanation:

Load balancing distributes traffic across multiple equal-cost paths to the same destination, optimizing network utilization and ensuring that no single path is overburdened, which improves overall performance.

Question: 176

What type of route would automatically be created for a network that is directly connected to a router interface?

A. Static route
B. Aggregate route
C. Generated route
D. Dynamic route

Answer: C
Explanation:

Generated routes are automatically created for networks directly connected to a router interface, allowing traffic to be routed locally without the need for manual configuration.

Question: 177

Which of the following statements about static routes is true?

A. Static routes automatically adjust to network changes
B. Static routes require manual configuration and do not change unless manually updated
C. Static routes can only be used for Layer 3 communication
D. Static routes are best used for networks with fluctuating paths

Answer: B
Explanation:

Static routes require manual configuration and do not automatically adjust to network changes. They remain fixed until the network administrator manually updates or removes them, making them reliable for stable networks.

Question: 178

Which of the following configurations could be used to ensure traffic is forwarded based on specific criteria such as source address in filter-based forwarding?

A. A routing policy
B. A static route configuration
C. An aggregate route
D. A dynamic routing protocol

Answer: A
Explanation:

A routing policy is used in filter-based forwarding (FBF) to define forwarding decisions based on specific criteria, such as source address, destination address, or application type. This allows for more precise control over traffic forwarding.

Question: 179

What is the benefit of using load balancing with multiple equal-cost paths in protocol-independent routing?

A. It improves network security by routing traffic to more secure paths
B. It optimizes the use of available bandwidth by distributing traffic evenly across paths
C. It reduces the size of the routing table by consolidating multiple routes into one
D. It enables routing based on traffic type and application

Answer: B
Explanation:

Load balancing optimizes the use of available bandwidth by distributing traffic evenly across multiple equal-cost paths, improving network performance and preventing any single path from becoming overloaded.

Question: 180

In which situation would you use an aggregate route in protocol-independent routing?

A. To summarize multiple smaller subnets into a single larger network
B. To route traffic to specific devices based on MAC addresses
C. To divide a large network into smaller, more manageable segments
D. To automatically route traffic between two VLANs

Answer: A
Explanation:

An aggregate route is used to summarize multiple smaller subnets into a single, larger network, which reduces the size of the routing table and simplifies routing by grouping related network destinations together.

Question: 181

Which of the following are characteristics of static routes in protocol-independent routing? (Choose two)

A. Static routes are manually configured by the network administrator
B. Static routes automatically update when network topology changes
C. Static routes are preferred over dynamic routes in all situations
D. Static routes do not change unless manually adjusted by the administrator

Answer: A, D **Explanation:**

Static routes are manually configured and do not change automatically based on network topology changes. They remain in the routing table until the administrator manually modifies them.

Question: 182

Which of the following are true about aggregate routes in protocol-independent routing? (Choose two)

A. Aggregate routes summarize multiple smaller networks into one larger route
B. Aggregate routes are dynamically learned through a routing protocol
C. Aggregate routes improve routing table efficiency by reducing the number of entries
D. Aggregate routes are used only for load balancing traffic across multiple paths

Answer: A, C
Explanation:

Aggregate routes combine multiple smaller network prefixes into a single, summarized route, which reduces the size of the routing table and improves routing efficiency.

Question: 183

Which of the following are valid use cases for generated routes in protocol-independent routing? (Choose two)

A. Manually configured to override dynamic routing decisions
B. Automatically created for locally connected networks
C. Created for static routes within a routing protocol
D. Automatically created for local interfaces to ensure reachability

Answer: B, D
Explanation:

Generated routes are automatically created for networks directly connected to a router, ensuring proper routing for locally connected interfaces. These routes are added to the routing table without manual configuration.

Question: 184

Which of the following statements are true regarding Martian addresses? (Choose two)

A. Martian addresses areused to route traffic between different private networks
B. Martian addresses fall within reserved address ranges and are not routable on the Internet
C. Martian addresses areused to identify and filter out invalid or unreachable IPaddresses
D. Martian addresses arepart of the public IP address range

Answer: B, C
Explanation:

Martian addresses fall within reserved or invalid address ranges and are not routable on the public Internet. These addresses are often used to detect and filter out invalid or unreachable addresses in the network.

Question: 185

Which of the following are features of filter-based forwarding (FBF) in protocol-independent routing? (Choose two)

A. FBF enables traffic to be forwarded based on specific criteria such as source or destination IP
B. FBF requires the use of a dynamic routing protocol for operation
C. FBF is applied to Layer 2 traffic only
D. FBF allows for more granular control of traffic based on policy

Answer: A, D
Explanation:

Filter-based forwarding allows for traffic to be forwarded based on specific filtering criteria, such as source or destination IP addresses, providing more granular control over traffic flow in the network.

Question: 186

Which of the following are valid benefits of using load balancing in protocol-independent routing? (Choose two)

A. It increases network security by isolating traffic from different VLANs
B. It ensures that all traffic is routed through the fastest available path
C. It allows traffic to be distributed across multiple paths with the same cost
D. It optimizes bandwidth utilization by preventing any single path from becoming overloaded

Answer: C, D
Explanation:

Load balancing allows for traffic to be distributed evenly across multiple paths with equal cost, optimizing bandwidth utilization and preventing any single path from being overwhelmed, thereby improving overall network performance.

Question: 187

Which of the following are examples of routes that can be found in a Routing Information Base (RIB)? (Choose two)

A. Generated routes for directly connected networks
B. Static routes manually configured by the network administrator
C. Dynamic routes learned through OSPF or BGP
D. Only routes learned via static routing protocols

Answer: A, B
Explanation:

The RIB contains routes that include manually configured static routes, as well as generated routes for locally connected networks. These routes are used for determining the best path for traffic forwarding.

Question: 188

Which of the following routing components can be used to manage traffic based on specific policies in protocol-independent routing? (Choose two)

A. Static routes
B. Dynamic routing protocols
C. Filter-based forwarding
D. Routing Information Base (RIB) groups

Answer: C, D
Explanation:

Filter-based forwarding allows for traffic to be managed and forwarded based on specific filtering policies. RIB groups help organize routing information based on different routing instances, providing more granular control over traffic forwarding.

Question: 189

Which of the following best describes the function of the Link-State Database (LSDB) in OSPF?

A. It only stores the best routes for direct connections

B. It stores routing table information for all IP subnets in the network
C. It contains a copy of the entire network's IP address table
D. It stores routing information and contains both the network topology and routes

Answer: D
Explanation:

The Link-State Database (LSDB) in OSPF contains the network topology, including all the link-state advertisements (LSAs) that describe the state of OSPF links. It helps OSPF routers build their routing tables by understanding the network's topology.

Question: 190

Which OSPF packet type is used to establish and maintain neighbor relationships between OSPF routers?

A. Link-State Advertisement (LSA)
B. Link-State Update (LSU) Packet
C. Database Description (DBD) Packet
D. Hello Packet

Answer: D
Explanation:

Hello packets in OSPF are used to establish and maintain neighbor relationships between routers. They are sent periodically and contain information that helps routers identify each other and form OSPF adjacencies.

Question: 191

In OSPF, what is the purpose of the Router ID?

A. To uniquely identify an OSPF router in the network
B. To indicate the router's preferred path
C. To define the router's primary network address
D. To specify the router's priority for becoming a Designated Router (DR)

Answer: A
Explanation:

The Router ID in OSPF is a unique identifier for each OSPF router within an area. It helps OSPF routers distinguish themselves from one another in the network and is essential for maintaining accurate routing tables.

Question: 192

What is the role of a Designated Router (DR) in OSPF?

A. To handle the OSPF protocol for all other routers in the network
B. To reduce the OSPF link-state flooding within a broadcast domain C. To maintain a backup routing table for other routers
D. To select the best path between routers in the OSPF domain

Answer: B
Explanation:

The Designated Router (DR) in OSPF is responsible for reducing the amount of OSPF link-state flooding within a broadcast domain. The DR acts as the central point for exchanging OSPF LSAs, preventing each router from directly exchanging LSAs with every other router.

Question: 193

Which OSPF router type is used to connect different OSPF areas?

A. Backbone Router
B. Internal Router
C. Area Border Router (ABR)
D. Autonomous System Boundary Router (ASBR)

Answer: C
Explanation:

The Area Border Router (ABR) connects different OSPF areas and maintains separate link-state databases for each area. The ABR exchanges routing information between areas to ensure proper OSPF route propagation across the network.

Question: 194

Which of the following is true about the OSPF Link-State Advertisement (LSA) types?

A. Type 1 LSAs describe the state of links between OSPF routers
B. Type 2 LSAs describe routes between different OSPF areas
C. Type 3 LSAs are used by OSPF external routes
D. Type 4 LSAs are used by the DR to advertise network information

Answer: A
Explanation:

Type 1 LSAs in OSPF describe the state of links between OSPF routers within the same area. They are used to build the link-state database and assist routers in understanding the topology of the area.

Question: 195

How does the OSPF router establish an adjacency with a neighbor?

A. By sending an LSR to request the best routes
B. By exchanging LSRs (Link-State Requests)
C. By directly connecting the routers with a shared network
D. By sending a Hello packet to initiate the OSPF negotiation

Answer: D
Explanation:

OSPF routers establish an adjacency by sending a Hello packet. This packet helps routers discover each other, verify compatibility, and form OSPF neighbor relationships.

Question: 196

What does the OSPF "Full" state signify in the context of OSPF neighbor relationships?

A. The routers have established a neighbor relationship but have not exchanged routing information
B. The routers have exchanged and synchronized their LSDBs
C. The router has successfully calculated the best path
D. The router is operating in a passive mode and does not exchange LSAs

Answer: B
Explanation:

The "Full" state in OSPF indicates that two routers have successfully exchanged and synchronized their Link-State Databases (LSDBs). This means they have completed the handshake and are fully exchanging OSPF routing information.

Question: 197

Which OSPF area type is used to connect an OSPF network to external networks?

A. Backbone Area
B. Stub Area
C. Not-So-Stubby Area (NSSA)
D. External Area

Answer: C
Explanation:

The Not-So-Stubby Area (NSSA) is an OSPF area type that allows external routes (from outside the OSPF domain) to be imported into the area. It is used when an area needs to connect to external networks but does not need full external route propagation.

Question: 198

What is the primary purpose of the BGP protocol?

A. To manage the flow of traffic within a single autonomous system
B. To provide routing decisions based on hop count
C. To exchange routing information between different autonomous systems
D. To enable multicast routing across the Internet

Answer: C
Explanation:

BGP (Border Gateway Protocol) is used to exchange routing information between different autonomous systems (ASes) on the Internet. It plays a critical role in inter-domain routing by selecting the best paths based on policies and attributes.

Question: 199

Which BGP message type is used to establish, maintain, or terminate a BGP session?

A. Open
B. Update
C. Notification
D. Keepalive

Answer: A
Explanation:

The Open message is used to establish, maintain, or terminate a BGP session between peers. It includes essential information such as the version of BGP being used and the router's AS number.

Question: 200

Which of the following is a key factor in BGP's route selection process?

A. Administrative distance
C. Subnet mask length
C. Link state
D. AS path length

Answer: D
Explanation:

The AS path length is a key factor in BGP's route selection process. BGP prefers the path with the shortest AS path, as it reflects fewer hops between ASes and is generally more reliable.

Question: 201

Which of the following BGP attributes is used to prevent routing loops by keeping track of the ASes a route has passed through?

A. Local Preference
D. MED (Multi-Exit Discriminator)
C. AS Path
D. Next Hop

Answer: C
Explanation:

The AS Path attribute is used by BGP to prevent routing loops by tracking the sequence of ASes a route has passed through. If a route is advertised back to the originating AS, it is rejected to avoid loops.

Question: 202

What is the role of the Next Hop attribute in BGP?

A. To specify the router from which the route was learned
B. To identify the destination IP address for the route
C. To prioritize routes within a BGP session
D. To determine the best exit point from an AS

Answer: D
Explanation:

The Next Hop attribute in BGP specifies the IP address of the next router to forward traffic to for a particular route. It determines the exit point from the AS and helps in routing traffic correctly.

Question: 203

Which of the following statements is true about IBGP (Internal BGP) and EBGP (External BGP)?

A. IBGP is used between routers within the same AS, while EBGP is used between routers in different ASes
B. IBGP is used to exchange routing information between different ASes
C. EBGP uses a different routing protocol to exchange information within the same AS
D. IBGP requires a higher administrative distance than EBGP

Answer: A
Explanation:

IBGP is used to exchange routing information between routers within the same AS, while EBGP is used between routers in different ASes. This separation helps manage inter-domain routing across the Internet.

Question: 204

What is the main purpose of the MED (Multi-Exit Discriminator) attribute in BGP?

A. To specify the next-hop IP address for a route
B. To select the best path when multiple exit points are available from an AS
C. To ensure that the routing table is up-to-date
D. To prevent routing loops between ASes

Answer: B
Explanation:

The MED attribute in BGP is used to influence routing decisions by specifying the preferred exit point when multiple routes are available from the same AS. A lower MED value is preferred over a higher one.

Question: 205

Which of the following BGP message types is used to update routing information?

A. Open
B. Keepalive

C. Notification
D. Update

Answer: D
Explanation:

The Update message in BGP is used to exchange routing information between BGP peers. It is sent when there are changes in the routing table, such as new routes or changes to existing routes.

Question: 206

What is the main purpose of the AS Path attribute in BGP?

A. To track the number of hops between routers in the same AS
B. To specify the originating router for a route
C. To define the preferred exit point from an AS
D. To avoid routing loops by listing all ASes a route has traversed

Answer: D
Explanation:

The AS Path attribute in BGP is used to prevent routing loops by keeping track of the ASes a route has passed through. If a route advertises an AS it has already traversed, the route is discarded to prevent loops.

Question: 207

Which of the following BGP attributes is used to determine the best route when multiple routes exist to the same destination?

A. Local Preference
B. AS Path
C. Next Hop
D. MED (Multi-Exit Discriminator)

Answer: A
Explanation:

Local Preference is used in BGP to determine the best route when multiple paths exist. It is a well-known discretionary attribute, and the route with the highest Local Preference is selected as the best route for outgoing traffic.

Question: 208

What is the primary role of OSPF's Link-State Advertisement (LSA)?

A. To share routing table information between OSPF routers
B. To exchange the network topology and link state information
C. To calculate the best path to each destination
D. To prioritize routes based on bandwidth

Answer: B
Explanation:

LSAs in OSPF are used to exchange network topology and link-state information between routers, which helps them build a complete map of the network. This allows OSPF routers to calculate the best paths to each destination.

Question: 209

Which OSPF packet type is responsible for carrying actual routing information between routers?

A. Hello Packet
B. Database Description (DBD) Packet
C. Link-State Advertisement (LSA) Packet
D. Link-State Update (LSU) Packet

Answer: D
Explanation:

Link-State Update (LSU) packets carry actual routing information between OSPF routers. These packets contain LSAs that update the routing table with new or changed information.

Question: 210

Which of the following best describes the OSPF router ID selection process?

A. The router ID is manually configured by the network administrator
B. The lowest IP address on a physical interface is selected as the router ID
C. The highest IP address on the loopback interface is selected as the router ID
D. The router ID is randomly selected at the time of OSPF router startup

Answer: C
Explanation:

The OSPF router ID is automatically selected based on the highest IP address on the loopback interface. If no loopback interfaces are configured, the highest IP address on any physical interface is chosen.

Question: 211

What is the function of the OSPF Designated Router (DR)?

A. To perform OSPF path calculations
B. To reduce the amount of OSPF link-state flooding on multi-access networks
C. To directly forward OSPF traffic to external networks
D. To advertise OSPF routes to other routers in the area

Answer: B
Explanation:

The Designated Router (DR) in OSPF reduces the amount of link-state flooding on multi-access networks, like Ethernet. The DR acts as the central point for exchanging OSPF LSAs, preventing all

routers from flooding LSAs to every other router.

Question: 212

Which OSPF router type is responsible for connecting OSPF areas to the backbone?

A. Backbone Router
B. Autonomous System Boundary Router (ASBR)
C. Area Border Router (ABR)
D. Internal Router

Answer: C
Explanation:

The Area Border Router (ABR) connects OSPF areas to the backbone (Area 0). It advertises routing information between areas and maintains separate link-state databases for each area it connects to.

Question: 213

What does OSPF's "Full" state indicate during the neighbor establishment process?

A. The routers have established a basic neighbor relationship but not exchanged LSAs

B. The routers have completed the LSR exchange and are now exchanging actual routes
C. The routers have fully synchronized their LSDBs and are exchanging LSAs
D. The routers have determined the best path for routing traffic

Answer: C
Explanation:

The "Full" state in OSPF indicates that two routers have completed the exchange of link-state information and fully synchronized their Link-State Databases (LSDBs), meaning they have successfully established a neighbor relationship.

Question: 214

In OSPF, which of the following packet types is used to check the OSPF router's neighbor status and maintain the OSPF session?

A. Link-State Advertisement (LSA)
B. Database Description (DBD)
C. Link-State Update (LSU)
D. Hello Packet

Answer: D
Explanation:

Hello packets are used to check the OSPF router's neighbor status and maintain the OSPF session. These packets also help in forming neighbor relationships by verifying that routers are compatible.

Question: 215

Which BGP attribute is used to avoid routing loops by listing the ASes a route has traversed?

A. MED (Multi-Exit Discriminator)
B. Next Hop
C. AS Path
D. Local Preference

Answer: C
Explanation:

The AS Path attribute in BGP is used to avoid routing loops by keeping track of the ASes a route has passed through. If a route is advertised back to the originating AS, it is discarded to prevent loops.

Question: 216

Which BGP message type is used to send routing updates between BGP peers?

A. Open

B. Update
C. Keepalive
D. Notification

Answer: B
Explanation:

The Update message in BGP is used to send routing updates between peers, including information about new routes or changes to existing routes.

Question: 217

Which BGP attribute is used to influence the selection of the best path when multiple routes are available?

A. MED (Multi-Exit Discriminator)
B. AS Path
C. Next Hop
D. Local Preference

Answer: D
Explanation:

Local Preference is a BGP attribute that influences path selection within an AS. The route with the highest Local Preference is selected as the preferred path for outgoing traffic.

Question: 218

In BGP, which of the following message types is used to maintain the connection between BGP peers?

A. Keepalive
B. Update
C. Open
D. Notification

Answer: A
Explanation:

The Keepalive message is used in BGP to maintain the connection between BGP peers by confirming that the session is still active and operational.

Question: 219

Which of the following is a valid reason to use IBGP (Internal BGP) within an AS?

A. To exchange routing information between routers in the same AS
B. To advertise routes from one AS to another

C. To propagate external routes to the Internet
D. To connect different routing tables in separate ASes

Answer: A
Explanation:

IBGP is used to exchange routing information between routers within the same AS. It helps ensure that routers within the AS have consistent routing information, particularly in larger networks with multiple routers.

Question: 220

What does the Next Hop attribute in BGP specify?

A. The AS path for a given route
B. The number of hops to reach the destination
C. The administrative distance for a route
D. The IP address of the router to which traffic should be sent

Answer: D
Explanation:

The Next Hop attribute in BGP specifies the IP address of the next router to which traffic should be sent for a specific route. This helps determine the correct exit point from an AS.

Question: 221

What is the function of the MED (Multi-Exit Discriminator) attribute in BGP?

A. To prioritize routes within the same Autonomous System (AS) based on certain metrics
B. To influence the selection of an exit point when multiple exit points are available for routing
C. To track and record the path taken by a route as it travels between Autonomous Systems (ASes)
D. To specify the next hop for a route, determining how traffic is forwarded across the network

Answer: B
Explanation:

The MED (Multi-Exit Discriminator) attribute in BGP is used to influence the choice of exit point when multiple exit points exist between ASes. A lower MED value is preferred over a higher value.

Question: 222

What type of BGP session is used to exchange routing information between routers in different ASes?

A. IBGP
B. EBGP
C. ASBR

D. iBGP

Answer: B
Explanation:

EBGP (External BGP) is used to exchange routing information between routers in different ASes, facilitating inter-domain routing across the Internet or between different networks.

Question: 223

Which BGP attribute is used to prevent routing loops between different ASes?

A. AS Path
B. MED
C. Local Preference
D. Next Hop

Answer: A
Explanation:

The AS Path attribute is used to prevent routing loops by keeping track of the sequence of ASes a route has traversed. If a route is advertised back to the originating AS, it is discarded to avoid a loop.

Question: 224

What is the primary function of the BGP Route Selection process?

A. To discard all routes that are not originated within the same AS
B. To route traffic between different OSPF areas
C. To determine the best route based on the BGP attributes
D. To calculate the shortest path based on hop count

Answer: C
Explanation:

The BGP Route Selection process determines the best route to a destination by evaluating various BGP attributes such as AS Path, Local Preference, and MED, selecting the most optimal route.

Question: 225

Which of the following is a characteristic of EBGP (External BGP) as compared to IBGP (Internal BGP)?

A. EBGP peers can reside within the same AS
B. EBGP peers use a different AS number
C. EBGP does not support the use of route maps
D. EBGP peers always have a directly connected relationship

Answer: C
Explanation:

EBGP peers must reside in different ASes and always have a directly connected relationship. This is a key distinction from IBGP, where peers reside within the same AS.

Question: 226

Which BGP attribute is used to select the best path when multiple paths exist to the same destination?

A. MED
B. Local Preference
C. AS Path
D. Next Hop

Answer: B
Explanation:

The Local Preference attribute is used in BGP to select the best path when multiple paths to the same destination exist. The path with the highest Local Preference is preferred for outgoing traffic from the AS.

Question: 227

In BGP, which of the following best describes the function of the AS Path attribute?

A. It specifies the next-hop router for a given route
B. It is used to determine the best path based on bandwidth
C. It lists the ASes a route has passed through, preventing routing loops
D. It assigns a preference value to routes within the same AS

Answer: C
Explanation:

The AS Path attribute in BGP lists the ASes that a route has passed through, helping prevent routing loops. If a route contains the AS number of the originating AS, it is discarded to avoid loops.

Question: 228

Which of the following are key functions of OSPF's Link-State Database (LSDB)? (Choose two)

A. Storing routing table information for all IP subnets
B. Storing all Link-State Advertisements (LSAs) received from OSPF neighbors
C. Representing the network topology of the OSPF area
D. Storing only the best routes for traffic forwarding

Answer: B, C
Explanation:

The OSPF Link-State Database (LSDB) stores all the LSAs received from OSPF neighbors, which collectively represent the network topology of the OSPF area. This database is essential for calculating the best paths.

Question: 229

Which of the following OSPF packet types are used to establish and maintain OSPF neighbor relationships? (Choose two)

A. Link-State Advertisement (LSA)
B. Link-State Update (LSU) Packet
C. Database Description (DBD) Packet
D. Hello Packet

Answer: C, D **Explanation:**

Hello packets are used to establish and maintain OSPF neighbor relationships. Database Description (DBD) packets are used during the initial exchange of OSPF routing information to describe the LSDB contents.

Question: 230

Which of the following BGP attributes are used to influence the selection of the best route? (Choose two)

A. Local Preference
B. AS Path
C. MED (Multi-Exit Discriminator)
D. Next Hop

Answer: A, C **Explanation:**

Local Preference and MED are BGP attributes used to influence the selection of the best route. Local Preference is used to prefer routes within an AS, while MED is used to prefer routes to the same destination when multiple exit points are available.

Question: 231

Which OSPF router types are responsible for handling traffic between OSPF areas? (Choose two)

A. Internal Router
B. Area Border Router (ABR)
C. Autonomous System Boundary Router (ASBR)
D. Backbone Router

Answer: B, D

Explanation:

Area Border Routers (ABRs) are responsible for handling routing between different OSPF areas. Backbone Routers are part of the backbone area (Area 0) and connect different OSPF areas to ensure inter-area communication.

Question: 232

Which of the following are characteristics of IBGP (Internal BGP)? (Choose two)

A. IBGP is used to exchange routes between different ASes
B. IBGP peers must be directly connected
C. IBGP does not require the next-hop IP address to be reachable
D. IBGP is used to propagate routes within the same AS

Answer: B, D Explanation:

IBGP is used to exchange routes within the same AS, and peers must be directly connected for the BGP session to work. Unlike EBGP, IBGP does not change the next-hop IP address and relies on internal routing to reach destinations.

Question: 233

Which OSPF packet types are used for advertising and requesting link-state information? (Choose two)

A. Hello Packet
B. Link-State Advertisement (LSA)
C. Link-State Update (LSU)
D. Database Description (DBD)

Answer: C, D
Explanation:

Link-State Update (LSU) packets are used to advertise link-state information, while Database Description (DBD) packets are used to describe the LSDB contents and request additional information during the OSPF handshake process.

Question: 234

Which of the following BGP message types are exchanged between BGP peers? (Choose two)

A. Open
B. Keepalive
C. Update
D. Hello

Answer: A, C Explanation:

Open messages are used to establish a BGP session between peers, and Update messages are used to exchange routing information. Keepalive messages ensure that the BGP session is still active, but they are not used to exchange routing information.

Question: 235

Which of the following OSPF areas have specific limitations or rules? (Choose two)

A. Backbone Area (Area 0)
B. Stub Area
C. Not-So-Stubby Area (NSSA)
D. External Area

Answer: B, C Explanation:

Stub Areas limit the type of routing information that can be received to reduce the size of the routing table. Not-So-Stubby Areas (NSSAs) are similar but allow external routes to be injected into the area, making them more flexible than standard Stub Areas.

Question: 236

Which BGP attributes help prevent routing loops and influence path selection? (Choose two)

A. AS Path
B. MED
C. Local Preference
D. Next Hop

Answer: A, B Explanation:

The AS Path attribute prevents routing loops by keeping track of the ASes a route has traversed, while the MED (Multi-Exit Discriminator) influences the selection of the exit point when multiple paths are available from the same AS.

Question: 237

Which of the following OSPF states are part of the OSPF neighbor relationship establishment process? (Choose two)

A. Full
B. Two-Way
C. Down

D. Init

Answer: B, D Explanation:

The "Two-Way" state is used when OSPF routers recognize each other and are ready to establish a full adjacency. The "Init" state indicates that OSPF has just discovered a potential neighbor but has not yet exchanged link-state information.

Question: 238

Which BGP attributes are used to influence outbound route selection within an AS? (Choose two)

A. AS Path
C. Next Hop
C. MED

D. Local Preference

Answer: B, D
Explanation:

Local Preference is used to prefer routes within an AS, while MED (Multi-Exit Discriminator) is used to influence the selection of the exit point when multiple exit points are available from the same AS. Both attributes affect outbound route selection.

Question: 239

Which of the following describes the function of a Link Aggregation Group (LAG) in high availability setups?

A. LAG combines multiple physical links into a single logical link for increased bandwidth and redundancy
B. LAG is used to detect failed links and automatically reroute traffic to the backup path
C. LAG aggregates routing information between multiple routers
D. LAG provides automatic failover for network devices

Answer: A
Explanation:

A Link Aggregation Group (LAG) combines multiple physical links into a single logical link, providing increased bandwidth and redundancy. This setup helps ensure high availability by distributing traffic across the aggregated links, allowing for failover if one link fails.

Question: 240

What is the primary purpose of Graceful Restart (GR) in high availability configurations?

A. To allow network devices to restart without causing downtime
B. To provide redundancy between multiple routers
C. To ensure minimal disruption during routing protocol reconvergence
D. To detect and resolve network loops during failures

Answer: C
Explanation:

Graceful Restart (GR) allows a router to restart its routing protocol processes without disrupting ongoing network operations. The router can retain its routing information during the restart, minimizing the disruption and allowing for quicker reconvergence.

Question: 241

How does Bidirectional Forwarding Detection (BFD) enhance high availability in networks?

A. By monitoring the physical layer for link status changes and detecting failures
B. By increasing the overall bandwidth of the network through link aggregation
C. By aggregating multiple network links into a single path for better efficiency

D. By providing fast detection of link failures and quickly rerouting traffic

Answer: D
Explanation:

Bidirectional Forwarding Detection (BFD) provides fast detection of link failures by monitoring the forwarding path between devices. It helps improve high availability by quickly notifying the network of a failure and enabling rapid traffic rerouting.

Question: 242

What is the role of a Virtual Chassis in high availability configurations?

A. To combine multiple routers into a single logical device for centralized management and failover
B. To enable routers to automatically synchronize routing tables for consistent forwarding
C. To aggregate multiple links between network devices to increase bandwidth and redundancy
D. To provide automatic detection of hardware failures and trigger a failover process to ensure uptime

Answer: A
Explanation:

A Virtual Chassis allows multiple physical devices (such as switches) to be grouped together and managed as a single logical device. This provides high availability by offering centralized management and seamless failover in case of device failure.

Question: 243

Which of the following is a key benefit of using Link Aggregation Groups (LAG) in a data center?

A. It reduces the number of required IP addresses
B. It increases the speed of routing protocol convergence
C. It simplifies routing protocols within the data center
D. It increases bandwidth and provides redundancy across multiple links

Answer: D
Explanation:

LAG increases bandwidth by combining multiple physical links into a single logical link. It also provides redundancy by enabling traffic to be distributed across the aggregated links, ensuring continued network operation if one link fails.

Question: 244

What is the purpose of using Bidirectional Forwarding Detection (BFD) in conjunction with OSPF or BGP?

A. To speed up the detection of OSPF neighbor relationships and improve convergence time
B. To ensure BGP sessions are securely established and maintained between peers
C. To detect link failures quickly and trigger fast rerouting

D. To prevent route flapping in BGP

Answer: C

Explanation:

BFD is used in conjunction with routing protocols like OSPF and BGP to detect link failures quickly. It provides rapid failure detection, which helps trigger fast rerouting to maintain high availability and minimize network downtime.

Question: 245

Which of the following is a valid method for monitoring the health of Link Aggregation Groups (LAG)?

A. Checking the interface status using the "show interfaces" command
B. Running the "ping" command to verify layer 2 connectivity
C. Using a router's routing table to track LAG health
D. Configuring BFD to monitor the status of each LAG member

Answer: A

Explanation:

The "show interfaces" command can be used to monitor the health of LAGs, displaying the status of each individual interface within the group and providing insights into the overall health of the aggregated link.

Question: 246

Which of the following tools is commonly used to troubleshoot high availability components like LAG, GR, or BFD?

A. Tracepath
B. Show commands and logs
C. Spanning Tree Protocol (STP)
D. RMON (Remote Monitoring)

Answer: B

Explanation:

Show commands and logs are commonly used to troubleshoot high availability components like LAG, GR, and BFD. They provide detailed diagnostic information about the status, configuration, and errors related to these components.

Question: 247

What is a potential advantage of using Graceful Restart (GR) during network maintenance?

A. It allows routing tables to be cleared and rebuilt automatically
B. It helps maintain routing information while the router restarts, reducing downtime
C. It prevents BGP session loss during maintenance, ensuring continuous communication between peers

D. It aggregates routing protocols across different Autonomous Systems (ASes) to optimize routing

Answer: B

Explanation:

Graceful Restart (GR) allows a router to retain its routing information while restarting its routing protocol processes. This reduces the impact of the restart and minimizes downtime by maintaining the existing routing state during maintenance.

Question: 248

Which of the following statements best describes how LAG contributes to high availability?

A. It ensures that multiple network devices are connected and operate together as a single unit for increased reliability
B. It prevents data loss by duplicating packets across all links
C. It provides redundancy and load balancing by aggregating links and maintaining active traffic paths
D. It helps optimize routing table convergence by limiting the number of active routes

Answer: C

Explanation:

LAG contributes to high availability by aggregating multiple physical links into a single logical link, providing redundancy and load balancing. If one link fails, traffic is automatically rerouted over the remaining active links.

Question: 249

Which high availability feature is used to ensure the network continues operating when a device or link fails?

A. Link aggregation B. Graceful Restart C. Virtual Chassis
D. All of the above

Answer: D

Explanation:

All of the listed features—Link Aggregation, Graceful Restart, and Virtual Chassis—help ensure high availability by providing redundancy, quick recovery, and failover mechanisms in case of device or link failures.

Question: 250

Which configuration command is used to enable Link Aggregation on a Junos OS device?

A. set interfaces ge-0/0/1 aggregation-mode
B. set interfaces ge-0/0/1 ether-channel
C. set interfaces ge-0/0/1 unit 0 family ethernet-switching
D. set interfaces ae0 aggregation-group

Answer: D

Explanation:

In Junos OS, the set interfaces ae0 aggregation-group command is used to configure Link Aggregation on an interface. The ae0 interface is a logical interface representing the aggregated links.

Question: 251

What does Bidirectional Forwarding Detection (BFD) provide in terms of high availability?

A. It reduces the convergence time of routing protocols during failure
B. It automatically assigns IP addresses to redundant links
C. It performs path selection based on multiple metrics
D. It decreases the administrative overhead in link management

Answer: A

Explanation:

BFD provides rapid detection of link failures, significantly reducing the convergence time of routing protocols. This helps maintain high availability by quickly rerouting traffic in case of a link failure.

Question: 252

In the context of high availability, what does the term "graceful restart" refer to?

A. A method for switching to a backup router
B. The ability to maintain routing information while a router or protocol restarts
C. A procedure to automatically update routing tables during maintenance
D. A process that optimizes the stability of BGP sessions

Answer: B

Explanation:

Graceful Restart refers to the ability to maintain routing information and minimize disruption while a router or routing protocol restarts. This feature allows for faster recovery and ensures that the network remains stable during maintenance.

Question: 253

What is a key feature of a Virtual Chassis in maintaining high availability?

A. It enables failover for routing protocols, ensuring continuous network connectivity
B. It provides path selection based on link cost, optimizing data traffic routing
C. It automatically generates backup IP addresses to ensure redundancy during failures
D. It allows multiple physical switches to act as a single logical switch

Answer: D

Explanation:

A Virtual Chassis allows multiple physical switches to operate as a single logical device, simplifying management and ensuring high availability by allowing for failover if one switch fails.

Question: 254

Which of the following does Link Aggregation (LAG) improve in high availability scenarios?

A. Routing table size
B. Redundancy and bandwidth utilization
C. The number of IP addresses required
D. Router CPU load

Answer: B

Explanation:

LAG improves redundancy by aggregating multiple links into a single logical interface and increases bandwidth utilization by distributing traffic across all the aggregated links, ensuring that the network remains operational if one link fails.

Question: 255

What role does the "Backup Designated Router" (BDR) play in OSPF?

A. It replaces the DR in case of failure and handles OSPF packet exchange
B. It synchronizes OSPF database information between routers
C. It detects link failures and triggers rerouting in OSPF networks
D. It prevents OSPF routing loops during network failure

Answer: A

Explanation:

The Backup Designated Router (BDR) in OSPF is ready to take over the role of the Designated Router (DR) if it fails. It listens to OSPF packets and keeps a copy of the network's LSDB (Link-State Database) to be able to assume the DR role immediately if needed.

Question: 256

Which of the following is a major advantage of using a Virtual Chassis configuration?

A. It increasesthe size of the routing table by combining all networkpaths
B. It allows formore complex routing configurations than standard switches
C. It simplifiesnetwork management by combining multiple physicalswitches intoone logical switch
D. It improvesthe network's security by isolating different devices inseparate VLANs

Answer: C

Explanation:

A Virtual Chassis configuration combines multiple physical switches into a single logical switch,

simplifying network management, reducing configuration overhead, and improving redundancy and failover capabilities.

Question: 257

Which of the following statements about Bidirectional Forwarding Detection (BFD) is correct?

A. BFD operates only on Layer 3 devices and routes IP traffic
B. BFD is a feature used for load balancing traffic between multiple routers
C. BFD is used to configure IP addresses in Layer 3 routing protocols
D. BFD detects network failures by sending control packets between routers to detect failures quickly

Answer: D
Explanation:

BFD is a fast failure detection protocol that sends control packets between routers to quickly detect network failures. It is used in high availability scenarios to speed up the rerouting process when a failure occurs.

Question: 258

Which of the following best describes the main purpose of Link Aggregation Groups (LAG) in high availability configurations?

A. To provide a backup link in case of a hardware failure
B. To dynamically adjust the network topology based on traffic load
C. To allow multiple IP addresses to be assigned to the same interface
D. To combine multiple physical links into one logical link for increased bandwidth and redundancy

Answer: D
Explanation:

Link Aggregation Groups (LAG) combine multiple physical links into one logical link, which provides increased bandwidth and redundancy. This ensures that if one physical link fails, the network can continue to operate using the remaining links.

Question: 259

What is the primary function of the Graceful Restart (GR) feature in high availability configurations?

A. To allow a router to restart its routing protocol processes without causing network disruption
B. To automatically reroute traffic during a network failure
C. To configure backup routes for failover scenarios
D. To synchronize routing information between different routers

Answer: A
Explanation:

Graceful Restart (GR) allows a router to restart its routing protocol processes while maintaining the

existing routing information. This minimizes disruptions during restarts, allowing for a smoother recovery process and maintaining network availability.

Question: 260

What is a typical use case for Link Aggregation Groups (LAG) in high availability networks?

A. To aggregate multiple Layer 3 routes into a single logical route
B. To combine multiple physical links between switches for bandwidth and redundancy
C. To enable fault tolerance for routing protocols
D. To segment traffic into different virtual networks based on IP addresses

Answer: B
Explanation:

LAG is used to combine multiple physical links between switches, increasing the total bandwidth and providing redundancy. If one link fails, the traffic is automatically rerouted over the remaining links,

ensuring high availability.

Question: 261

What is the benefit of using Bidirectional Forwarding Detection (BFD) in a network?

A. BFD reduces the complexity of routing protocols
B. BFD aggregates multiple paths into a single logical link
C. BFD allows for faster convergence by quickly detecting link failures
D. BFD dynamically adjusts IP address assignments to improve routing

Answer: C
Explanation:

BFD allows for faster convergence by providing rapid detection of link failures. It reduces the time it takes for routing protocols to react to changes in the network, helping to maintain high availability and minimize downtime.

Question: 262

Which of the following is a key advantage of using a Virtual Chassis configuration in a high availability setup?

A. It increases the capacity of routing protocols to handle large amounts of traffic
B. It improves the speed of network convergence after a link failure
C. It eliminates the need for redundant devices in the network
D. It reduces the number of devices in the network by combining multiple switches into one logical device

Answer: D
Explanation:

A Virtual Chassis configuration allows multiple switches to function as a single logical device, simplifying network management and providing failover capabilities. This setup improves high availability by centralizing control and reducing the impact of hardware failures.

Question: 263

Which of the following are benefits of using Link Aggregation Groups (LAG) in high availability configurations? (Choose two)

A. Provides redundancy by distributing traffic across multiple physical links
B. Increases network performance by aggregating multiple IP addresses
C. Helps ensure continuous connectivity by automatically rerouting traffic if a link fails
D. Reduces the need for dynamic routing protocols

Answer: A, C

Explanation:

LAG provides redundancy by aggregating multiple physical links into one logical link, which increases bandwidth and ensures that traffic continues to flow even if one link fails. It improves overall network performance and resilience.

Question: 264

Which of the following components are monitored by Bidirectional Forwarding Detection (BFD) to detect link failures? (Choose two)

A. BFD monitors Layer 2 traffic to detect failures
B. BFD checks the status of forwarding paths between routers
C. BFD tracks the status of IP routing tables
D. BFD sends packets between devices to verify that the path is operational

Answer: B, D

Explanation:

BFD monitors the status of forwarding paths between routers and sends control packets to verify that the path is still operational. It helps quickly detect failures and reroute traffic to maintain high availability.

Question: 265

Which of the following high availability features provide redundancy and failover in case of hardware or link failures? (Choose two)

A. Routing Information Protocol (RIP)
B. Graceful Restart (GR)
C. Virtual Chassis
D. Link Aggregation Groups (LAG)

Answer: C, D

Explanation:

LAG provides redundancy by combining multiple physical links into a single logical link, and Virtual Chassis enables multiple physical switches to act as one logical device, both of which improve network resilience and failover capabilities in high availability scenarios.

Question: 266

Which of the following are key characteristics of Graceful Restart (GR) in high availability configurations? (Choose two)

A. GR allows a router to retain its routing information during a restart
B. GR reduces the complexity of link aggregation
C. GR minimizes downtime during the restart of routing protocols
D. GR enables automatic rerouting to backup links

Answer: A, C
Explanation:

Graceful Restart (GR) allows a router to retain its routing information during a restart, minimizing downtime and reducing the impact of restarts on network performance and stability.

Question: 267

Which of the following configurations are typically used to monitor and ensure high availability? (Choose two)

A. Monitoring interface status using SNMP (Simple Network Management Protocol)
B. Using redundancy protocols like VRRP (Virtual Router Redundancy Protocol)
C. Using QoS (Quality of Service) to prioritize traffic
D. Monitoring routing protocol neighbor relationships using BGP and OSPF

Answer: A, B Explanation:

SNMP can be used to monitor the status of interfaces, while redundancy protocols like VRRP provide automatic failover for routers, ensuring continuous connectivity in case of device failure.

Question: 268

Which of the following are true about Bidirectional Forwarding Detection (BFD) in high availability networks? (Choose two)

A. BFD works independently of Layer 3 protocols like OSPF and BGP
B. BFD improves the speed of route convergence by detecting link failures faster than routing protocols
C. BFD is used to detect physical layer link failures only
D. BFD operates between routers and can trigger fast rerouting when a failure is detected

Answer: B, D Explanation:

BFD improves the speed of route convergence by quickly detecting link failures, and it operates between routers to trigger fast rerouting, ensuring minimal downtime during link failures.

Question: 269

Which of the following OSPF router types participate in the exchange of routing information between OSPF areas? (Choose two)

A. Internal Router
B. Area Border Router (ABR)
C. Backbone Router
D. Autonomous System Boundary Router (ASBR)

Answer: B, C

Explanation:

Area Border Routers (ABRs) are responsible for routing information between OSPF areas, while Backbone Routers connect different OSPF areas through the backbone (Area 0) to ensure routing continuity across the network.

Question: 270

Which of the following high availability techniques are used to ensure seamless traffic forwarding during link or device failures? (Choose two)

A. Dynamic Routing Protocols (e.g., OSPF, BGP)
B. Layer 2 VPN
C. Bidirectional Forwarding Detection (BFD)
D. Virtual Chassis

Answer: C, D

Explanation:

Virtual Chassis provides redundancy by allowing multiple switches to act as a single logical unit, ensuring high availability during a device failure. BFD helps ensure high availability by detecting link failures quickly and rerouting traffic without significant delay.

Question: 271

Which of the following components can be used to troubleshoot and ensure high availability in a network? (Choose two)

A. SNMP traps for device status monitoring
B. Spanning Tree Protocol (STP) to prevent loops
C. Using ping to monitor Layer 3 connectivity
D. Using debug commands to view real-time protocol status

Answer: A, D

Explanation:

SNMP traps provide real-time monitoring of device status, which can help detect issues and ensure high availability. Using debug commands allows for detailed inspection of protocol behavior and can assist in troubleshooting network issues in high availability setups.