Number: JN0-280
Passing Score: 800
Time Limit: 120
File Version: 5.0

Exam Code: JN0-280

Exam Name: Data Center, Associate

Exam A

QUESTION 1
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.

Correct Answer: B
Section:
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 2
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.

Correct Answer: A, C, E
Section:
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
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 3
Which signaling protocol is used for EVPN?

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

Correct Answer: D
Section:
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 4
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

Correct Answer: D
Section:
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 5
Exhibit:

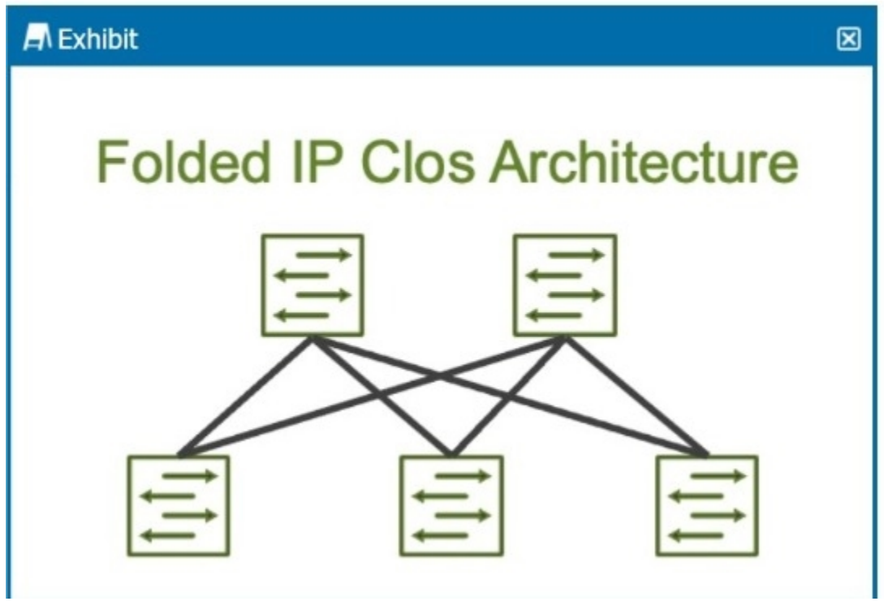

How many stages are shown in the exhibit?

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

Correct Answer: D
Section:
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 Explanation:
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 6
Exhibit:

Exhibit

```
[edit routing-options]
user@Router# show
static {
    route 0.0.0.0/0 {
        next-hop 172.25.11.254;
        qualified-next-hop 172.25.11.200 {
            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

Correct Answer: A
Section:
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 7
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.

Correct Answer: B, C
Section:
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 8
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.

Correct Answer: A
Section:
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=ReferenceBandwidthLinkBandwidth\text{Cost} = \frac{\text{Reference Bandwidth}}{\text{Link Bandwidth}}Cost=LinkBandwidthReferenceBandwidth
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 9
Which state in the adjacency process do OSPF routers check the MTU size?

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

Correct Answer: B
Section:
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 10
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

Correct Answer: B
Section:
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 11
What are two consequences of having all network devices in a single collision domain? (Choose two.)

A. The amount of network resource consumption does not change.
B. The chance of packet collision is decreased.
C. The chance of packet collision is increased.
D. The amount of network resource consumption is increased.

Correct Answer: C, D
Section:
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 12
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.

Correct Answer: B
Section:
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 13
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)

Correct Answer: A, B, E
Section:
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 14
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.

Correct Answer: D
Section:
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 15
Which statement is correct about aggregate routes?

A. The default next hop is discard.
B. The default next hop is readvertise.
C. The default next hop is resolve.
D. The default next hop is reject.

Correct Answer: D
Section:
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 Explanation:
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 16
Which Junos OS routing table stores IPv6 addresses?

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

Correct Answer: D
Section:
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 Explanation:
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 17
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

Correct Answer: C
Section:
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 18
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.

Correct Answer: C, D

Section:
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 19
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.

Correct Answer: C
Section:
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 20
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.

Correct Answer: D
Section:
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 21
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.

Correct Answer: A
Section:
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 22
A routing policy has been created to advertise OSPF routes in BGP. Which statement is correct in this scenario?

A. Apply the policy as an export policy within BGP.
B. Apply the policy as an export policy within OSPF.
C. Apply the policy as an import policy within BGP.
D. Apply the policy as an import policy within OSPF.

Correct Answer: A
Section:
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 23
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

Correct Answer: A, B
Section:
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 24
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

Correct Answer: A
Section:
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 25
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.

Correct Answer: C
Section:
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 26
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

Correct Answer: C

Section:
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 27
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.

Correct Answer: B, D
Section:
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 28
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.

Correct Answer: A, C
Section:
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 29
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

Correct Answer: A
Section:
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 30
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.

Correct Answer: C, D
Section:
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 31
What information in the Ethernet header is used to populate the bridging table?

A. destination address
B. source address
C. type
D. protocol

Correct Answer: B
Section:
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 32
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.

Correct Answer: B, C
Section:
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.