Juniper.JN0-351.v2026-03-02.q75

Exam Code:	JN0-351
Exam Name:	Enterprise Routing and Switching, Specialist (JNCIS-ENT)
Certification Provider:	Juniper
Free Question Number:	75
Version:	v2026-03-02
# of views:	107
# of Questions views:	750
https://www.freeqas.com/qa/Juniper/JN0-351/Juniper.JN0-351.v2026-03-02.q75.html	

NEW QUESTION: 1
You are an operator for a network running IS-IS. Two routers are failing to form an adjacency.
What are two reasons for this problem? (Choose two.)
A. There are mismatched router IDs on the L2 routers.
B. There is no configured ISO address on any IS-IS interface.
C. There is a mismatched area ID between the L2 routers.
D. The family iso configuration is missing from the adjacency interface.
Answer: B,D (LEAVE A REPLY)
The two reasons for the failure to form an adjacency in a network running IS-IS could be:
B) There is no configured ISO address on any IS-IS interface. IS-IS requires each router interface to have an ISO address configured. Without this address, the routers cannot form an adjacency.
D) The family iso configuration is missing from the adjacency interface. The 'family iso' configuration is essential for IS-IS to function correctly. If this configuration is missing from the adjacency interface, it could prevent the formation of an adjacency.
These explanations are based on the Enterprise Routing and Switching Specialist (JNCIS-ENT) documents and learning resources available at Juniper Networks.

NEW QUESTION: 2
You deployed a new EX Series switch with DHCP snooping enabled and you do not see any entries in the snooping databases for an interface. Which two Juniper configurations for that interface caused this issue?
(Choose two.)
A. The interface is configured as a disabled port.
B. MAC limiting is enabled on the interface.
C. The interface is configured as a trunk port.
D. Dynamic ARP inspection is enabled on the interface.
Answer: (SHOW ANSWER)

A is correct because the interface is configured as a disabled port. A disabled port does not forward any traffic, including DHCP packets. Therefore, DHCP snooping cannot learn any MAC addresses or lease information from a disabled port1.
C is correct because the interface is configured as a trunk port. By default, all trunk ports on the switch are trusted for DHCP snooping2. This means that DHCP snooping does not inspect or filter any DHCP packets received on a trunk port. Therefore, DHCP snooping does not add any entries to the snooping database for a trunk port2.

NEW QUESTION: 3
A new network requires multiple topology support. You decide to use IS-IS in this situation. Which three protocol topologies are supported in this scenario? (Choose three.)
A. IPsec
B. anycast
C. IPv6
D. multicast
E. IPv4
Answer: (SHOW ANSWER)
Explanation
IS-IS (Intermediate System to Intermediate System) is a routing protocol that is designed to move information efficiently within a computer network12. It supports multiple protocol topologies, including IPv4, IPv6, and multicast12. Therefore, options C, E, and D are correct.

NEW QUESTION: 4
Referring to the exhibit, why is the route for 10.5.5.5 hidden?

```
user@host> show route hidden detail
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
127.0.0.1/32 (1 entry, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Hidden Martian Int>
                Local AS:     1
                Age: 4:27:37
                Task: IF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.5.5.5/32 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Route Distinguisher: 10.4.4.4:4
                Next hop type: Unusable
                Next-hop reference count: 6
                State: <Secondary Hidden Int Ext>
                Local AS:     1 Peer AS:     1
                Age: 3:45:09
                Task: BGP_1.10.4.4.4+2493
                AS path: 100 I
                Communities: target:1:999
                VPN Label: 100064
                Localpref: 100
                Router ID: 10.4.4.4
                Primary Routing Table bgp.l3vpn.0
```


A. It has an invalid community.

B. The next hop cannot be resolved.

C. It is a martian route.

D. It is an L3VPN route.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 5

You are receiving multiple BGP routes from an upstream neighbor and only want to advertise a single summarized prefix to your internal OSPF neighbors. This route should only be advertised when you are receiving these BGP routes from this neighbor.

In this scenario, which type of route should you create?

A. aggregate route

B. static route using the resolve feature
C. generate route
D. static route using qualified next hops
Answer: A (LEAVE A REPLY)
Explanation
In this scenario, you should create an 1. Aggregate routes are used for advertising summarized network prefixes1. They help minimize the number of routing tables in an IP network by consolidating selected multiple routes into a single route advertisement1. This approach is in contrast to non-aggregation routing, in which every routing table contains a unique entry for each route1.
Therefore, option A is correct. Options B, C, and D are not correct because:
Static route using the resolve feature: This type of route uses the resolve feature to install a static route in the routing table only if a specific condition is met1. However, it does not provide the capability to summarize multiple routes into a single prefix.
Generate route: This type of route generates a route that is always present in the routing table and can be used to summarize routes. However, it does not have the capability to only advertise the route when specific BGP routes are being received from a neighbor1.
Static route using qualified next hops: This type of route allows for the specification of multiple next-hop addresses for a static route1. However, it does not provide the capability to summarize multiple routes into a single prefix.

NEW QUESTION: 6
You enable persistent MAC learning on your Juniper switch. In this scenario, which statement is correct?
A. You can enable persistent MAC learning on an interface where MAC learning is disabled.
B. You can enable persistent MAC learning on an interface that is part of a redundant trunk group.
C. You can only enable persistent MAC learning on an interface in access mode.
D. You can only enable persistent MAC learning on an interface on which 802.1x authentication is configured.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 7
You are deploying a trunk port with multiple VLANs. In this scenario, which statement is true about the native VLAN?
A. The native VLAN is the highest tagged VLAN ID on the port.
B. The native VLAN must be the default VLAN on the port.
C. The native VLAN is automatically configured through DHCP.
D. The native VLAN is an untagged VLAN ID on the port.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 8
Which two statements are correct about using firewall filters on EX Series switches? (Choose two.)
A. You can deploy only stateless firewall filters on an EX Series switch.
B. You can only apply firewall filters to Layer 2 traffic on an EX Series switch.
C. You can apply firewall filters to both Layer 2 and Layer 3 traffic on an EX Series switch.
D. You can deploy both stateless and stateful firewall filters on an EX Series switch.
Answer: (SHOW ANSWER)
A is correct because you can deploy only stateless firewall filters on an EX Series switch. A stateless firewall filter is a filter that evaluates each packet individually based on the header information, such as source and destination addresses, protocol, and port numbers. A stateless firewall filter does not keep track of the state or context of a packet flow, such as the sequence number, flags, or session information. EX Series switches support only stateless firewall filters, which are also called access control lists (ACLs) or packet filters.
C is correct because you can apply firewall filters to both Layer 2 and Layer 3 traffic on an EX Series switch. Layer 2 traffic is traffic that is switched within a VLAN or a bridge domain, while Layer 3 traffic is traffic that is routed between VLANs or networks. EX Series switches support three types of firewall filters: port (Layer 2) firewall filters, VLAN firewall filters, and router (Layer 3) firewall filters. You can apply these filters to different interfaces and directions to control the traffic entering or exiting the switch.

NEW QUESTION: 9
Exhibit.

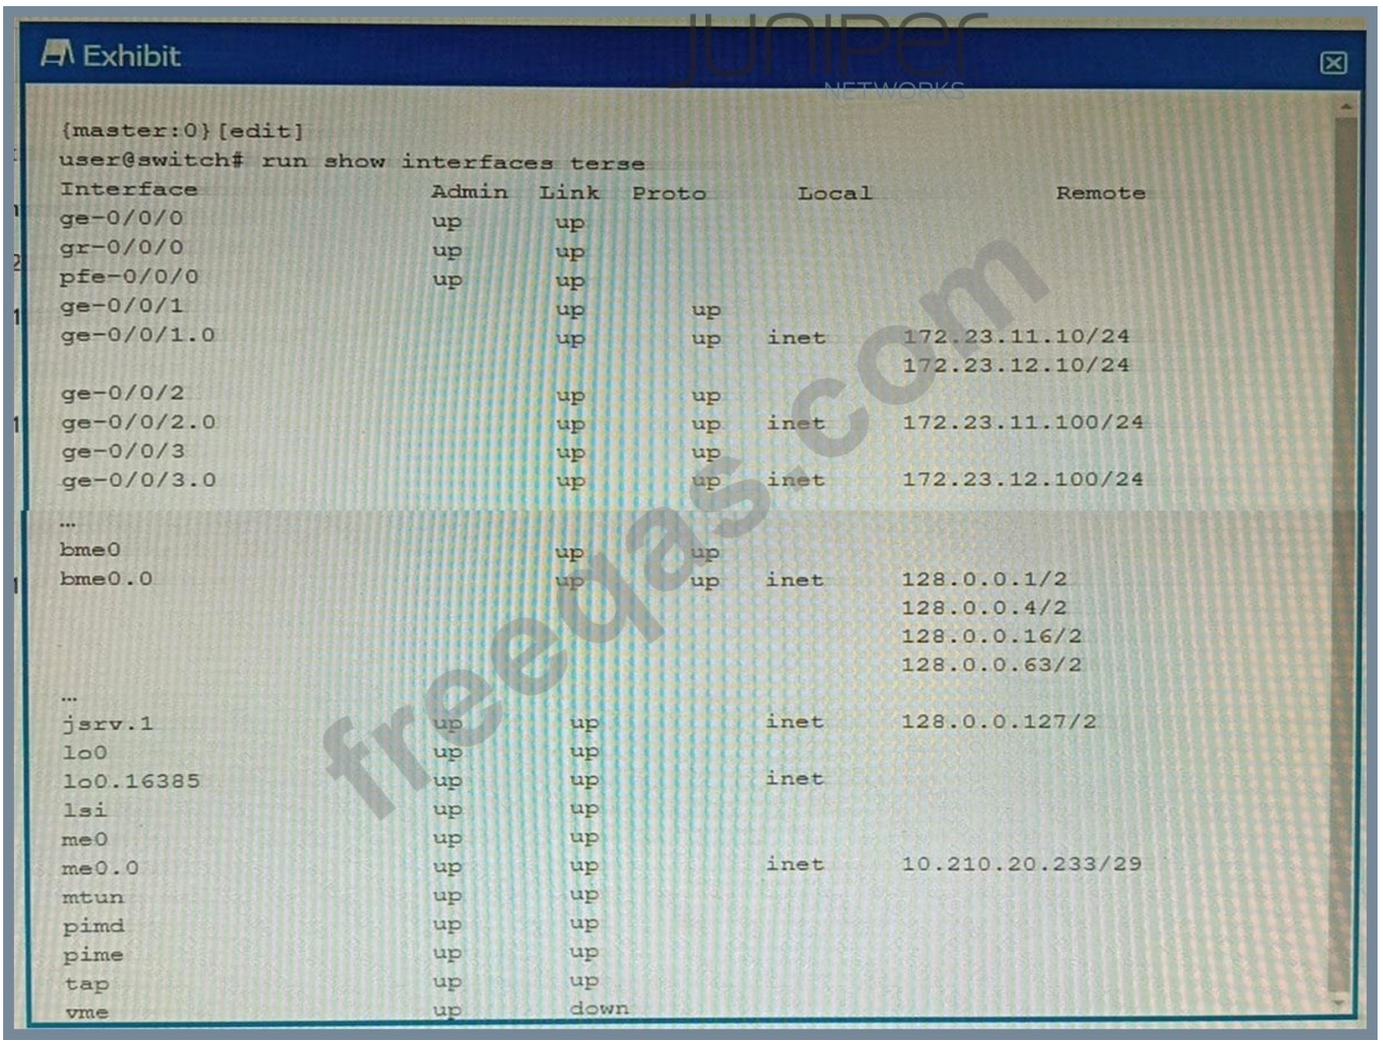

What is the management IP address of the device shown in the exhibit?

A. 10.210.20.233

B. 172.23.12.100

C. 128.0.0.1

D. 172.23.11.10

Answer: (SHOW ANSWER)

Explanation

The management IP address of a device is the IP address that is used to access the device for configuration and monitoring purposes. It is usually assigned to a dedicatedmanagement interface that is separate from the data interfaces. The management interface can be accessed via SSH, Telnet, HTTP, or other protocols.

In the exhibit, the list of interfaces and their statuses shows that the management interface isme0. This interface has an admin status ofup, a protocol status ofinet, a local address of172.23.12.100/24, and a remote address ofunspecified. This means that the me0 interface is active, has an IPv4 address assigned, and is not connected to another device.

Therefore, the management IP address of the device shown in the exhibit is172.23.12.100.

References:

[Management Interfaces Overview] : [Displaying Interface Status Information]

NEW QUESTION: 10
A host is connected to your EX Series switch using an interface with persistent MAC learning enabled. When you connect the host to another interface on the same switch, all traffic generated by the host is dropped. Which action should you perform to solve the problem?
A. Enter the clear ethernet switching table persistent learningcommand.
B. Configure the interface as a trunk port.
C. Enter the clear dhcp snooping databasecommand.
D. Disable dynamic ARP inspection.
Answer: (SHOW ANSWER)

NEW QUESTION: 11
Which two statements about redundant trunk groups on EX Series switches are correct? (Choose two.)
A. Redundant trunk groups use spanning tree to provide loop-free redundant uplinks.
B. Redundant trunk groups load balance traffic across two designated uplink interfaces.
C. Layer 2 control traffic is permitted on the secondary link.
D. If the active link fails, then the secondary link automatically takes over.
Answer: C,D (LEAVE A REPLY)
C is correct because Layer 2 control traffic is permitted on the secondary link of a redundant trunk group (RTG) on EX Series switches. Layer 2 control traffic includes protocols such as LLDP, LACP, and STP, which are used to exchange information and coordinate actions between switches1. According to the Juniper Networks documentation2, Layer 2 control traffic is allowed to pass through both the active and the secondary links of an RTG, but data traffic is only forwarded through the active link. This allows the switches to maintain their Layer 2 adjacencies and monitor the link status on both links.
D is correct because if the active link fails, then the secondary link automatically takes over in an RTG on EX Series switches. An RTG consists of two trunk links: an active or primary link, and a secondary or backup link2. The active link is used to forward data traffic, while the secondary link is in standby mode. If the active link fails or becomes unavailable, the secondary link immediately transitions to a forwarding state and takes over the data traffic without waiting for normal STP convergence2. This provides fast recovery and redundancy for the network.

NEW QUESTION: 12
Which two statements are true about the default VLAN on Juniper switches? (Choose two.)
A. The default VLAN is set to a VLAN ID of 1 by default
B. The default VLAN ID is not assigned to any interface.
C. The default VLAN ID is not visible.
D. The default VLAN ID can be changed.
Answer: A,D (LEAVE A REPLY)

On Juniper switches, the default VLAN is set to a VLAN ID of 1 by default. This means that all interfaces on the switch are members of VLAN until they are specifically assigned to another VLAN. Therefore, option A is correct.
The default VLAN ID can be changed. This allows network administrators to configure the switch to use a different VLAN as the default, if necessary. Therefore, option D is correct.

NEW QUESTION: 13
In RSTP, which three port roles are associated with the discarding state? (Choose three.)
A. root
B. backup
C. alternate
D. disabled
E. designated
Answer: B,C,D (LEAVE A REPLY)
Explanation
In Rapid Spanning Tree Protocol (RSTP), there are several port roles that determine the behavior of the port in the spanning tree123. The roles include root, designated, alternate, backup, and disabled123.
The discarding state is associated with the backup, alternate, and disabled roles123. In a stable topology with consistent port roles throughout the network, RSTP ensures that every root port and designated port immediately transition to the forwarding state while all alternate and backup ports are always in the discarding state2. Disabled ports are also in the discarding state3.
Therefore, options B, C, and D are correct.

NEW QUESTION: 14
You are asked to configure a redundant trunk group (RTG). Which two requirements would accomplish this task? (Choose two.)
A. The RTG must be participating in a Spanning Tree topology.
B. Both interfaces of an RTG must be configured to service the same VLANS.
C. Interfaces within an RTG can be configured as access ports.
D. Interfaces within an RTG must be configured as trunk ports.
Answer: B,D (LEAVE A REPLY)

NEW QUESTION: 15
You are using tunnels in your network. It is important that the routes be specific enough to ensure that the tunnels are established. You need to prevent a route that is new to the network from being used. In this scenario, which type of route should be used?
A. aggregate
B. multicast
C. static
D. anycast

Answer: C (LEAVE A REPLY)

NEW QUESTION: 16
Which two statements about BGP facilitate the prevention of routing loops between two autonomous systems?
(Choose two.)
A. EBGP routers will append their AS number when advertising routes to their neighbors.
B. EBGP routers will only accept routes that contain their own AS number in the AS_PATH.
C. EBGP routers will drop routes that contain their own AS number in the AS_PATH
D. EBGP routers will prepend their AS number when advertising routes to their neighbors
Answer: A,C (LEAVE A REPLY)
Explanation
BGP (Border Gateway Protocol) is a protocol designed to exchange routing and reachability information among autonomous systems (AS) on the internet1.
Option A is correct. When an EBGP router advertises routes to its neighbors, it appends its AS number to the AS_PATH attribute1. This is a key mechanism in BGP to prevent routing loops1.
Option C is correct. BGP has a built-in loop prevention mechanism whereby if a BGP router detects its own AS in the AS_PATH attribute, it will drop the prefix and will not continue to advertise it2. This helps to prevent routing loops2.
Option B is incorrect. EBGP routers do not accept routes that contain their own AS number in the AS_PATH2. Instead, they drop such routes as part of the loop prevention mechanism2.
Option D is incorrect. While it's true that EBGP routers append their AS number when advertising routes, they do not prepend their AS number1. The term "prepend" in BGP usually refers to a technique used to influence path selection by artificially lengthening the AS_PATH3.

NEW QUESTION: 17
You want to enable redundancy for the EBGP peering between the two routers shown in the exhibit.
Which three actions will you perform in this scenario? (Choose three.)

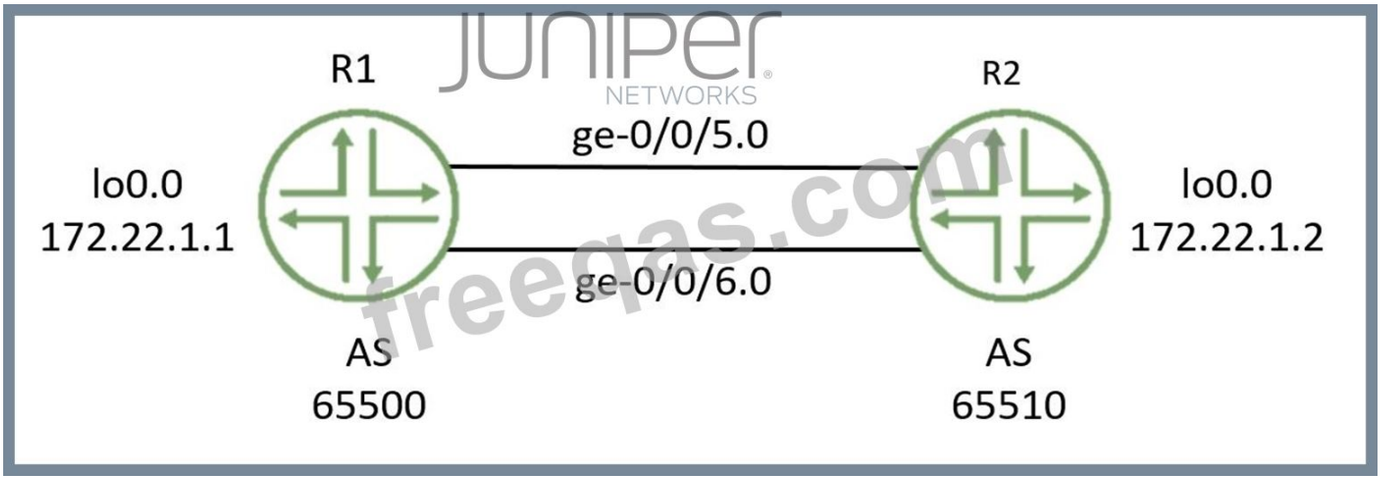

A. Configure BGP multihop.
B. Configure loopback interface peering.
C. Configure routes for the peer loopback interface IP addresses.
D. Configure an MD5 peer authentication.
E. Configure a cluster ID.
Answer: A,B,C (LEAVE A REPLY)
A is correct because you need to configure BGP multihop to enable redundancy for the EBGP peering between the two routers. BGP multihop is a feature that allows BGP peers to establish a session over multiple hops, instead of requiring them to be directly connected. By default, EBGP peers use a time-to-live (TTL) value of 1 for their packets, which means that they can only reach adjacent neighbors. However, if you configure BGP multihop with a higher TTL value, you can allow EBGP peers to communicate over multiple routers in between. This can provide redundancy in case of a link failure or a router failure between the EBGP peers.
B is correct because you need to configure loopback interface peering to enable redundancy for the EBGP peering between the two routers. Loopback interface peering is a technique that uses loopback interfaces as the source and destination addresses for BGP sessions, instead of physical interfaces. Loopback interfaces are virtual interfaces that are always up and reachable as long as the router is operational. By using loopback interface peering, you can avoid the dependency on a single physical interface or link for the BGP session, and use multiple paths to reach the loopback address of the peer. This can provide redundancy and load balancing for the EBGP peering.
C is correct because you need to configure routes for the peer loopback interface IP addresses to enable redundancy for the EBGP peering between the two routers. Routes for the peer loopback interface IP addresses are necessary to ensure that the routers can reach each other's loopback addresses over multiple hops. You can use static routes or dynamic routing protocols to advertise and learn the routes for the peer loopback interface IP addresses. Without these routes, the routers will not be able to establish or maintain the BGP session using their loopback interfaces.

NEW QUESTION: 18
Referring to the output shown in the exhibit, which two statements are true?

```
user@host> show route 0/0 exact detail
inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
0.0.0.0/0 (1 entry, 1 announced)
        *Aggregate Preference: 130
                Next hop type: Router, Next hop index: 546
                Next-hop reference count: 4
Next hop: 172.30.25.1 via ge-0/0/1.100, selected
State: <Active Int Ext>
Local AS: 65400
Age: 1:03:46
Task: Aggregate
Announcement bits (2): 0-KRT 2-OSPF
AS path: I
Flags: Generate Depth: 0 Active
Contributing Routes (1):
10.0.0.0/16 proto BGP
```

A. The route is an aggregate route

B. The route is a generate route

C. The route is active

D. The route is not active

Answer: (SHOW ANSWER)

NEW QUESTION: 19

You want to redistribute the static routes into BGP, but they do not show as BGP routes.

What causes this result to occur?

```
user@router> show route
inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0           *[Static/5] 9w0d 02:07:24
                     >  to 10.1.1.254 via xe-0/0/0.0
10.1.1.0/24         *[Direct/0] 9w0d 02:07:24
                     >  via xe-0/0/0.0
10.16.0.0/24        *[Direct/0] 9w0d 02:07:24
                     >  via xe-0/0/2.0
10.123.255.64/26 *[Static/5] 9w0d 02:07:24
                     >  via xe-0/0/2.0
172.16.0.0/24 *[Direct/0] 9w0d 02:07:24
                     >  via xe-0/0/3.0
...
```

A. The no-readvertisefeature is specified.

B. The as-pathparameter must be specified.

C. The static routes do not have a valid next hop.

D. A community string must be configured.

Answer: C (LEAVE A REPLY)

NEW QUESTION: 20
You are troubleshooting a BGP routing issue between your network and a customer router and are reviewing the BGP routing policies. Which two statements are correct in this scenario? (Choose two.)
A. Export policies are applied to routes in the RIB-In table.
B. Import policies are applied to routes in the RIB-Local table.
C. Import policies are applied after the RIB-In table.
D. Export policies are applied after the RIB-Local table.
Answer: C,D (LEAVE A REPLY)
Explanation
In BGP, routing policies are used to control the flow of routing information between BGP peers1. Option C suggests that import policies are applied after the RIB-In table. This is correct because import policies in BGP are applied to routes that are received from a BGP peer, before they are installed in the local BGP Routing Information Base (RIB-In)1. The RIB-In is a database that stores all the routes that are received from all peers1.
Option D suggests that export policies are applied after the RIB-Local table. This is correct because export policies in BGP are applied to routes that are being advertised to a BGP peer, after they have been selected from the local BGP Routing Information Base (RIB-Local)1. The RIB-Local is a database that stores all the routes that the local router is using1.
Therefore, options C and D are correct.

NEW QUESTION: 21
Exhibit.

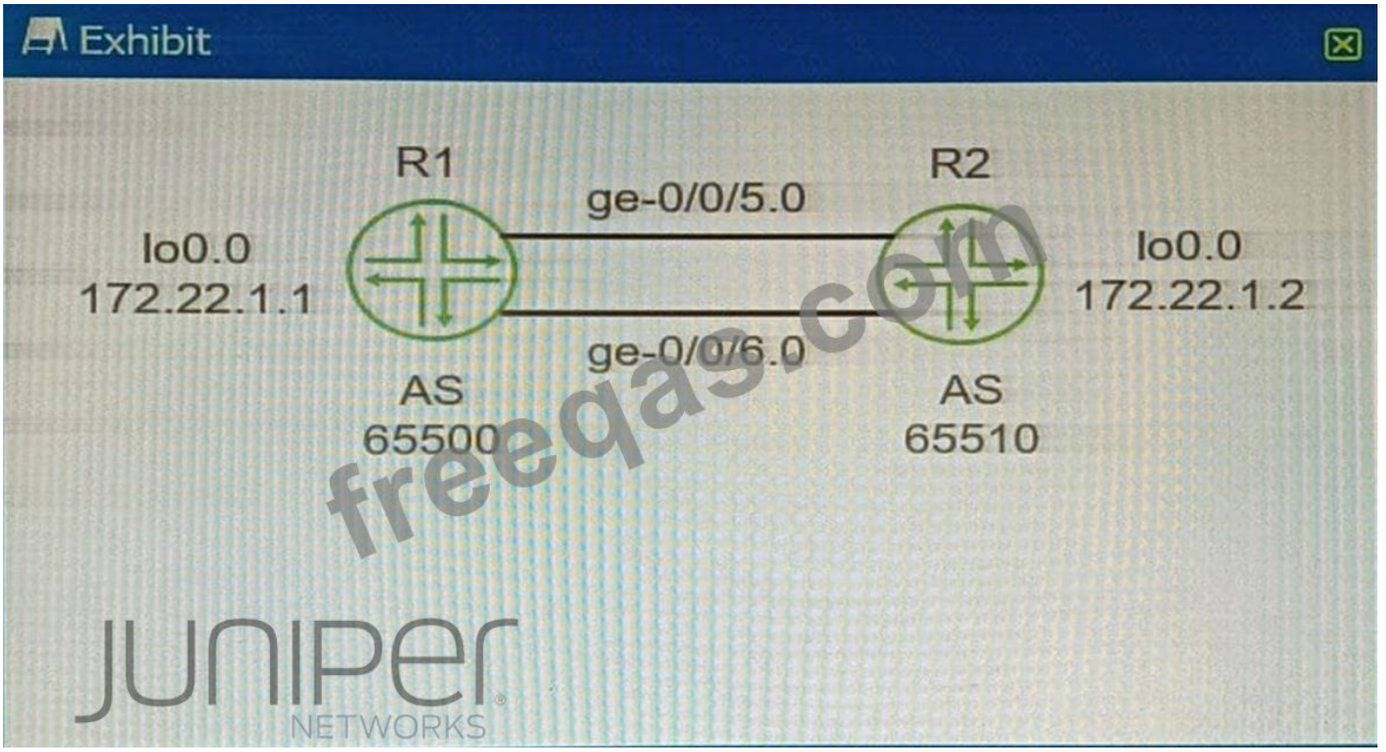

You want to enable redundancy for the EBGP peering between the two routers shown in the exhibit. Which three actions will you perform in this scenario? (Choose three.)

A. Configure BGP multihop.
B. Configure loopback interface peering.
C. Configure routes for the peer loopback interface IP addresses.
D. Configure an MD5 peer authentication.
E. Configure a cluster ID.
Answer: (SHOW ANSWER)
A is correct because you need to configure BGP multihop to enable redundancy for the EBGP peering between the two routers. BGP multihop is a feature that allows BGP peers to establish a session over multiple hops, instead of requiring them to be directly connected1. By default, EBGP peers use a time-to-live (TTL) value of 1 for their packets, which means that they can only reach adjacent neighbors1. However, if you configure BGP multihop with a higher TTL value, you can allow EBGP peers to communicate over multiple routers in between1. This can provide redundancy in case of a link failure or a router failure between the EBGP peers.
B is correct because you need to configure loopback interface peering to enable redundancy for the EBGP peering between the two routers. Loopback interface peering is a technique that uses loopback interfaces as the source and destination addresses for BGP sessions, instead of physical interfaces2. Loopback interfaces are virtual interfaces that are always up andreachable as long as the router is operational2. By using loopback interface peering, you can avoid the dependency on a single physical interface or link for the BGP session, and use multiple paths to reach the loopback address of the peer2. This can provide redundancy and load balancing for the EBGP peering.
C is correct because you need to configure routes for the peer loopback interface IP addresses to enable redundancy for the EBGP peering between the two routers. Routes for the peer loopback interface IP addresses are necessary to ensure that the routers can reach each other's loopback addresses over multiple hops2. You can use static routes or dynamic routing protocols to advertise and learn the routes for the peer loopback interface IP addresses2. Without these routes, the routers will not be able to establish or maintain the BGP session using their loopback interfaces.

NEW QUESTION: 22
Referring to the exhibit, which router will become the OSPF BDR if all routers are powered on at the same time?

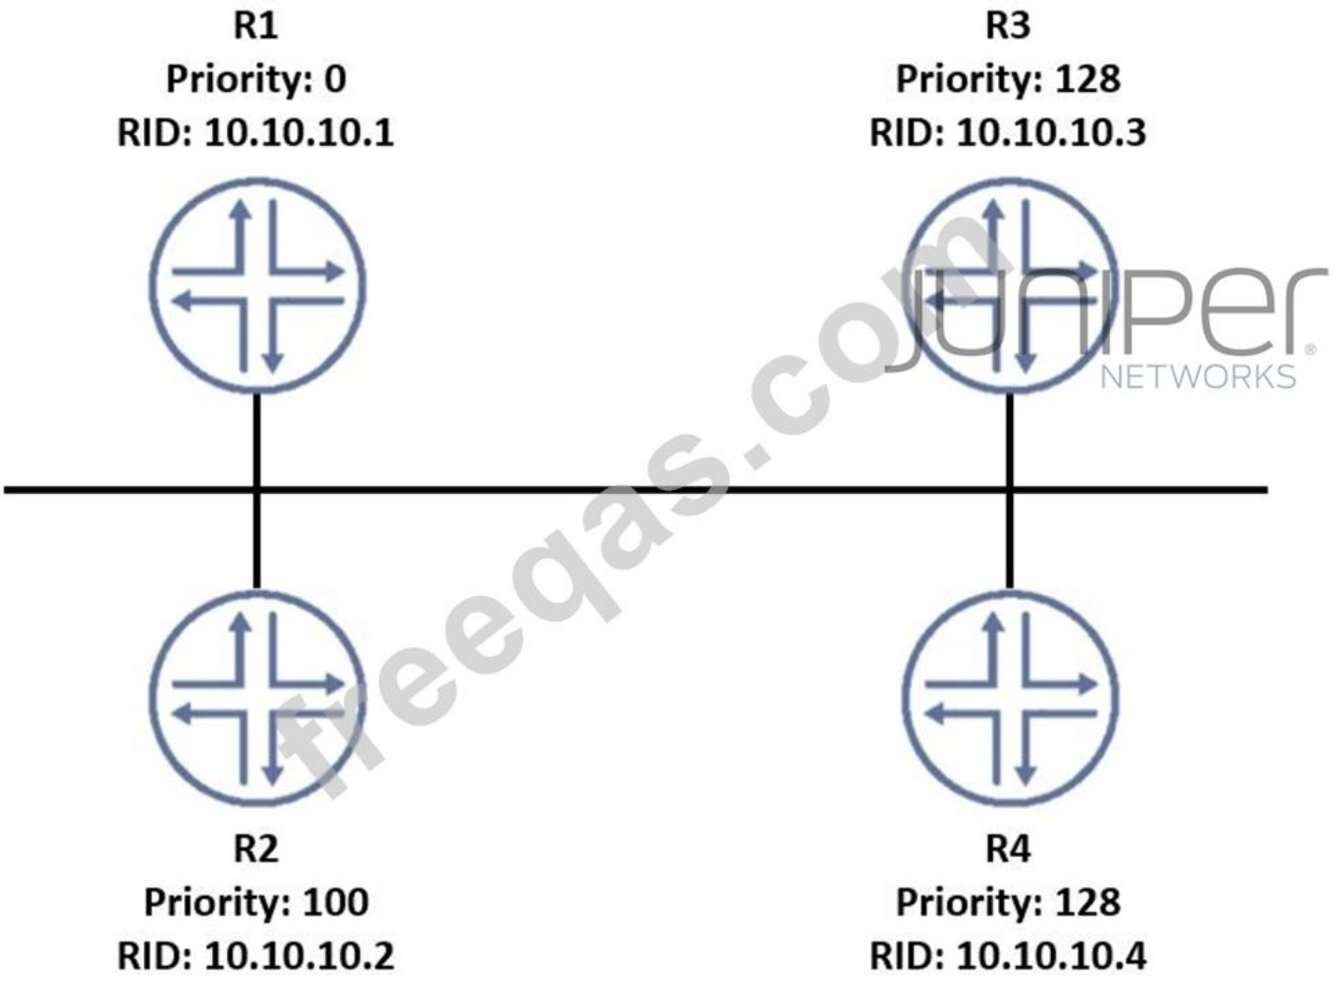

A. R4
B. R1
C. R3
D. R2
Answer: C (LEAVE A REPLY)
Priority set to 0 will never become DR or BDR so R1 is excluded.
Higher priority is better in OSPF so R2 is lowest than both R3 and R4.
R3 and R4 has the same priority so RID will decide.
R4 has the highest RID so it will become DR and R3 will then become BDR.

NEW QUESTION: 23
Which three protocols support BFD? (Choose three.)
A. RSTP
B. BGP
C. OSPF
D. LACP
E. FTP
Answer: B,C,D (LEAVE A REPLY)

Explanation
BFD is a protocol that can be used to quickly detect failures in the forwarding path between two adjacent routers or switches. BFD can be integrated with various routing protocols and link aggregation protocols to provide faster convergence and fault recovery.
According to the Juniper Networks documentation, the following protocols support BFD on Junos OS devices1:
BGP: BFD can be used to monitor the connectivity between BGP peers and trigger a session reset if a failure is detected. BFD can be configured for both internal and external BGP sessions, as well as for IPv4 and IPv6 address families2.
OSPF: BFD can be used to monitor the connectivity between OSPF neighbors and trigger a state change if a failure is detected. BFD can be configured for both OSPFv2 and OSPFv3 protocols, as well as for point-to-point and broadcast network types3.
LACP: BFD can be used to monitor the connectivity between LACP members and trigger a link state change if a failure is detected. BFD can be configured for both active and passive LACP modes, as well as for static and dynamic LAGs4.
Other protocols that support BFD on Junos OS devices are:
IS-IS: BFD can be used to monitor the connectivity between IS-IS neighbors and trigger a state change if a failure is detected. BFD can be configured for both level 1 and level 2 IS-IS adjacencies, as well as for point-to-point and broadcast network types.
RIP: BFD can be used to monitor the connectivity between RIP neighbors and trigger a route update if a failure is detected. BFD can be configured for both RIP version 1 and version 2 protocols, as well as for IPv4 and IPv6 address families.
VRRP: BFD can be used to monitor the connectivity between VRRP routers and trigger a priority change if a failure is detected. BFD can be configured for both VRRP version 2 and version 3 protocols, as well as for IPv4 and IPv6 address families.
The protocols that do not support BFD on Junos OS devices are:
RSTP: RSTP is a spanning tree protocol that provides loop prevention and rapid convergence in layer 2 networks. RSTP does not use BFD to detect link failures, but relies on its own hello mechanism that sends BPDU packets every 2 seconds by default.
FTP: FTP is an application layer protocol that is used to transfer files between hosts over a TCP connection. FTP does not use BFD to detect connection failures, but relies on TCP's own retransmission and timeout mechanisms.
References:
1: [Configuring Bidirectional Forwarding Detection] 2: [Configuring Bidirectional Forwarding Detection for BGP] 3: [Configuring Bidirectional Forwarding Detection for OSPF] 4: [Configuring Bidirectional Forwarding Detection for Link Aggregation Control Protocol] : [Configuring Bidirectional Forwarding Detection for IS-IS] : [Configuring Bidirectional Forwarding Detection for RIP] : [Configuring Bidirectional Forwarding Detection for VRRP] : [Understanding Rapid Spanning Tree Protocol] : [Understanding FTP]

NEW QUESTION: 24

Which statement is correct about the IS-IS ISO NET address?
A. An ISO NET address defined with a system ID of 0000.0000.0000 must be selected as the DIS.
B. An ISO NET address must be unique for each device in the network.
C. You can only define a single ISO NET address per device.
D. The Area ID must match on all devices within a L2 area.
Answer: (SHOW ANSWER)
An ISO NET address is a type of network address used by the IS-IS routing protocol. It identifies a point of connection to the network, such as a router interface, and is also called a Network Service Access Point (NSAP).
An ISO NET address consists of three parts: an area ID, a system ID, and a selector. The area ID identifies the IS-IS area to which the device belongs. The system ID uniquely identifies the device within the area. The selector identifies a specific service or function on the device, such as routing or management.
An ISO NET address must be unique for each device in the network, because it is used by IS-IS to establish adjacencies, exchange routing information, and compute shortest paths. If two devices have the same ISO NET address, they will not be able to communicate with each other or with other devices in the network. Therefore, it is important to assign different ISO NET addresses to each device in the network.

NEW QUESTION: 25
What is the default MAC age-out timer on an EX Series switch?
A. 30 minutes
B. 30 seconds
C. 300 minutes
D. 300 seconds
Answer: D (LEAVE A REPLY)
The default MAC age-out timer on an EX Series switch is 300 seconds. The MAC age-out timer is the maximum time that an entry can remain in the MAC table before it "ages out," or is removed. This configuration can influence efficiency of network resource use by affecting the amount of traffic that is flooded to all interfaces. When traffic is received for MAC addresses no longer in the Ethernet routing table, the router floods the traffic to all interfaces.

NEW QUESTION: 26
You are concerned about spoofed MAC addresses on your LAN. Which two Layer 2 security features should you enable to minimize this concern? (Choose two.)
A. dynamic ARP inspection
B. IP source guard
C. DHCP snooping
D. static ARP
Answer: A,C (LEAVE A REPLY)

A is correct because dynamic ARP inspection (DAI) is a Layer 2 security feature that prevents ARP spoofing attacks. ARP spoofing is a technique that allows an attacker to send fake ARP messages to associate a spoofed MAC address with a legitimate IP address. This can result in traffic redirection, man-in-the-middle attacks, or denial-of-service attacks. DAI validates ARP packets by checking the source MAC address and IP address against a trusted database, which is usually built by DHCP snooping. DAI discards any ARP packets that do not match the database or have invalid formats. C is correct because DHCP snooping is a Layer 2 security feature that prevents DHCP spoofing attacks. DHCP spoofing is a technique that allows an attacker to act as a rogue DHCP server and offer fake IP addresses and other network parameters to unsuspecting clients. This can result in traffic redirection, man-in-the-middle attacks, or denial-of-service attacks. DHCP snooping filters DHCP messages by classifying switch ports as trusted or untrusted. Trusted ports are allowed to send and receive any DHCP messages, while untrusted ports are allowed to send only DHCP requests and receive only valid DHCP replies from trusted ports. DHCP snooping also builds a database of MAC addresses, IP addresses, lease times, and binding types for each client.

NEW QUESTION: 27
You are a network operator who wants to add a second ISP connection and remove the default route to the existing ISP You decide to deploy the BGP protocol in the network.
What two statements are correct in this scenario? (Choose two.)
A. IBGP updates the next-hop attribute to ensure reachability within an AS.
B. IBGP peers advertise routes received from EBGP peers to other IBGP peers.
C. IBGP peers advertise routes received from IBGP peers to other IBGP peers.
D. EBGP peers advertise routes received from IBGP peers to other EBGP peers.
Answer: A,B (LEAVE A REPLY)
A is correct because IBGP updates the next-hop attribute to ensure reachability within an AS. This is because the next-hop attribute is the IP address of the router that advertises the route to a BGP peer. If the next-hop attribute is not changed by IBGP, it would be the IP address of an external router, which may not be reachable by all routers within the AS. Therefore, IBGP updates the next-hop attribute to the IP address of the router that received the route from an EBGP peer1. B is correct because IBGP peers advertise routes received from EBGP peers to other IBGP peers. This is because BGP follows the rule of advertising only the best route to a destination, and EBGP routes have a higher preference than IBGP routes. Therefore, IBGP peers advertise routes learned from an EBGP peer to all BGP peers, including both EBGP and IBGP peers1.

NEW QUESTION: 28
Referring to the exhibit, which two configuration changes must you apply for packets to reach from R1 to R3 using IS-IS? (Choose two.)

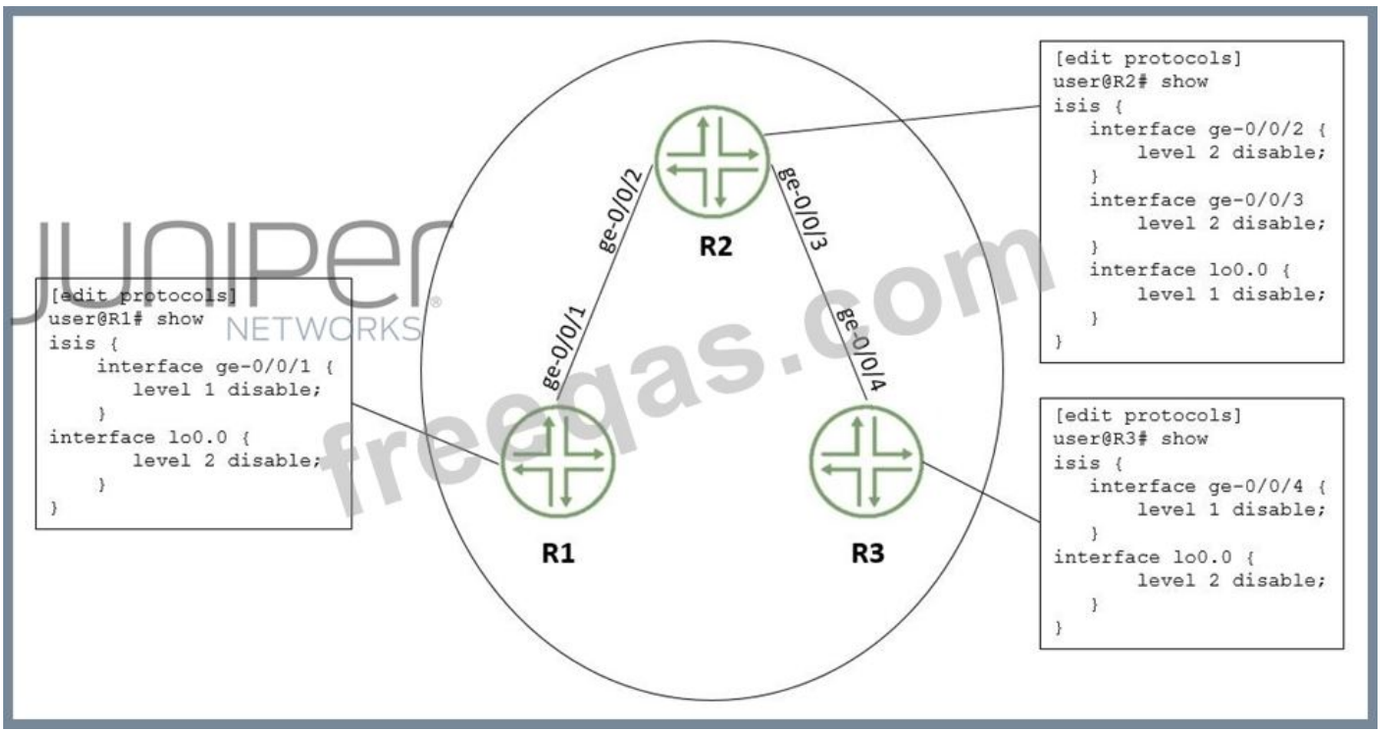

A. On R1, enable Level 1 on the ge-0/0/1 interface.
B. On R3 disable Level 2 on the ge-0/0/4 interface.
C. On R1, disable Level 2 on the ge-0/0/1 interface.
D. On R3 enable Level 1 on the ge-0/0/4 interface
Answer: A,D (LEAVE A REPLY)
A) On R1, enable Level 1 on the ge-0/0/1 interface. In IS-IS, both levels (Level 1 and Level 2) are enabled by default when you enable IS-IS on an interface1. Level 1 systems route within an area. If the destination is outside an area, Level 1 systems route toward a Level 2 system. Therefore, enabling Level 1 on the ge-0/0/1 interface on R1 would allow packets to reach from R1 to R3.
D) On R3 enable Level 1 on the ge-0/0/4 interface Similarly, enabling Level 1 on the ge-0/0/4 interface on R3 would allow packets to reach from R1 to R3. These explanations are based on the IS-IS configuration documents and learning resources available at Juniper Networks1 and Cisco.

NEW QUESTION: 29
You are deploying an EX Series switch with a Wi-Fi access point. The access point needs all untagged traffic to use a specific VLAN. In this scenario, which feature should you enable on the switch port?
A. flexible VLAN tagging
B. voice VLAN
C. VLAN tagging
D. native VLAN
Answer: D (LEAVE A REPLY)

NEW QUESTION: 30

What is the default keepalive time for BGP?

A. 10 seconds

B. 60 seconds

C. 30 seconds

D. 90 seconds

Answer: C (LEAVE A REPLY)

The default hold-time is 90 seconds, meaning that the default frequency for keepalive messages is 30 seconds.

https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/ref/statement/precision-timers-edit-protocols- bgp.html#:~:text=BGP%20on%20the%20local%20routing,keepalive%20messages%20is%2030

%20seconds

NEW QUESTION: 31

Which command solves the error that you see in the exhibit?

```
{master:0}[edit]
user@ex2# show interfaces ge-0/0/2
unit 0 {
    family ethernet-switching {
        interface-mode access;
        vlan {
            members [ default v11 ];
        }
    }
}

master:0}[edit]
user@ex2# commit
[edit interfaces ge-0/0/2 unit 0 family ethernet-switching]
  'vlan'
    Access interface can be part of only one vlan
error: configuration check-out failed

{master:0}[edit]
user@ex2#
```

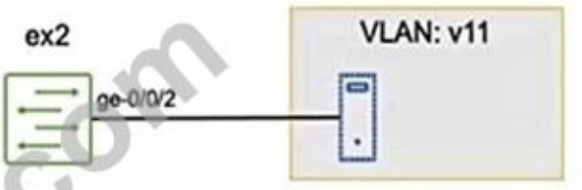

A. delete ge-0/0/2 unit 0 family ethernet-switching vlan members default

B. set ge-0/0/2 unit 0 family ethernet-switching vlan members [default v11]

C. delete ge-0/0/2 unit 0 family ethernet-switching vlan members all

D. set ge-0/0/2 unit 0 family ethernet-switching vlan members all

Answer: A (LEAVE A REPLY)

The error in the exhibit occurs because an access interface can only be a member of one VLAN, but it is currently configured with the "default" VLAN and is trying to be added to VLAN "v11." To resolve the issue, the correct command is:

delete ge-0/0/2 unit 0 family ethernet-switching vlan members default

This command removes the interface from the "default" VLAN, allowing it to be part of VLAN

"v11" as intended.

NEW QUESTION: 32
Which three actions are required for filter-based forwarding? (Choose three.)
A. Create a forwarding option for load balancing.
B. Create a RIB group for next-hop entries.
C. Create a firewall filter to match desired traffic.
D. Create routing policies for source addresses.
E. Create routing instances for individual destinations.
Answer: B,C,D (LEAVE A REPLY)

NEW QUESTION: 33
What is a purpose of using a spanning tree protocol?
A. to look up MAC addresses
B. to eliminate broadcast storms
C. to route IP packets
D. to tunnel Ethernet frames
Answer: B (LEAVE A REPLY)
A broadcast storm is a network condition where a large number of broadcast packets are sent and received by multiple devices, causing congestion and performance degradation1. A broadcast storm can occur when there are loops in the network topology, meaning that there are multiple paths between two devices2.
A spanning tree protocol is a network protocol that prevents loops from being formed when switches or bridges are interconnected via multiple paths. It does this by creating a logical tree structure that spans all the devices in the network, and disabling or blocking the links that are not part of the tree, leaving a single active path between any two devices3.
By eliminating loops, a spanning tree protocol also eliminates broadcast storms, as broadcast packets will not be forwarded endlessly along the looped paths. Instead, broadcast packets will be sent only along the tree structure, reaching each device once and avoiding congestion3.

NEW QUESTION: 34

You want to ensure traffic is routed through a GRE tunnel. In this scenario, which two statements will satisfy this requirement? (Choose two.)
A. Tunnel endpoints must have a route that directs traffic into the tunnel.
B. All intermediary devices must have a route to the tunnel endpoints.
C. Keepalives must be used on stateless tunneling protocols.
D. BFD must be used on the stateless tunneling protocols.
Answer: A,B (LEAVE A REPLY)
Option A is correct. For traffic to be sent through a GRE tunnel, there must be a route that directs the traffic into the tunnel. This is typically accomplished through the use of a static route or a dynamic routing protocol.
Option B is correct. All intermediary devices must have a route to the tunnel endpoints. In real-world scenarios, the tunnel endpoints for a tunnel going over the Internet must have globally reachable internet addresses. Otherwise, intermediate routers in the Internet cannot forward the tunneled packets.

NEW QUESTION: 35
You are asked to explore adding BGP route reflectors to allow for increased scale in your core network. Which statement is correct in this scenario?
A. Route reflectors are required whenever internal BGP (IBGP) sessions are needed.
B. The cluster ID is required to prevent routing loops.
C. The cluster ID is not required to use route reflectors.
D. Route reflectors require that the network is connected in a physical full mesh.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 36
What are two reasons for creating multiple areas in OSPF? (Choose two.)
A. to reduce the convergence time
B. to increase the number of adjacencies in the backbone
C. to increase the size of the LSDB
D. to reduce LSA flooding across the network
Answer: A,D (LEAVE A REPLY)
Option A is correct. Creating multiple areas in OSPF can help to reduce the convergence time . This is because changes in one area do not affect other areas, so fewer routers need to run the SPF algorithm in response to a change.
Option D is correct. Creating multiple areas in OSPF can help to reduce Link State Advertisement (LSA) flooding across the network. This is because LSAs are not flooded out of their area of origin.

NEW QUESTION: 37
Referring to the exhibit, which three statements are correct? (Choose three.)

A. The IP source guard is enabled.
B. Dynamic ARP inspection is enabled.
C. DHCP snooping is enabled.
D. The DHCP snooping database is protected.
E. DHCP snooping is enabled for IPv6 traffic.
Answer: (SHOW ANSWER)
The IP source guard is enabled.
The configuration shows that ip-source-guardis explicitly enabled under dhcp-security.
Dynamic ARP inspection is enabled.
The configuration also includes arp-inspectionunder dhcp-security, which indicates that Dynamic ARP Inspection (DAI) is enabled.
DHCP snooping is enabled.
DHCP snooping is implied as enabled because both IP source guard and ARP inspection are dependent on the DHCP snooping database to function properly. Even though DHCP snooping is not explicitly mentioned, its presence is required for these features to work.

NEW QUESTION: 38
You are troubleshooting a BGP routing issue between your network and a customer router and are reviewing the BGP routing policies. Which two statements are correct in this scenario? (Choose two.)
A. Export policies are applied to routes in the RIB-In table.
B. Import policies are applied to routes in the RIB-Local table.
C. Import policies are applied after the RIB-In table.
D. Export policies are applied after the RIB-Local table.
Answer: C,D (LEAVE A REPLY)
In BGP, routing policies are used to control the flow of routing information between BGP peers.

Option C suggests that import policies are applied after the RIB-In table. This is correct because import policies in BGP are applied to routes that are received from a BGP peer, before they are installed in the local BGP Routing Information Base (RIB-In). The RIB-In is a database that stores all the routes that are received from all peers.
Option D suggests that export policies are applied after the RIB-Local table. This is correct because export policies in BGP are applied to routes that are being advertised to a BGP peer, after they have been selected from the local BGP Routing Information Base (RIB-Local). The RIB-Local is a database that stores all the routes that the local router is using.

NEW QUESTION: 39
Which statement is correct about controlling the routes installed by a RIB group?
A. An import policy is applied to the RIB group.
B. Only routes in the last table are installed.
C. A firewall filter must be configured to install routes in the RIB groups.
D. An export policy is applied to the RIB group.
Answer: (SHOW ANSWER)
Explanation
A RIB group is a configuration that allows a routing protocol to install routes into multiple routing tables in Junos OS. A RIB group consists of an import-rib statement,which specifies the source routing table, and an export-rib statement, which specifies the destination routing table or group. A RIB group can also include an import-policy statement, which specifies one or more policies to control which routes are imported into the destination routing table or group1.
An import policy is a policy statement that defines the criteria for accepting or rejecting routes from the source routing table. An import policy can also modify the attributes of the imported routes, such as preference, metric, or community. An import policy can be applied to a RIB group by using the import-policy statement under the [edit routing-options rib-groups] hierarchy level1.
Therefore, option A is correct, because an import policy is applied to the RIB group to control which routes are installed in the destination routing table or group. Option B is incorrect, because all routes in the source routing table are imported into the destination routing table or group, unless filtered by an import policy.
Option C is incorrect, because a firewall filter is not used to install routes in the RIB groups; a firewall filter is used to filter packets based on various criteria. Option D is incorrect, because an export policy is not applied to the RIB group; an export policy is applied to a routing protocol to control which routes are advertised to other devices.
References:
1: rib-groups | Junos OS | Juniper Networks

NEW QUESTION: 40
Which two statements about the firewall filter terms shown in the exhibit are true? (Choose two.)

```
edit firewall
ser@switch# show
amily ethernet-switching {
    filter mac-address {
        term one {
            from {
                source-mac-address 88:05:00:29:3c:de/48
            }
            then {
                log;
            }
            term two {
                then accept;
            }
        }
    }
```


A. Traffic that matches the fromstatement in term oneis discarded.
B. Traffic that matches the fromstatement in term oneis accepted.
C. All traffic not matching the fromstatement in term oneis accepted by term two.
D. Traffic matching the fromstatement in term twois logged.
Answer: B,C (LEAVE A REPLY)

NEW QUESTION: 41
Which two types of tunnels are able to be created on all Junos devices? (Choose two.)
A. STP
B. GRE
C. IP-IP
D. IPsec
Answer: (SHOW ANSWER)
GRE tunnels can be created on all Junos devices. GRE is a tunneling protocol that encapsulates a wide variety of network layer protocols inside virtual point-to-point links.
IP-IP tunnels can also be created on all Junos devices. IP-IP tunneling is used to encapsulate IP packets within IP packets, providing a simple way to tunnel IP traffic.

NEW QUESTION: 42
You are troubleshooting an issue where traffic to 192.168.10.0/24 is being sent to R1 instead of your desired path through R2.
Referring to the exhibit, what is the reason for the problem?

```
R1 - 10.100.24.2
R2 - 10.100.25.2

user@router# run show route protocol bgp 192.168.10.0/24
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)
+ = ActiveRoute, - = Last Active, * = Both
192.168.10.0/24     * [BGP/170] 00:00:30, localpref 500
                        AS path: 64533 I, validation-state: unverified
                      > to 10.100.24.2 via ge-0/0/0.0
                      [BGP/170] 00:00:00, localpref 100
                        AS path: 64533 64533 64533 64533 64544 ?, validation-
state: unverified
                      > to 10.100.25.2 via ge-0/0/1.0
```


A. R2's route is not the best path due to loop prevention.
B. R2's route is not the best path due to a lower origin code.
C. R1's route is the best path due to a higher local preference
D. R1's route is the best path due to the shorter AS path.
Answer: C (LEAVE A REPLY)
The exhibit shows the output of the command show ip bgp, which displays information about the BGP routes in the routing table. The output shows two routes for the destination 192.168.10.0/24, one from R1 and one from R2.
The route from R1 has a local preference of 200, while the route from R2 has a local preference of 100. Local preference is a BGP attribute that indicates the degree of preference for a route within an autonomous system (AS). A higher local preference means a more preferred route.
BGP uses a best path selection algorithm to choose the best route for each destination among multiple paths. The algorithm compares different attributes of the routes in a specific order of precedence. The first attribute that is compared is weight, which is a Cisco-specific attribute that is local to the router. If the weight is equal or not set, the next attribute that is compared is local preference.
In this case, both routes have the same weight of 0, which means that they are learned from external BGP (eBGP) peers. Therefore, the next attribute that is compared is local preference. Since R1's route has a higher local preference than R2's route, it is chosen as the best path and installed in the routing table. The other attributes, such as origin code and AS path, are not considered in this case.

NEW QUESTION: 43
Which statement is correct about the storm control feature?
A. The storm control feature is enabled in the factory-default configuration on EX Series switches.
B. The storm control feature requires a special license on EX Series switches.
C. The storm control feature is not supported on aggregate Ethernet interfaces.
D. The storm control configuration only applies to traffic being sent between the forwarding and control plane.
Answer: A (LEAVE A REPLY)

The storm control feature is enabled in the factory-default configuration on EX Series switches. On EX2200, EX3200, EX3300, EX4200, and EX6200 switches, the factory default configuration enables storm control for broadcast and unknown unicast traffic on all switch interfaces. On EX4300 switches, the factory default configuration enables storm control on all Layer 2 switch interfaces.

NEW QUESTION: 44
Referring to the exhibit, Router-1 and Router-2 are failing to form an IS-IS adjacency.

```
[edit]
user@Router-1# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.10.10.33/24;
        }
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.0.254/24;
        }
        family iso {
            address 49.0003.0192.0168.0113.00;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.1.11/32;
        }
        family iso {
            address 49.0002.0192.0168.0111.00;
        }
    }
}

[edit]
user@Router-1# show protocols
isis {
    overload;
    level 2 disable;
    interface all;
}

[edit]
```

```
user@Router-2# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.10.10.34/24;
        }
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.0.1/16;
        }
        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.1.12/32;
        }
        family iso {
            address 49.0001.0192.0168.0112.00;
        }
    }
}

[edit]
user@Router-2# show protocols
isis {
    interface all;
}
```

What should you do to solve the problem?

A. Remove the ISO address from ge-0/0/2 on Router-1.

B. Change the IP subnet masks to match on the ge-0/0/2 interfaces of both routers.

C. Remove the overloaded statement from Router-1.

D. Change the ISO areas on the lo0 interfaces to match on both routers.

Answer: (SHOW ANSWER)

NEW QUESTION: 45
In a LAN segment, which tie-breaker would spanning tree consider if equal-cost paths exist between two or more switches to the root bridge?
A. port priority
B. bridge ID
C. interface number
D. port ID
Answer: (SHOW ANSWER)

NEW QUESTION: 46
You deployed a new EX Series switch with DHCP snooping enabled and you do not see any entries in the snooping databases for an interface. Which two Juniper configurations for that interface caused this issue? (Choose two.)
A. The interface is configured as a disabled port.
B. MAC limiting is enabled on the interface.
C. The interface is configured as a trunk port.
D. Dynamic ARP inspection is enabled on the interface.
Answer: (SHOW ANSWER)
A is correct because the interface is configured as a disabled port. A disabled port does not forward any traffic, including DHCP packets. Therefore, DHCP snooping cannot learn any MAC addresses or lease information from a disabled port.
C is correct because the interface is configured as a trunk port. By default, all trunk ports on the switch are trusted for DHCP snooping. This means that DHCP snooping does not inspect or filter any DHCP packets received on a trunk port. Therefore, DHCP snooping does not add any entries to the snooping database for a trunk port.

NEW QUESTION: 47
Which statement is true about Layer 2 firewall filters on EX Series switches?
A. They are stateless and evaluated by the forwarding plane.
B. They are stateful and evaluated by the control plane.
C. They are stateful and evaluated by the forwarding plane.
D. They are stateless and evaluated by the control plane.

Answer: A (LEAVE A REPLY)

NEW QUESTION: 48
Which two statements are true about the default VLAN on Juniper switches? (Choose two.)
A. The default VLAN is set to a VLAN ID of 1 by default
B. The default VLAN ID is not assigned to any interface.
C. The default VLAN ID is not visible.
D. The default VLAN ID can be changed.
Answer: A,D (LEAVE A REPLY)
Explanation
On Juniper switches, the default VLAN is set to a VLAN ID of 1 by default12. This means that all interfaces on the switch are members of VLAN 1 until they are specifically assigned to another VLAN12. Therefore, option A is correct.
The default VLAN ID can be changed12. This allows network administrators to configure the switch to use a different VLAN as the default, if necessary12. Therefore, option D is correct.

NEW QUESTION: 49
Which two statements about redundant trunk groups on EX Series switches are correct? (Choose two.)
A. Redundant trunk groups load-balance traffic across two designated uplink interfaces.
B. If the active link fails, then the secondary link automatically takes over.
C. Layer 2 control traffic is permitted on the secondary link
D. Redundant trunk groups must be connected to the same aggregation switch.
Answer: B,D (LEAVE A REPLY)
Explanation
Redundant Trunk Groups (RTGs) on EX Series switches provide a simple solution for network recovery when a trunk port on a switch goes down1. They are configured on the access switch and contain two links: a primary or active link, and a secondary link1. Therefore, option B is correct because if the active link fails, the secondary link automatically starts forwarding data traffic without waiting for normal spanning-tree protocol convergence1.
Option D is also correct. In a typical enterprise network composed of distribution and access layers, RTGs are used where one Access switch is connected to two different uplink switches2. This implies that RTGs must be connected to the same aggregation switch2.

NEW QUESTION: 50
Referring to the output shown in the exhibit, which statement is correct?

```
user@host> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.26.1.1       ge-0/0/3.0             2Way      192.168.1.1      128  31
```

A. The state is normal for a DR neighbor.

B. The state is normal for a DRother neighbor
C. An MTU mismatch exists between the OSPF neighbors.
D. An area ID mismatch exists between the OSPF neighbors
Answer: (SHOW ANSWER)
In OSPF, the state of the neighbor relationship is determined by the exchange of OSPF packets between routers. The state "2Way" as shown in the exhibit indicates that bi-directional communication has been established between the two OSPF routers. This is the normal state for a neighbor that is not the Designated Router (DR) or Backup Designated Router (BDR) on a broadcast, non-broadcast multi-access (NBMA), or point-to-multipoint network. These neighbors are often referred to as "DRothers". Therefore, option B is correct.

NEW QUESTION: 51
Which two statements about BGP facilitate the prevention of routing loops between two autonomous systems? (Choose two.)
A. EBGP routers will append their AS number when advertising routes to their neighbors.
B. EBGP routers will only accept routes that contain their own AS number in the AS_PATH.
C. EBGP routers will drop routes that contain their own AS number in the AS_PATH
D. EBGP routers will prepend their AS number when advertising routes to their neighbors
Answer: A,C (LEAVE A REPLY)
When an EBGP router advertises a route to a neighbor in a different AS, it appends its own AS number to the AS_PATH attribute. This helps other routers recognize the path the route has taken and prevents routing loops.
If an EBGP router receives a route advertisement that already contains its own AS number in the AS_PATH attribute, it will drop the route. This prevents the router from accepting a route that would create a routing loop.

NEW QUESTION: 52
What is the default keepalive time for BGP?
A. 10 seconds
B. 60 seconds
C. 30 seconds
D. 90 seconds
Answer: B (LEAVE A REPLY)
Explanation
The default keepalive time for BGP is 60 seconds1. The keepalive time is the interval at which BGP sends keepalive messages to maintain the connection with its peer1. If the keepalive message is not received within the hold time, the connection is considered lost1. By default, the hold time is three times the keepalive time, which is 180 seconds1.

NEW QUESTION: 53
Which statement is true about IP-IP tunnels?

A. The packet header is replaced before entering the tunnel.
B. The time-to-live value of the original packet is decremented.
C. The packet is encapsulated unchanged before entering the tunnel.
D. IP-IP tunnels are protocol agnostic.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 54
Exhibit

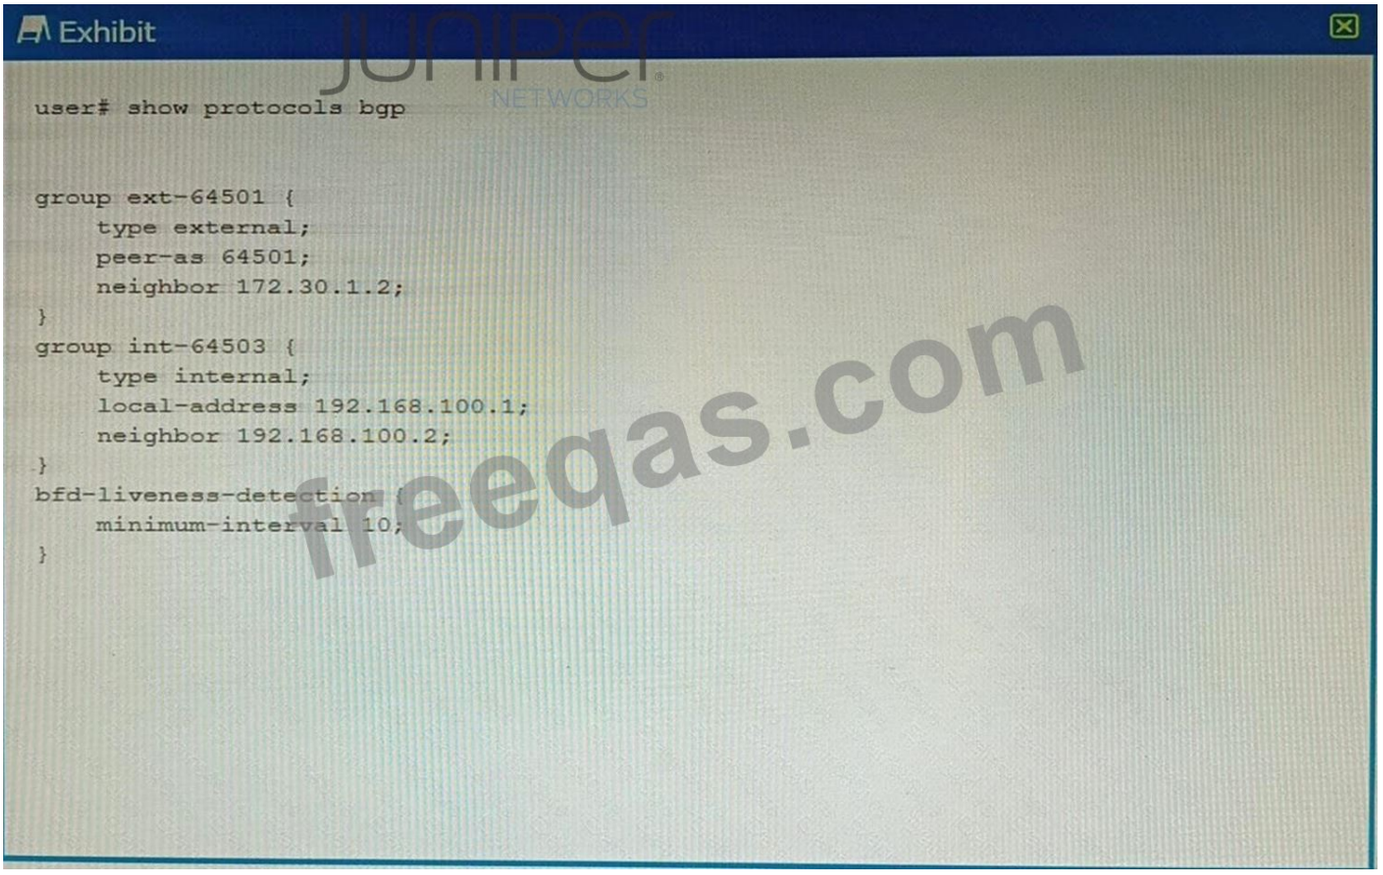

Your BGP neighbors, one in the USA and one in France, are not establishing a connection with each other.
Referring to the exhibit, which statement is correct?
A. The BFD liveness is set too low.
B. The BFD liveness must be configured on the BGP neighbor.
C. The BFD liveness must be configured on the BGP group.
D. The BFD liveness is set too high.
Answer: B (LEAVE A REPLY)
Explanation
The exhibit shows the configuration of BFD liveness detection for BGP at the global level, which applies to all BGP neighbors by default1. However, this configuration does not specify the session mode, which determines whether BFD uses single-hop or multihop mode to communicate with a neighbor2.
For single-hop BGP neighbors, which are directly connected on the same subnet, the session mode can be either automatic or single-hop. For multihop BGPneighbors, which are not directly

connected and require multiple hops to reach, the session mode must be multihop2.
Since your BGP neighbors are in different countries, they are likely to be multihop neighbors.
Therefore, you need to configure the session mode as multihop for each neighbor individually at the [edit protocols bgp group group-name neighbor address bfd-liveness-detection] hierarchy level2. For example:
protocols { bgp { group usa { neighbor 192.0.2.1 { bfd-liveness-detection { session-mode multihop; } } } group france { neighbor 198.51.100.1 { bfd-liveness-detection { session-mode multihop; } } } } } If you do not configure the session mode for multihop neighbors, BFD will use the default mode of automatic, which will try to use single-hop mode and fail to establish a BFD session with the remote neighbor2. This will prevent BGP from using BFD to detect liveliness and failover.
Therefore, the answer B is correct, as you need to configure the BFD liveness detection on the BGP neighbor level with the appropriate session mode for multihop neighbors.

NEW QUESTION: 55
Which command displays the output shown in the exhibit?

```
Routing table: default.ethernet-switching
ETHERNET-SWITCHING:

Destination             Type    RtRef  Next hop  Type   Index   NhRef  Netif

default                 perm        0            dscd      66       1
2, *                    user        0            comp    1304       2
2, *                    intf        0            rslv    1302       1
2, 00:26:88:02:74:86    user        0            ucst    1303       3  ge-0/0/6.0
2, 00:26:88:02:74:87    user        0            ucst    1305       3  ge-0/0/7.0
2, 00:26:88:02:74:88    user        0            ucst    1306       3  ge-0/0/8.0
```


A. show route forwarding-table
B. show ethernet-switching table
C. show ethernet-switching table extensive
D. show route forwarding-table family ethernet-switching
Answer: (SHOW ANSWER)
This command shows the Ethernet switching table, which includes information about MAC addresses, their associated VLANs, and the interfaces they are learned on, similar to the output shown in your exhibit.

NEW QUESTION: 56
You need to configure a LAG between your switches. In this scenario, which two statements are correct? (Choose two.)
A. Duplex and speed settings are not required to match on both participating devices.
B. Duplex and speed settings are required to match on both participating devices.
C. Member links are not required to be contiguous ports.
D. Member links are required to be contiguous ports.
Answer: (SHOW ANSWER)
Up to 64 Ethernet interfaces can be grouped to form a LAG, and In a Junos Fusion, up to 1,000

LAGs are supported on QFX10002 switches acting as aggregation devices.
The LAG must be configured on both sides of the link.
The interfaces on either side of the link must be set to the same speed and be in full-duplex mode.

NEW QUESTION: 57
Two routers share the same highest priority and start time. In this situation, what is evaluated next when determining the designated router?
A. The router with the lowest router ID become the DR.
B. The router with the highest router ID becomes the DR
C. The routers perform another DR election.
D. The router with the highest MAC address become the DR
Answer: B (LEAVE A REPLY)
According to the OSPF protocol, the designated router (DR) is the router that acts as the focal point for exchanging routing information on a multi-access network segment, such as a LAN. The DR election process is based on the following criteria, in order of precedence:
The router with the highest OSPF priority becomes the DR. The default priority is 1, and a priority of 0 means the router will not participate in the election.
If there is a tie in priority, the router with the highest router ID becomes the DR. The router ID is a 32-bit number that uniquely identifies a router in an OSPF domain. It can be manually configured or automatically derived from the highest IP address of a loopback interface or a physical interface. If there is a tie in router ID, the router that was first to become an OSPF neighbor becomes the DR. In your scenario, two routers share the same highest priority and start time. This means that they have equal chances of becoming the DR based on the first and third criteria. Therefore, the second criterion will be used to break the tie, which is the router ID. The router with the highest router ID will become the DR, and the other router will become the backup designated router (BDR), which is ready to take over the role of DR if it fails.
NEW QUESTION: 58
You are asked to connect an IP phone and a user computer using the same interface on an EX Series switch.
The traffic from the computer does not use a VLAN tag, whereas the traffic from the IP phone uses a VLAN tag.
Which feature enables the interface to receive both types of traffic?
A. voice VLAN
B. native VLAN
C. MAC limiting
D. DHCP snooping
Answer: A (LEAVE A REPLY)
Explanation
The feature that enables an interface on an EX Series switch to receive both untagged traffic (from the computer) and tagged traffic (from the IP phone) is the voice VLAN12.

The voice VLAN feature in EX-series switches enables access ports to accept both data (untagged) and voice (tagged) traffic and separate that traffic into different VLANs12. This allows the switch to differentiate between voice and data traffic, ensuring that voice traffic can be treated with a higher priority12. Therefore, option D is correct.

NEW QUESTION: 59
You are asked to connect an IP phone and a user computer using the same interface on an EX Series switch. The traffic from the computer does not use a VLAN tag, whereas the traffic from the IP phone uses a VLAN tag.
Which feature enables the interface to receive both types of traffic?
A. native VLAN
B. DHCP snooping
C. MAC limiting
D. voice VLAN
Answer: (SHOW ANSWER)
The feature that enables an interface on an EX Series switch to receive both untagged traffic (from the computer) and tagged traffic (from the IP phone) is the voice VLAN. The voice VLAN feature in EX-series switches enables access ports to accept both data (untagged) and voice (tagged) traffic and separate that traffic into different VLANs. This allows the switch to differentiate between voice and data traffic, ensuring that voice traffic can be treated with a higher priority.

NEW QUESTION: 60
Which statement is correct about GRE tunnels?
A. GRE tunnels only support encapsulating non-IP traffic.
B. There are 20 bytes of overhead with GRE encapsulation.
C. The time-to-live field in the tunnel payload is decremented during transit.
D. GRE tunnels only support encapsulating IP traffic.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 61
You are receiving the BGP route shown in the exhibit from four different upstream ISPs.
Referring to the exhibit, which ISP will be selected as the active path?

Route	Next-hop	AS-Path	Origin	Local preference
172.27.0.0/24	ISP 1	65010 65520 65512	I	100
172.27.0.0/24	ISP 2	65112	E	100
172.27.0.0/24	ISP 3	64599 65532 65520 65512	?	150
172.27.0.0/24	ISP 4	65000 65512	E	150

A. ISP 1
B. ISP 3
C. ISP 4

D. ISP 2

Answer: (SHOW ANSWER)

In BGP, the path selection process is based on a set of attributes. The process starts by preferring the path with the highest weight, then the highest local preference, then the locally originated routes, and so on. If all these attributes are the same, then it prefers the path with the shortest AS path.

Referring to the exhibit, all four ISPs have the same weight, local preference, and origin. However, ISP 4 has the shortest AS path. Therefore, ISP 4 will be selected as the active path. So, option C is correct.

NEW QUESTION: 62

What does the * indicate in the output shown in the exhibit?

A. The switch ports have a router attached.

B. The interface is down.

C. The interface is active.

D. All interfaces have elected a root bridge.

Answer: C (LEAVE A REPLY)

The exhibit shows the output of the command show vlans brief, which displays brief information about VLANs and their associated interfaces.

The output has four columns: Routing instance, VLAN name, Interfaces, and Tagging. The * symbol indicates that the interface is active, meaning that it is up and forwarding traffic. This can

be verified by the command show interfaces terse, which displays the status of the interfaces.

NEW QUESTION: 63
Which two statements are correct about link aggregation? (Choose two.)
A. IP traffic is hashed using source and destination MAC addresses.
B. Member links must use the same MTU.
C. All RE-generated traffic traverses the lowest member link.
D. LAGs provide physical layer redundancy.
Answer: B,D (LEAVE A REPLY)

NEW QUESTION: 64
A number of reports from end users indicate that internal and external communications are intermittent and not reliable. You verified the status of the switch ports and have determined that they are up and operational. You also noticed a very high level of link bandwidth utilization on those same ports. The current topology of the affected environment is shown in the exhibit.

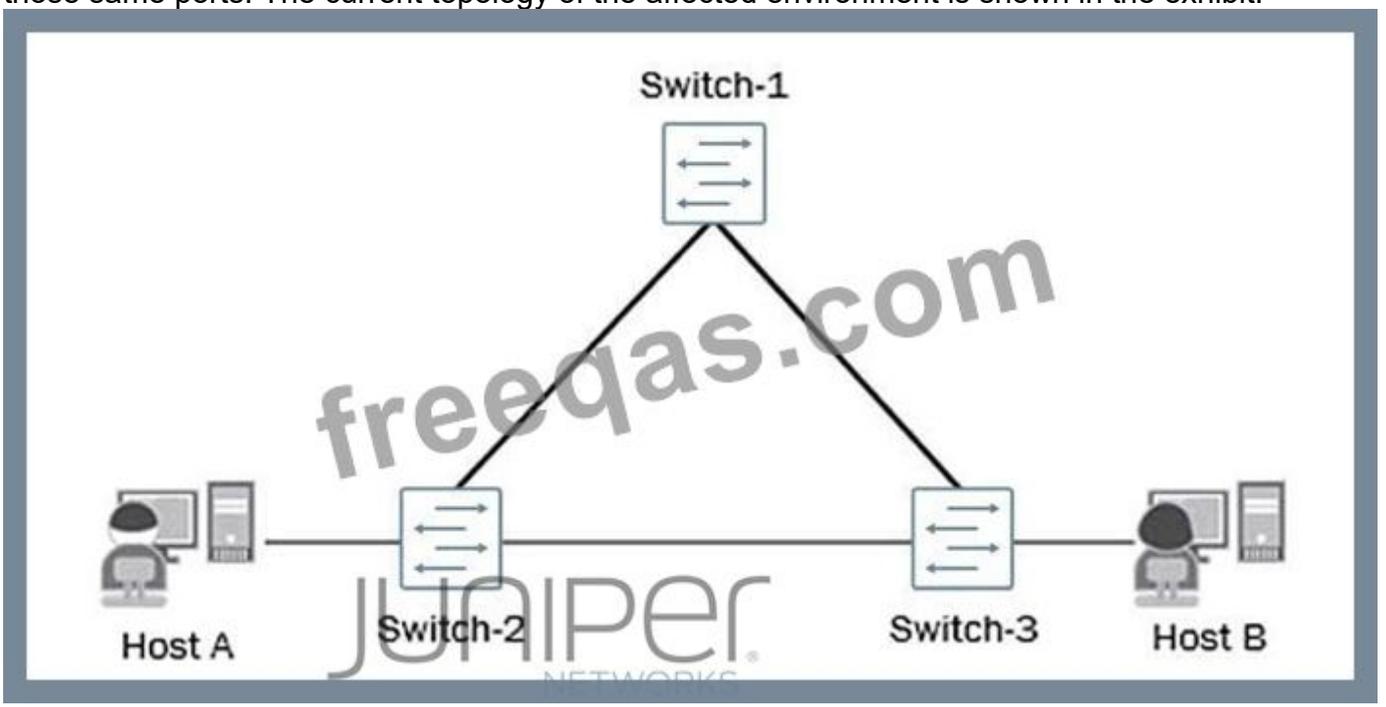

What would be the cause of the reported issues?
A. A misconfigured interior gateway protocol (IGP).
B. A malformed route-based ACL improperly filtering traffic flows.
C. A lack of port-based ACLs filtering the traffic flows.
D. A lack of a loop-prevention mechanism or protocol.
Answer: (SHOW ANSWER)

NEW QUESTION: 65
Which three statements are correct about OSPF packet types? (Choose three.)
A. Database description packets are used during the adjacency formation process.
B. Link-state acknowledgment packets are sent every 30 seconds.

C. Link-state update packets contain one or more LSAs.
D. Hello packets are sent every 30 seconds by default.
E. Link-state request packets are sent after receiving an LSA header not in the OSPF database.
Answer: (SHOW ANSWER)

NEW QUESTION: 66
Your network has two ISPs available. You want to ensure that all outbound traffic is routed out ISP-1. If the connection to ISP-1 fails, all outbound traffic is routed to the backup ISP, ISP-2. In this scenario, how should you configure BGP on your network?
A. Configure the gateway for ISP-1 with a higher peer ID than the gateway for ISP-2.
B. Set the local-preference attribute to a higher value for ISP-2 than ISP-1.
C. Configure the gateway for ISP-1 with a higher origin code than the gateway for ISP-2.
D. Set the local-preference attribute to a higher value for ISP-1 than ISP-2.
Answer: D (LEAVE A REPLY)
NEW QUESTION: 67
You need to configure a LAG between your switches. In this scenario, which two statements are correct?
(Choose two.)
A. Duplex and speed settings are not required to match on both participating devices.
B. Duplex and speed settings are required to match on both participating devices.
C. Member links are not required to be contiguous ports.
D. Member links are required to be contiguous ports.
Answer: (SHOW ANSWER)
B is correct because duplex and speed settings are required to match on both participating devices. According to the Juniper Networks documentation1, all the interfaces in a LAG must have the same speed and be in full-duplex mode. This ensures that the LAG can operate as a single logical link without any performance or compatibility issues.
C is correct because member links are not required to be contiguous ports. According to the Juniper Networks documentation2, you can group any Ethernet interfaces on a switch into a LAG, regardless of their physical location or slot number. This provides flexibility and scalability for configuring LAGs on switches.

NEW QUESTION: 68
Referring to the exhibit, which configuration change is needed for an IS-IS Level 1 adjacency between R1 and R2?

A. Configure the lo0 family ISO address 49.0002.0010.0042.0002.00 on R1

B. Disable Level 2 on R2's ge-0/0/1 interface

C. Configure the lo0 family ISO address 49.0002.0010.0042.0002.00 on R2

D. Enable Level 2 on R1's ge-0/0/1 interface

Answer: C (LEAVE A REPLY)

NEW QUESTION: 69

Referring to the exhibit, which two statements are correct? (Choose two.)

```
[edit protocols rstp]
user@switch# show
bridge-priority 8k;
interface ge-0/0/10 {
    disable;
}
interface ge-0/0/13 {
    cost 20000;
    mode point-to-point;
}
```


A. This device must be selected as the root bridge.
B. The ge-0/0/10 interface will be part of the RSTP topology but will block incoming BPDUs.
C. The ge-0/0/13 interface will be selected as the forwarding interface.
D. The ge-0/0/10 interface will not participate in the RSTP topology.
Answer: (SHOW ANSWER)

NEW QUESTION: 70
You have configured OSPF routing as shown in the exhibit. You notice that all interfaces have formed full adjacencies, with the exception of the interfaces connecting R3 and R4 with a status of 2Way.

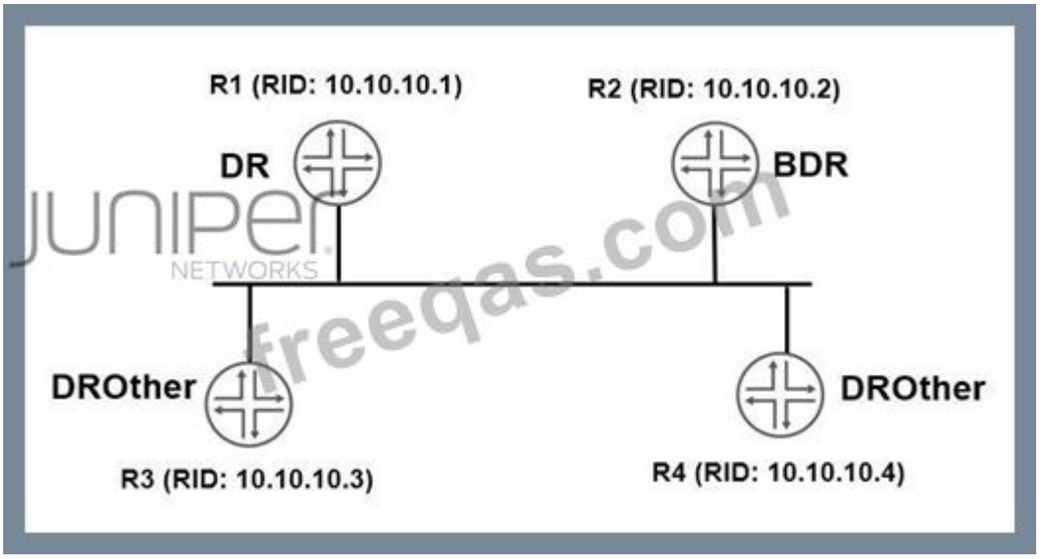

What is the reason for this status?
A. DROther routers will not form a full adjacency with each other.
B. The two routers must both be configured as DR routers.
C. The two routers must be configured in different areas.
D. The interface-type is not configured as p2p.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 71
You have two OSPF routers forming an adjacency. R1 has a priority of 32 and a router ID of 192.168.1.2. R2 has a priority of 64 and a router ID of 192.168.1.1. The routers were started at

the same time and all other OSPF settings are the default settings.
Which statement is correct in this scenario?
A. At least three routers are required for a DR/BDR election
B. Router IDs must match for an adjacency to form.
C. R2 will be the BDR.
D. R1 will be the BDR.
Answer: (SHOW ANSWER)
Explanation
In OSPF, the Designated Router (DR) and Backup Designated Router (BDR) are elected based on the priority of the routers1. The router with the highest priority becomes the DR, and the router with the second highest priority becomes the BDR1. If there is a tie in priority, then the router with the highest Router ID is chosen1.
In this scenario, R2 has a higher priority (64) than R1 (32), so R2 will become the DR1. Since R1 has the second highest priority, it will become the BDR1. Therefore, option D is correct.

NEW QUESTION: 72
You have two OSPF routers forming an adjacency. R1 has a priority of 32 and a router ID of 192.168.1.2. R2 has a priority of 64 and a router ID of 192.168.1.1. The routers were started at the same time and all other OSPF settings are the default settings.
Which statement is correct in this scenario?
A. At least three routers are required for a DR/BDR election
B. Router IDs must match for an adjacency to form.
C. R2 will be the BDR.
D. R1 will be the BDR.
Answer: (SHOW ANSWER)
The router with the highest OSPF priority becomes the DR. In this case, R2 has a higher priority (64) compared to R1's priority (32). Hence, R2 will become the DR.
The router with the next highest priority becomes the BDR. Since R1 has the next highest priority after R2, R1 will become the BDR.

NEW QUESTION: 73
Which two statements are correct about generated routes? (Choose two.)
A. Generated routes require a contributing route.
B. Generated routes show a next hop in the routing table.
C. Generated routes appear in the routing table as static routes
D. Generated routes cannot be redistributed into dynamic routing protocols.
Answer: A,B (LEAVE A REPLY)
A is correct because generated routes require a contributing route. A contributing route is a route that matches the destination prefix of the generated route and has a valid next hop1. A generated route is only installed in the routing table if there is at least one contributing route available2. This ensures that the generated route is reachable and useful. If there is no contributing route, the

generated route is not added to the routing table2.
B is correct because generated routes show a next hop in the routing table. A generated route inherits the next hop of its primary contributing route, which is the most preferred route among all the contributing routes2. The next hop of the generated route can be either an IP address or an interface name, depending on the type of the contributing route2. The next hop of the generated route can also be modified by a routing policy3.

NEW QUESTION: 74
Your BGP neighbors, one in the USA and one in France, are not establishing a connection with each other.
Referring to the exhibit, which statement is correct?

```
user# show protocols bgp

group ext-64501 {
    type external;
    peer-as 64501;
    neighbor 172.30.1.2;
}
group int-64503 {
    type internal;
    local-address 192.168.100.1;
    neighbor 192.168.100.2;
}
bfd-liveness-detection {
    minimum-interval 10;
{
```


A. The BFD liveness is set too low.
B. The BFD liveness must be configured on the BGP neighbor.
C. The BFD liveness must be configured on the BGP group.
D. The BFD liveness is set too high.

Answer: (SHOW ANSWER)
In the configuration shown, BFD (Bidirectional Forwarding Detection) liveness detection is globally configured, but for BFD to work with specific BGP neighbors, it must be configured under each BGP neighbor or BGP group. Without specific BFD configuration at the neighbor or group level, BFD will not be used to monitor the BGP sessions, which might be causing the BGP neighbors not to establish a connection.

NEW QUESTION: 75
Which two multicast IP addresses are used for link-state update packets? (Choose two.)
A. 224.0.0.9
B. 224.0.0.10
C. 224.0.0.6
D. 224.0.0.5
Answer: C,D (LEAVE A REPLY)