Juniper.JN0-351.v2026-05-27.q43

Exam Code:	JN0-351
Exam Name:	Enterprise Routing and Switching, Specialist (JNCIS-ENT)
Certification Provider:	Juniper
Free Question Number:	43
Version:	v2026-05-27
# of views:	105
# of Questions views:	443
https://www.freecram.net/torrent/Juniper.JN0-351.v2026-05-27.q43.html	

NEW QUESTION: 1
Refer to the exhibit.

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

Referring to the output shown in the exhibit, which statement is correct?
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 1 . The state "2Way" as shown in the exhibit indicates that bi-directional communication has been established between the two OSPF routers 1 . 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 1 . These neighbors are often referred to as " DRothers " 1 .
Therefore, option B is correct.

NEW QUESTION: 2
An update to your organization ' s network security requirements document requires management traffic to be isolated in a non-default routing-instance. You want to implement this requirement on your Junos-based devices.
Which two commands enable this behavior? (Choose two.)
A. set routing-instances mgmtjunoa interface ge-0/0/0.0
B. set routing-instances mgmt_junos interface em1
C. set system management-instance
D. set routing-instances mgmt_junos

Answer: (SHOW ANSWER)
To isolate management traffic in a non-default routing-instance on Junos-based devices, you can use the set system management-instance and set routing-instances mgmt_junos commands 1 2 .
* set system management-instance : This command associates the management interface (usually named fxp0 or em0 for Junos OS, or re0:mgmt-* or re1:mgmt-* for Junos OS Evolved) with the non-default virtual routing and forwarding (VRF) instance 1 . After you configure the non-default management VRF instance, management traffic no longer has to share a routing table with other control traffic or protocol traffic 1 .
* set routing-instances mgmt_junos : This command creates a new routing instance named mgmt_junos
. The name of the dedicated management VRF instance is reserved and hardcoded as mgmt_junos ; you cannot configure any other routing instance by the name mgmt_junos 1 .
Therefore, options C and D are correct. Options A and B are not correct because they attempt to assign an interface to the mgmt_junos routing instance, which is not necessary for isolating management traffic 1 .

NEW QUESTION: 3
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)
Junos devices support various types of tunnels for different purposes 1 2 .
* Option B is correct. Generic Routing Encapsulation (GRE) is a tunneling protocol that can encapsulate a wide variety of network layer protocols inside virtual point-to-point links over an Internet Protocol network 1 . Junos devices support GRE tunnels 1 .
* Option D is correct. IPsec (Internet Protocol Security) is a protocol suite for securing Internet Protocol (IP) communications by authenticating and encrypting each IP packet of a communication session 1
. Junos devices support IPsec tunnels 1 .
* Option A is incorrect. Spanning Tree Protocol (STP) is not a type of tunnel. It's a network protocol designed to prevent loops in a bridged Ethernet local area network 2 .
* Option C is incorrect. While Junos devices do support IP-IP (also known as IP tunneling), it's not supported on all Junos devices 1 .

NEW QUESTION: 4
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: (SHOW ANSWER)
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: 5

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. static

B. multicast

C. aggregate

D. anycast

Answer: (SHOW ANSWER)

NEW QUESTION: 6

Exhibit.

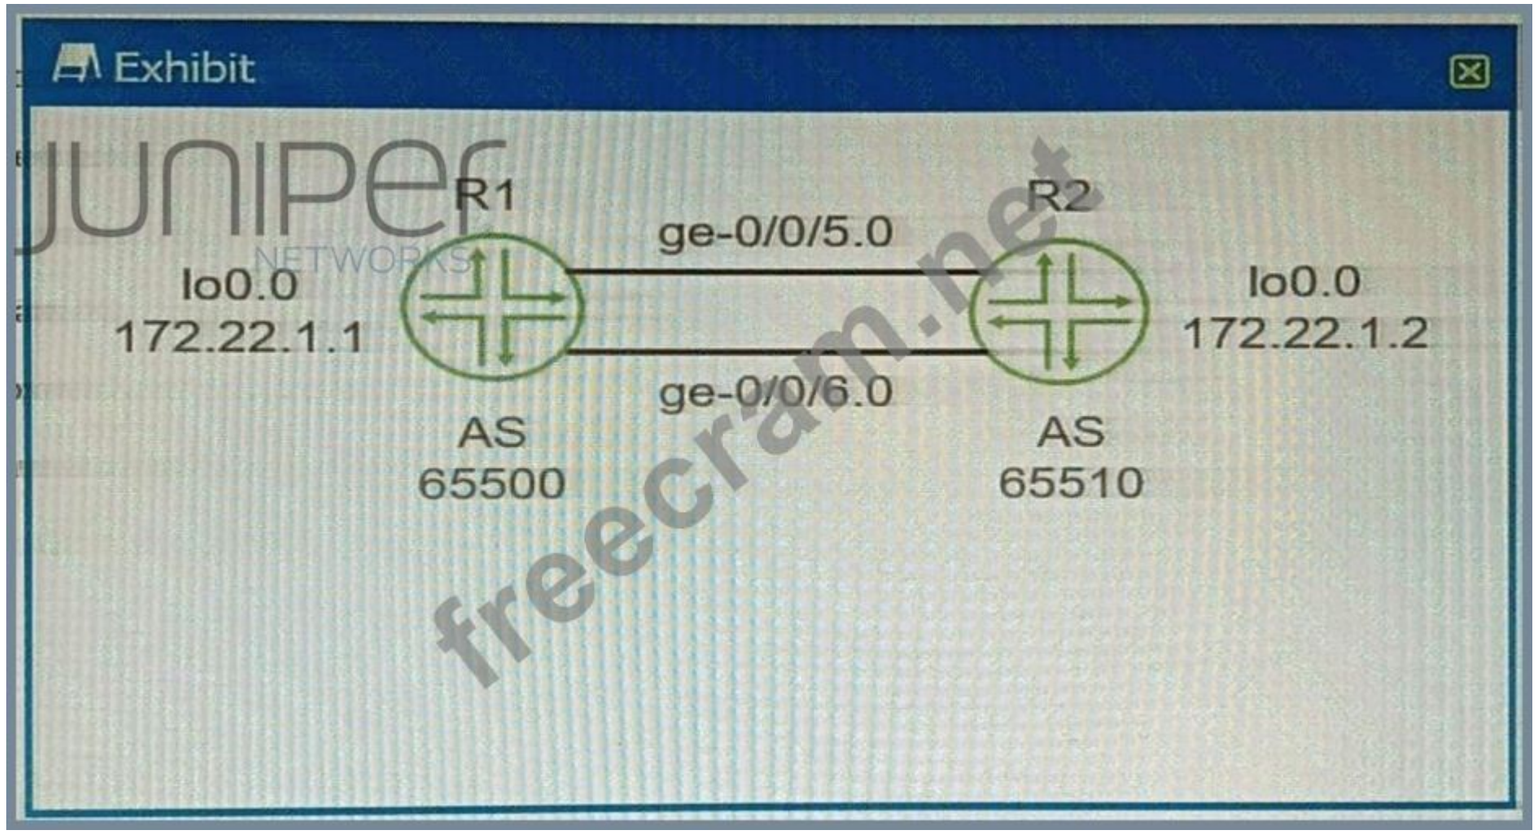

You want to enable redundancy for the EBGP peering between the two routers shown in the exhibit. Which three actions will you perform in this scenario? (Choose three.)

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: 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 1 . 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 1

. However, if you configure BGP multihop with a higher TTL value, you can allow EBGP peers to communicate over multiple routers in between 1 . 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 2
. Loopback interfaces are virtual interfaces that are always up and reachable as long as the router is operational 2 . 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 2
. 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 2 . You can use static routes or dynamic routing protocols to advertise and learn the routes for the peer loopback interface IP addresses 2 . Without these routes, the routers will not be able to establish or maintain the BGP session using their loopback interfaces.

NEW QUESTION: 7
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: 8
You are asked to create a new firewall filter to evaluate Layer 3 traffic that is being sent between VLANs. In this scenario, which two statements are correct? (Choose two.)
A. You should create a family Ethernet-switching firewall filter with the appropriate match criteria and actions.
B. You should apply the firewall filter to the appropriate VLAN.
C. You should create a family inet firewall filter with the appropriate match criteria and actions.
D. You should apply the firewall filter to the appropriate IRB interface.
Answer: (SHOW ANSWER)
A firewall filter is a configuration that defines the rules that determine whether to forward or discard packets at specific processing points in the packet flow. A firewall filter can also modify the attributes of the packets, such as priority, marking, or logging. A firewall filter can be applied to various interfaces, protocols, or routing instances on a Juniper device 1 .
A firewall filter has a family attribute, which specifies the type of traffic that the filter can evaluate. The family attribute can be one of the following: inet, inet6, mpls, vpls, iso, or ethernet-switching 2 . The family inet firewall filter is used to evaluate IPv4 traffic, which is the most common type of Layer 3 traffic on a network.
To create a family inet firewall filter, you need to specify the appropriate match criteria and actions for each term in the filter. The match criteria can include various fields in the IPv4 header, such as source address, destination address, protocol, port number, or DSCP value. The actions can include accept, discard, reject, count, log, policer, or next term 3 .
To apply a firewall filter to Layer 3 traffic that is being sent between VLANs, you need to apply the filter to the appropriate IRB interface. An IRB interface is an integrated routing and bridging interface that provides Layer 3 functionality for a VLAN on a Juniper device. An IRB interface has an IP address that acts as the default gateway for the hosts in the VLAN. An IRB interface can also participate in routing protocols and forward packets to other VLANs or networks 4 .

Therefore, option C is correct, because you should create a family inet firewall filter with the appropriate match criteria and actions. Option D is correct, because you should apply the firewall filter to the appropriate IRB interface.
Option A is incorrect, because you should not create a family ethernet-switching firewall filter with the appropriate match criteria and actions. A family ethernet-switching firewall filter is used to evaluate Layer 2 traffic on a Juniper device. A family ethernet-switching firewall filter can only match on MAC addresses or VLAN IDs, not on IP addresses or protocols 5 .
Option B is incorrect, because you should not apply the firewall filter to the appropriate VLAN. A VLAN is a logical grouping of hosts that share the same broadcast domain on a Layer 2 network. A VLAN does not have an IP address or routing capability. A firewall filter cannot be applied directly to a VLAN; it must be applied to an interface that belongs to or connects to the VLAN 6 .
References:
1 : Firewall Filters Overview 2 : Configuring Firewall Filters 3 : Configuring Firewall Filter Match Conditions and Actions 4 : Understanding Integrated Routing and Bridging Interfaces 5 : Configuring Ethernet-Switching Firewall Filters 6 : Understanding VLANs
https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/task/firewall-filter-qfx- series-cli.html

NEW QUESTION: 9
Exhibit

Exhibit

```
{master:0}
user@switch> show vlans brief
Routing instance        VLAN name               Tag         Interfaces
default-switch          default                 1
                                                            ge-0/0/0.0*
                                                            ge-0/0/1.0*
                                                            ge-0/0/2.0*
                                                            ge-0/0/3.0*
                                                            ge-0/0/4.0*
                                                            ge-0/0/5.0*
```


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: (SHOW ANSWER)

* The exhibit shows the output of the command show vlans brief , which displays brief information about VLANs and their associated interfaces 1 .

* 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 1 . This can be verified by the command show interfaces terse , which displays the status of the interfaces 2 .

NEW QUESTION: 10

You are configuring a router as a new Level 1 router in an existing Level 1 area. The NET address for an existing Level 1 router is 49.0011.1111.2222.3333.00. The new router must form a Level 1 adjacency with the existing router.
What should be the area ID in this scenario?

A. 3333

B. 0011

C. 00

D. 49

Answer: (SHOW ANSWER)

NEW QUESTION: 11

Which two statements are correct about IS-IS? (Choose two.)

A. IS-IS only allows you to configure two areas.

B. IS-IS has two metric types and Junos sends both by default.

C. IS-IS sends a maximum metric value of 63 by default.

D. IS-IS uses IPv6 as its transport protocol in the Junos OS Implementation.

Answer: (SHOW ANSWER)

NEW QUESTION: 12

Exhibit.

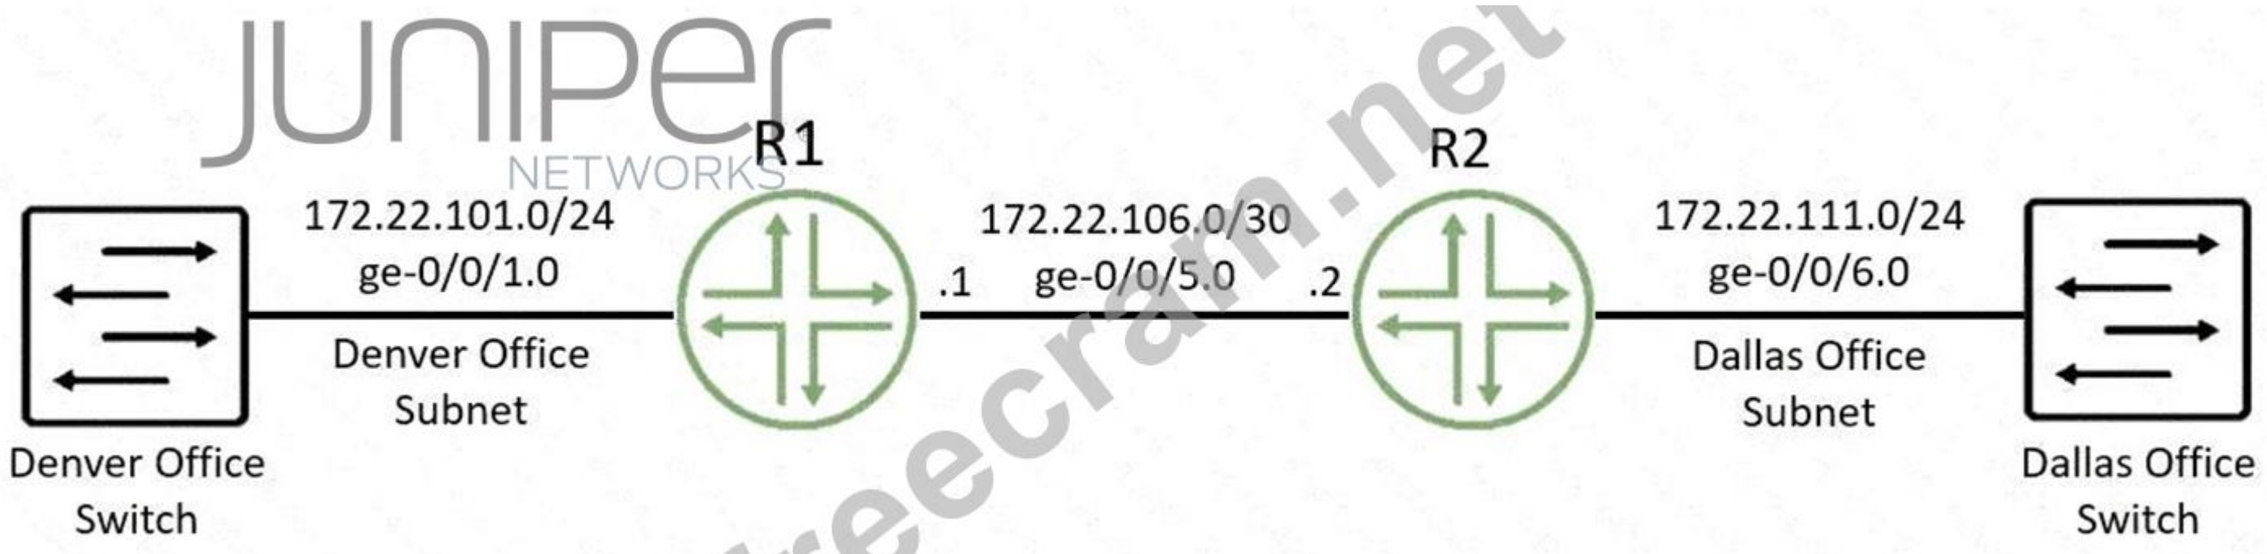

You are using OSPF to advertise the subnets that are used by the Denver and Dallas offices. The routers that are directly connected to the Dallas and Denver subnets are not advertising the connected subnets.

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

A. Create static routes on the switches using the local vMX router ' s loopback interface for the next hop.

B. Configure and apply a routing policy that redistributes the Dallas and Denver subnets using Type 5 LSAs.

C. Configure and apply a routing policy that redistributes the connected Dallas and Denver subnets.

D. Enable the passive option on the OSPF interfaces that are connected to the Dallas and Denver subnets.

Answer: (SHOW ANSWER)

The routers that are directly connected to the Dallas and Denver subnets are not advertising the connected subnets. This can be resolved by redistributing the connected subnets into OSPF 1 .

Option C suggests to configure and apply a routing policy that redistributes the connected Dallas and Denver subnets. This is correct because redistribution allows routes from one routing protocol to be communicated to another, and in this case, it allows the connected subnets to be advertised through OSPF 1 .

Option D suggests enabling the passive option on the OSPF interfaces that are connected to the Dallas and Denver subnets. This is also correct because in OSPF, a passive interface is an interface that belongs to the OSPF router, but does not send OSPF Hello packets 1 . It's typically used on an interface that you don't want to use for OSPF adjacencies, but you still want to advertise its IP address 1 .

Therefore, enabling passive interface can help in advertising the Dallas and Denver subnets.

NEW QUESTION: 13

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: (SHOW ANSWER)

* Option A is correct. The storm control feature is enabled in the factory-default configuration on EX Series switches 1 2 . 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 2 . On EX4300 switches, the factory default configuration enables storm control on all Layer 2 switch interfaces 1 .

* Option B is incorrect. The storm control feature does not require a special license on EX Series switches 3 4 .
* Option C is incorrect. There's no information available that suggests the storm control feature is not supported on aggregate Ethernet interfaces.
* Option D is incorrect. The storm control configuration applies to traffic at the ingress of an interface 5 , not just between the forwarding and control plane.

NEW QUESTION: 14
You have deployed a Juniper EX Series switch in the network. The switch receives a unicast frame destined to a MAC address that is not registered in the bridging table.
Which statement describes the behavior of the switch?
A. The frame is flooded out all ports in all VLANs configured on the switch.
B. The frame is flooded out all ports that are part of the same VLAN, except the receiving port.
C. The switch performs an ARP request to discover the MAC address of the destination host.
D. The switch sends an error message to the sender declaring that the host is unreachable.
Answer: (SHOW ANSWER)
According to the Juniper documentation1, when a switch receives a unicast frame destined to a MAC address that is not registered in the bridging table, it floods the frame out all ports that are part of the same VLAN, except the receiving port. This is the default behavior of a switch when it does not know the destination MAC address. The switch does not perform an ARP request, send an error message, or flood the frame to all VLANs. References: Understanding Unicast Flooding | Junos OS | Juniper Networks

NEW QUESTION: 15
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) 1 .
* An ISO NET address consists of three parts: an area ID, a system ID, and a selector 2 . 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 2 .
* 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 2 . 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: 16
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: 17
Which two statements about MACsec are true? (Choose two.)
A. MACsec functions on point-to-point connections.
B. MACsec functions on point-to-multipoint connections.
C. MACsec is an IEEE standard.
D. MACsec does not provide data integrity.
Answer: (SHOW ANSWER)

NEW QUESTION: 18
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: (SHOW ANSWER)
In Rapid Spanning Tree Protocol (RSTP), there are several port roles that determine the behavior of the port in the spanning tree 1 2 3 . The roles include root, designated, alternate, backup, and disabled 1 2 3 .
The discarding state is associated with the backup, alternate, and disabled roles 1 2 3 . 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 state 2 . Disabled ports are also in the discarding state 3 .
Therefore, options B, C, and D are correct.

NEW QUESTION: 19
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: (SHOW ANSWER)
Both aggregate and generated routes would work, but the best choice is generated routes.

NEW QUESTION: 20
Which two statements are correct about tunnels? (Choose two.)
A. BFD cannot be used to monitor tunnels.

B. Tunnel endpoints must have a valid route to the remote tunnel endpoint.
C. IP-IP tunnels are stateful.
D. Tunnels add additional overhead to packet size.
Answer: B,D (LEAVE A REPLY)
A tunnel is a connection between two computer networks, in which data is sent from one network to another through an encrypted link. Tunnels are commonly used to secure data communications between two networks or to connect two networks that use different protocols.
Option B is correct, because tunnel endpoints must have a valid route to the remote tunnel endpoint. A tunnel endpoint is the device that initiates or terminates a tunnel connection. For a tunnel to be established, both endpoints must be able to reach each other over the underlying network. This means that they must have a valid route to the IP address of the remote endpoint 1 .
Option D is correct, because tunnels add additional overhead to packet size. Tunnels work by encapsulating packets: wrapping packets inside of other packets. This means that the original packet becomes the payload of the surrounding packet, and the surrounding packet has its own header and trailer. The header and trailer of the surrounding packet add extra bytes to the packet size, which is called overhead. Overhead can reduce the efficiency and performance of a network, as it consumes more bandwidth and processing power 2 .
Option A is incorrect, because BFD can be used to monitor tunnels. 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. BFD can also be used to monitor the connectivity of tunnels, such as GRE, IPsec, or MPLS.
Option C is incorrect, because IP-IP tunnels are stateless. IP-IP tunnels are a type of tunnels that use IP as both the encapsulating and encapsulated protocol. IP-IP tunnels are simple and easy to configure, but they do not provide any security or authentication features. IP-IP tunnels are stateless, which means that they do not keep track of the state or status of the tunnel connection. Stateless tunnels do not require any signaling or negotiation between the endpoints, but they also do not provide any error detection or recovery mechanisms.
References:
1 : What is Tunneling? | Tunneling in Networking 2 : What Is Tunnel In Networking, Its Types, And Its Benefits? : [Configuring Bidirectional Forwarding Detection] : [IP-IP Tunneling]

NEW QUESTION: 21
Exhibit.

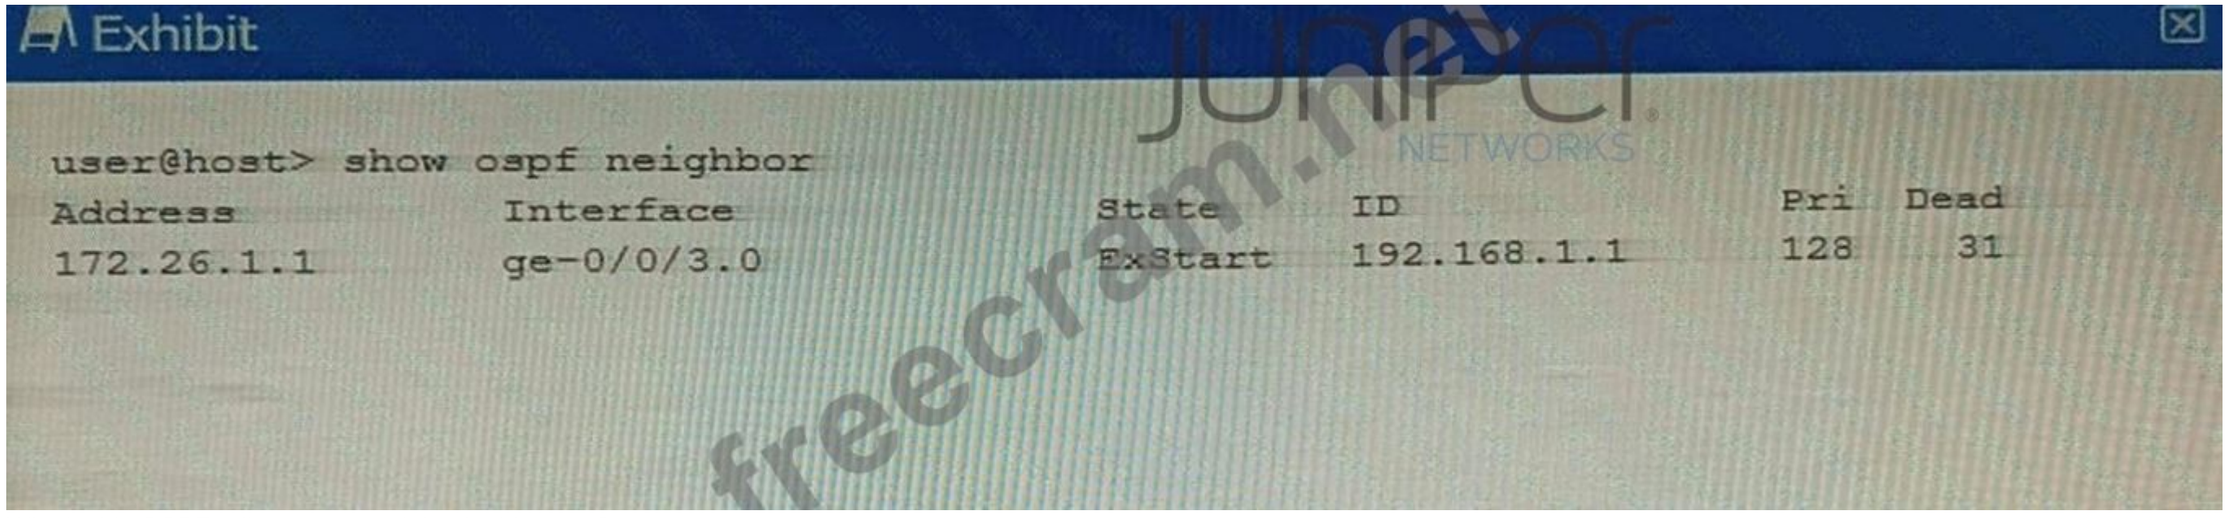

Why is this OSPF adjacency remaining in this state?
A. A subnet mask mismatch exists between the OSPF neighbors.
B. An MTU mismatch exists between the OSPF neighbors.
C. A hello interval mismatch exists between the OSPF neighbors.
D. An area ID mismatch exists between the OSPF neighbors
Answer: (SHOW ANSWER)

* The exhibit shows the output of the command show ospf neighbor , which displays information about the OSPF neighbors on a router 1 .
* The output shows that the OSPF neighbor with the address 172.26.1.1 and the interface ge-0/0/3.0 is in the Exstart state 1 .
* The Exstart state is the fourth state in the OSPF neighbor formation process, after Down, Init, and 2- Way states 2 . In this state, the OSPF neighbors establish a master-slave relationship and exchange database description (DBD) packets, which contain summaries of their link-state databases 2 .
* The most common reason for OSPF neighbors to be stuck in the Exstart state is an MTU mismatch between the interfaces 3 . MTU stands for maximum transmission unit, which is the largest size of a packet that can be transmitted on a network segment 4 . If the MTU values of two OSPF neighbors are different, they may not be able to exchange DBD packets successfully, as some packets may be dropped or fragmented due to their size exceeding the MTU limit 3 .
* To solve this problem, you need to ensure that the MTU values of both OSPF neighbors are the same or compatible. You can use the command show interfaces to display the MTU value of an interface 5 .
You can also use the command ping with the do-not-fragment option to test the MTU size between two routers. You can change the MTU value of an interface by using the command set interfaces interface-name mtu mtu-value in configuration mode 5 .

NEW QUESTION: 22
You are configuring an IS-IS IGP network and do not see the IS-IS adjacencies established. In this scenario, what are two reasons for this problem? (Choose two.)
A. MTU is not at least 1492 bytes.
B. IP subnets are not a /30 address.
C. The Level 2 routers have mismatched areas.
D. The lo0 interface is not included as an IS-IS interface.
Answer: (SHOW ANSWER)
Option A suggests that the MTU is not at least 1492 bytes. This is correct because IS-IS requires a minimum MTU of 1492 bytes to establish adjacencies 1 . If the MTU is less than this, IS-IS adjacencies will not be established 1 .
Option D suggests that the lo0 interface is not included as an IS-IS interface. This is also correct because the loopback interface (lo0) is typically used as the router ID in IS-IS 1 . If the loopback interface is not included in IS-IS, it could prevent IS-IS adjacencies from being established 1 .
Therefore, options A and D are correct.

NEW QUESTION: 23
You implemented the MAC address limit feature with the shutdown action on all interfaces on your switch.
In this scenario, which statement is correct when a violation occurs?
A. By default, you must manually clear the violation for the interface to send and receive traffic again.
B. By default, the violation will automatically be cleared after 300 seconds and the interface will resume sending and receiving traffic for all learned devices.
C. By default, devices that are learned before the violation occurs are still allowed to send and receive traffic through the specific interface.
D. By default, the interface will continue to send and receive traffic for all connected devices after a violation has occurred.
Answer: (SHOW ANSWER)
When the MAC address limit feature with the shutdown action is implemented on a switch, if a violation occurs, the interface is disabled and a system log entry is generated 1 . If the switch has been configured with the port-error-disable statement, the disabled interface recovers automatically upon expiration of the specified disable timeout 1 . However, if the switch has not been configured for auto-recovery from port error disabled conditions, you must manually clear the violation by running the clear ethernet-switching port-error command for the interface to send and receive traffic again 1 . This explanation is based on the Enterprise Routing and Switching Specialist (JNCIS-ENT) documents and learning resources available at Juniper Networks 1 .

NEW QUESTION: 24
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. Set the local-preference attribute to a higher value tor ISP-2 than ISP-1.
B. Configure the gateway for ISP-1 with a higher peer ID than the gateway for ISP-2.
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: 25
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: (SHOW ANSWER)
* 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 1 . DAI discards any ARP packets that do not match the database or have invalid formats 1 .
* 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 2 . DHCP snooping also builds a database of MAC addresses, IP addresses, lease times, and binding types for each client 2 .

NEW QUESTION: 26
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. Route reflectors require that the network is connected in a physical full mesh.
D. The cluster ID is not required to use route reflectors.
Answer: (SHOW ANSWER)
A route reflector is a BGP router that is allowed to break the iBGP loop avoidance rule and advertise routes received from an iBGP peer to another iBGP peer under specific conditions 1 . A route reflector can have three types of peerings: EBGP neighbor, IBGP client neighbor, and IBGP non-client neighbor 2 . A cluster is a group of route reflectors and their clients that are fully meshed 3 . The cluster ID is a 4-byte value that identifies the route reflector cluster and is carried in the ORIGINATOR_ID attribute of the BGP update message 4 . The cluster ID is used to prevent routing loops among route reflectors and their clients. If a route reflector receives a route with its own cluster ID, it discards the route 5 . Therefore, the cluster ID is required to prevent routing loops when using route reflectors.
References:
1 : What is BGP Route Reflector - Explanation & Discussion (Cisco Example)
2 : BGP Route Reflector - NetworkLessons.com
3 : BGP Route Reflectors (RR) - Cisco Community
4 : Junos OS Routing Protocols Library for Routing Devices
5 : Junos OS Routing Protocols Library for Routing Devices

NEW QUESTION: 27
Click the Exhibit button.

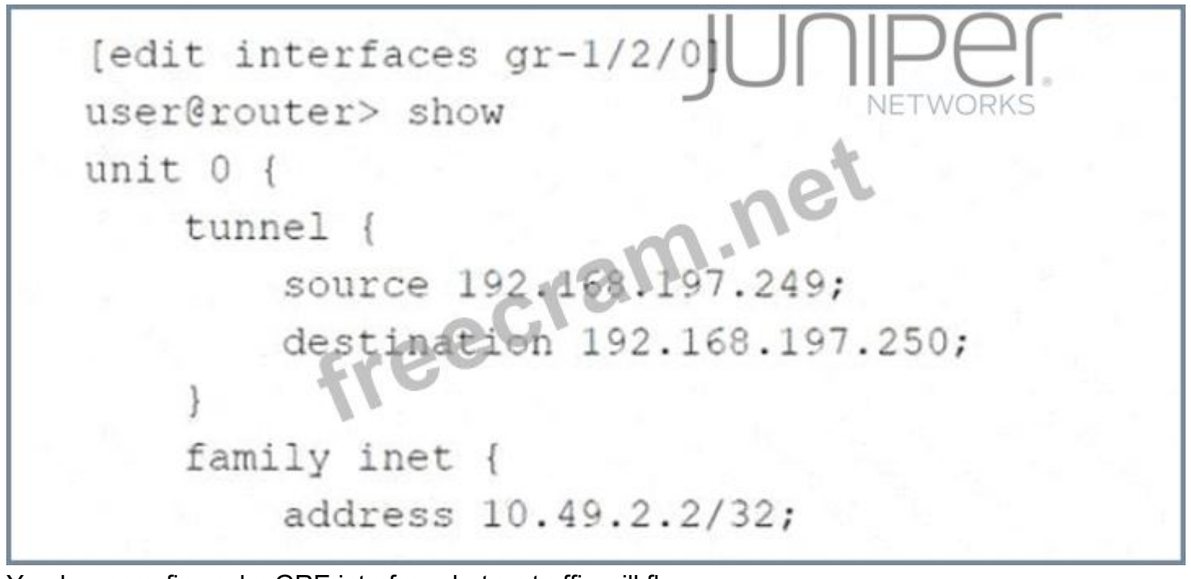

You have configured a GRE interface, but no traffic will flow.
Referring to the exhibit, which statement is true?
A. The IP subnet mask cannot be a /32.
B. The gr-1/2/0 interface is an invalid GRE interface.
C. The unit number cannot be 0.
D. The source and destination addresses cannot have the same third octet.
Answer: (SHOW ANSWER)
According to the Juniper documentation 1 , the source and destination addresses of a GRE tunnel interface must be in different subnets. If they are in the same subnet, the router will not be able to route the GRE packets correctly. In the exhibit, the source and destination addresses are both in the 192.168.197.0/24 subnet, which is invalid for GRE tunneling. Therefore, the correct answer is D.
References: Configuring GRE Tunnel Interfaces | Junos OS

NEW QUESTION: 28
Exhibit

Exhibit

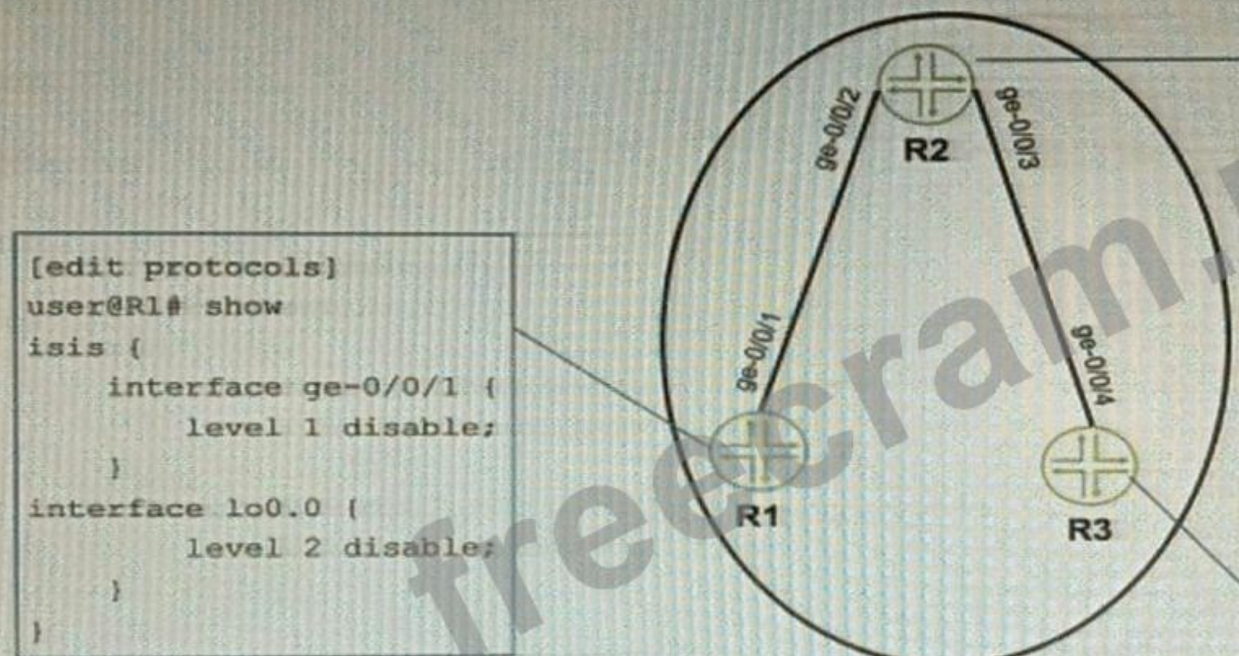


```
[edit protocols]
user@R1# show
isis {
    interface ge-0/0/1 {
        level 1 disable;
    }
interface lo0.0 {
        level 2 disable;
    }
}
```

```
[edit protocols]
user@R2# show
isis {
    interface ge-0/0/2 {
        level 2 disable;
    }
    interface ge-0/0/3 {
        level 2 disable;
    }
    interface lo0.0 {
        level 1 disable;
    }
}
```

```
[edit protocols]
user@R3# show
isis {
    interface ge-0/0/4 {
        level 1 disable;
    }
interface lo0.0 {
        level 2 disable;
    }
}
```


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 interface 1 . Level 1 systems route within an area 2 . If the destination is outside an area, Level 1 systems route toward a Level 2 system 2 . 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 Networks 1 and Cisco 3 4 .

NEW QUESTION: 29
You have DHCP snooping enabled but no entries are automatically created in the snooping database for an interface on your EX Series switch. What are two reasons for the problem? (Choose two.)
A. The device that is connected to the interface has performed a DHCPRELEASE.
B. MAC limiting is enabled on the interface.
C. The device that is connected to the interface has a static IP address.
D. Dynamic ARP inspection is enabled on the interface.
Answer: (SHOW ANSWER)
The DHCP snooping feature in Juniper Networks' EX Series switches works by building a binding database that maps the IP address, MAC address, lease time, binding type, VLAN number, and interface information 1
. This database is used to filter and validate DHCP messages from untrusted sources 1 .
However, there are certain conditions that could prevent entries from being automatically created in the snooping database for an interface:
* MAC limiting: If MAC limiting is enabled on the interface, it could potentially interfere with the operation of DHCP snooping. MAC limiting restricts the number of MAC addresses that can be learned on a physical interface to prevent MAC flooding attacks 1 . This could inadvertently limit the number of DHCP clients that can be learned on an interface, thus preventing new entries from being added to the DHCP snooping database.
* Static IP address : If the device connected to the interface is configured with a static IP address, it will not go through the DHCP process and therefore will not have an entry in the DHCP snooping database
1 . The DHCP snooping feature relies on monitoring DHCP messages to build its database 1 , so devices with static IP addresses that do not send DHCP messages will not have their information added.
Therefore, options B and C are correct. Options A and D are not correct because performing a DHCPRELEASE would simply remove an existing entry from the database 1 , and Dynamic ARP inspection (DAI) uses the information stored in the DHCP snooping binding database but does not prevent entries from being created 1 .

NEW QUESTION: 30
Exhibit

```
user@R1> show bgp neighbor
Peer: 10.32.1.2+63645 AS 65401 Local: 10.32.1.1+179 AS 65400
  Description: EBGP peering to 10.32.1.2
  Group: IPCLOS_eBGP       Routing-Instance: master
  Forwarding routing-instance: master
  Type: External   State: Established   Flags: <Sync>Last
  State: OpenConfirm   Last Event: RecvkeepAlive
  Last Error: None
  Export: [ IPCLOS_BGP_EXP ] Import: [ IPCLOS_BGP_IMP ]
  Options: <Preference PeerAS Multipath LocalAS Refresh>
  Options: <VpnApplyExport MtuDiscovery MultipathAs BfdEnabled>
  Holdtime: 90 Preference: 170 Local AS: 65400 Local System AS: 0
  Number of flaps: 0
  Peer ID: 10.52.100.2  Local ID: 10.52.100.1  Active Holdtime: 90
  Keepalive Interval: 30  Group index: 0   Peer index: 0   SNMP
index: 0
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: enabled, up
  Local Interface: ge-0/0/1.0
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65401)
  Peer does not support Addpath
  Table inet.0 Bit: 20000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              6
    Received prefixes:            9
    Accepted prefixes:            9
    Suppressed due to damping:    0
    Advertised prefixes:          22
  Last traffic (seconds): Received 22 Sent 10 Checked 69617
  Last Input messages: Total 2568 Updates 4 Refreshes 0 Octets 48991
  Output messages:     Total 2572 Updates 8 Refreshes 0 Octets 49362
  Output Queue[1]: 0           (inet.0, inet-unicast)
```

You are a network operator troubleshooting BGP connectivity.

Which two statements are correct about the output shown in the exhibit? (Choose two.)

A. Peer 10.32.1.2 is configured for AS 63645.
B. The BGP session is not established.
C. The R1 is configured for AS 65400.
D. The routers are exchanging IPv4 routes.
Answer: (SHOW ANSWER)
The output of the show bgp neighbor command shows the following information:
* The local AS number of R1 is 65400, as indicated by the line "Local: 10.32.1.1+179 AS 65400".
* The peer AS number of 10.32.1.2 is 63645, as indicated by the line "Peer: 10.32.1.2+179 AS 63645".
* The BGP session is established, as indicated by the line "State: Established".
* The routers are exchanging IPv4 routes, as indicated by the line "Type: External State: Established Flags: < Sync > " and the table showing the number of prefixes received and sent for inet.0 routing table. References: Enterprise Routing and Switching, Specialist (JNCIS-ENT) - Juniper Networks ,
[show bgp neighbor - Technical Documentation - Support - Juniper Networks]

NEW QUESTION: 31
What is the default keepalive time for BGP?
A. 10 seconds
B. 60 seconds
C. 30 seconds
D. 90 seconds
Answer: (SHOW ANSWER)
Default BGP keepalive timer in Junos is 30 seconds
The default keepalive time for BGP is 30 seconds, which means that BGP peers send keepalive messages to each other every 30 seconds to verify the liveness of the BGP session. If no keepalive message is received from a peer within the hold time, which is 90 seconds by default, the peer is considered to be unavailable and the BGP session is closed. The keepalive time and the hold time can be configured manually for each BGP peer or peer group, but they must match on both ends of the BGP session. The keepalive time should be one third of the hold time for optimal performance. Therefore, option C is correct and the other options are incorrect. References: BGP keepalive messages and hold time - Juniper Networks

NEW QUESTION: 32
What is the maximum allowable MTU size for a default GRE tunnel without IPv4 traffic fragmentation?
A. 1496 bytes
B. 1480 bytes
C. 1500 bytes
D. 1476 bytes
Answer: (SHOW ANSWER)
The maximum allowable MTU size for a default GRE tunnel without IPv4 traffic fragmentation is 1476 bytes

1 . This is because GRE packets are formed by the addition of the original packets and the required GRE headers 1 . These headers are 24-bytes in length and since these headers are added to the original frame, depending on the original size of the packet we may run into IP MTU problems 1 . The most common IP MTU is 1500-bytes in length (Ethernet) 1 . When the tunnel is created, it deducts the 24-bytes it needs to encapsulate the passenger protocols and that is the IP MTU it will use 1 . For example, if we are forming a tunnel over FastEthernet (IP MTU 1500) the IOS calculates the IP MTU on the tunnel as: 1500-bytes from Ethernet - 24-bytes for the GRE encapsulation = 1476-Bytes 1 .

NEW QUESTION: 33
Click the Exhibit button.

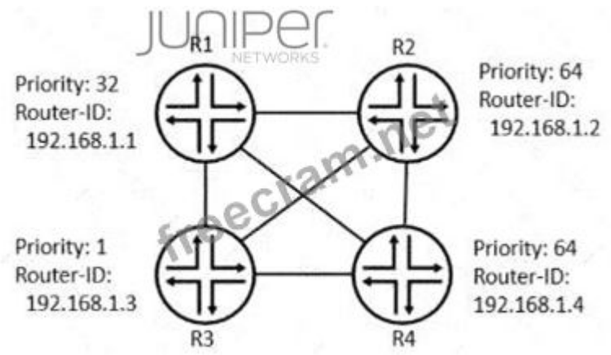

R1 was started at 4:00 PM. R2 and R3 were started at 6:00 PM. R4 was started at 10:00 PM. All routers are in Area 0 and there are no point-to-point links.
Referring to the exhibit, which router is the designated router?
A. R1
B. R4
C. R3
D. R2
Answer: (SHOW ANSWER)

NEW QUESTION: 34
Which statement is correct about IP-IP tunnels?
A. IP-IP tunnels only support encapsulating IP traffic.
B. IP-IP tunnels only support encapsulating non-IP traffic.
C. The TTL in the inner packet is decremented during transit to the tunnel endpoint.
D. There are 24 bytes of overhead with IP-IP encapsulation.
Answer: (SHOW ANSWER)
IP-IP tunnels encapsulate IP packets within other IP packets, allowing them to traverse networks that wouldn ' t normally support the original IP packets. This encapsulation involves the TTL (Time-to-Live) of the inner packet being decremented as it traverses through the tunnel towards its endpoint.

NEW QUESTION: 35
Exhibit

```
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
                     > 10.100.25.2 via ge-0/0/1.0
```


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?

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: (SHOW ANSWER)

* The exhibit shows the output of the command show ip bgp , which displays information about the BGP routes in the routing table 1 . 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) 2 . A higher local preference means a more preferred route 2 .
* 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 3
. The first attribute that is compared is weight, which is a Cisco-specific attribute that is local to the router 3 . If the weight is equal or not set, the next attribute that is compared is local preference 3 .
* In this case, both routes have the same weight of 0, which means that they are learned from external BGP (eBGP) peers 3 . 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 3 . The other attributes, such as origin code and AS path, are not considered in this case.

NEW QUESTION: 36

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)
In OSPF, a DR (Designated Router) and a BDR (Backup Designated Router) are elected on each multi-access network, such as Ethernet or Frame Relay. The DR and BDR are responsible for exchanging routing information with other routers on the same network segment. The election is based on two criteria: the OSPF priority and the router ID. The router with the highest priority becomes the DR, and the router with the second-highest priority becomes the BDR. If the priorities are equal, the router with the highest router ID becomes the DR, and the router with the second-highest router ID becomes the BDR. The default priority is 1, and the router ID is either manually configured or derived from the highest IP address on any interface12.
In this scenario, R1 has a priority of 32 and a router ID of 192.168.1.2, and R2 has a priority of 64 and a router ID of 192.168.1.1. Since R2 has a higher priority than R1, it will become the BDR. The router ID does not matter in this case, because the priority is different. Therefore, the correct answer is C.
The other options are not correct because:
A). At least two routers are required for a DR/BDR election, not three2.
B). Router IDs do not have to match for an adjacency to form, they just have to be unique3.
D). R1 will not be the BDR, because it has a lower priority than R212.
References: Understanding OSPF Areas | Junos OS, OSPF DR/BDR Election explained - NetworkLessons.
com, OSPF adjacency - Cisco Community

NEW QUESTION: 37
Which statement about aggregate routes is correct?
A. Aggregate routes can only be used for static routing but not for dynamic routing protocols.
B. Aggregate routes are automatically generated for all of the subnets in a routing table.
C. Aggregate routes are always preferred over more specific routes, even when the specific routes have a better path.
D. Aggregate routes are used for advertising summarized network prefixes.
Answer: (SHOW ANSWER)
Aggregate routes are used for advertising summarized network prefixes 1 2 . They help minimize the number of routing tables in an IP network by consolidating selected multiple routes into a single route advertisement 1
. This approach is in contrast to non-aggregation routing, in which every routing table contains a unique entry for each route 1 .
Therefore, option D is correct. Options A, B, and C are not correct because:
* Aggregate routes can be used with both static routing and dynamic routing protocols 1 .
* Aggregate routes are not automatically generated for all of the subnets in a routing table. They need to be manually configured 1 .
* Aggregate routes are not always preferred over more specific routes. The route selection process in Junos OS considers several factors, including route preference and metric, before determining the active route 1 .

NEW QUESTION: 38
Exhibit

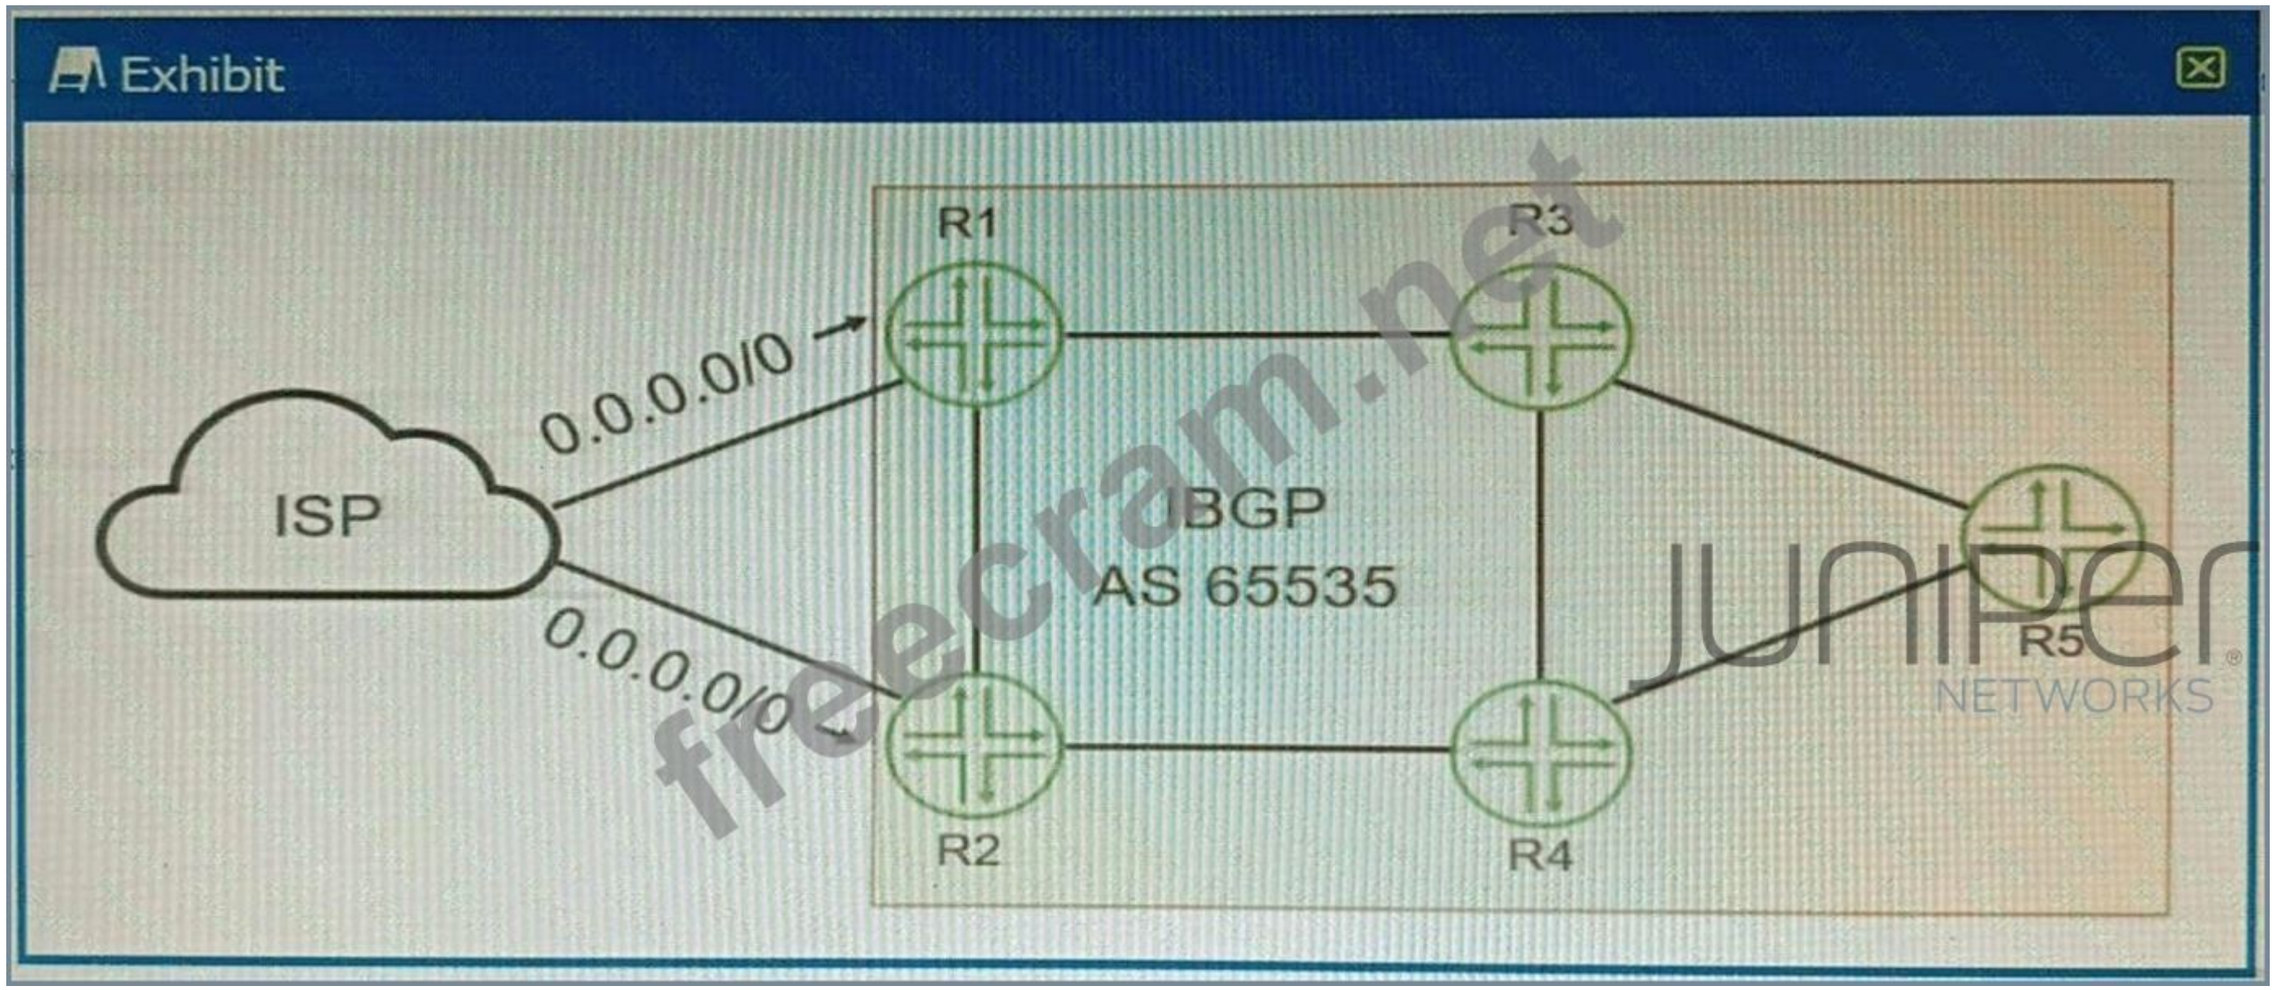

Your ISP is announcing a default route to both R1 and R2. You want your network routers to forward all Internet traffic through the R1 device Which BGP attribute would you use?

A. MED

B. next-hop

C. local preference

D. origin

Answer: C (LEAVE A REPLY)

The BGP attribute that you would use to forward all Internet traffic through the R1 device is the local preference 1 .

The local preference is an attribute that is used within an autonomous system (AS) and exchanged between iBGP routers 1 . It is used to select an exit point from the AS 1 . The path with the highest local preference is preferred 1 . By setting a higher local preference for the routes received from R1, you can make R1 the preferred exit point for all Internet traffic 1 .

NEW QUESTION: 39

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)

IS-IS (Intermediate System to Intermediate System) is a routing protocol that is designed to move information efficiently within a computer network 1 2 . It supports multiple protocol topologies, including IPv4, IPv6, and multicast 1 2 . Therefore, options C, E, and D are correct.

NEW QUESTION: 40

What are two purposes of an aggregate route? (Choose two.)

A. to decrease the number of route advertisements

B. to hide internal routes from external peers

C. to allow external peers to see internal routes

D. to increase the number of route advertisements

Answer: (SHOW ANSWER)

An aggregate route is a route that summarizes multiple more specific routes into a single route advertisement.

For example, an aggregate route to 10.0.0.0/8 can represent many routes to 10.x.x.x/24 subnets. An aggregate route has several purposes, such as:

To decrease the number of route advertisements: By consolidating multiple routes into one, an aggregate route reduces the size of the routing table and the amount of routing updates that need to be exchanged between routers. This can improve the scalability and efficiency of the network12.

To hide internal routes from external peers: By advertising only the aggregate route and not the more specific routes, an aggregate route can provide a level of abstraction and security for the internal network topology. This can prevent external peers from knowing the details of the internal routes and avoid potential attacks or misconfigurations13.

The other options are not correct because:

C). To allow external peers to see internal routes: This is the opposite of what an aggregate route does. An aggregate route hides the internal routes from external peers, not exposes them13.

D). To increase the number of route advertisements: This is also the opposite of what an aggregate route does. An aggregate route decreases the number of route advertisements, not increases them12.

References: Configuring Route Aggregation | Junos OS, JunOS Aggregate and Generate Routes - Network Direction, Aggregate Routes vs Generated Routes. | Junos OS - Juniper Networks

NEW QUESTION: 41

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: (SHOW ANSWER)

The default MAC age-out timer on an EX Series switch is 300 seconds 1 2 . 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 3 1 . This configuration can influence efficiency of network resource use by affecting the amount of traffic that is flooded to all interfaces 1 . When traffic is received for MAC addresses no longer in the Ethernet routing table, the router floods the traffic to all interfaces 1 .

NEW QUESTION: 42

What are two characteristics of RSTP alternate ports? (Choose two.)

A. RSTP alternate ports block traffic while receiving superior BPDUs from a neighboring switch.

B. RSTP alternate ports provide an alternate lower cost path to the root bridge.

C. RSTP alternate ports provide an alternate higher cost path to the root bridge.

D. RSTP alternate ports are active ports used to forward frames toward the root bridge.

Answer: (SHOW ANSWER)

* A is correct because RSTP alternate ports block traffic while receiving superior BPDUs from a neighboring switch. An alternate port is a backup port for a root port, which means it receives better BPDUs from another bridge than the current root port 1 . However, an alternate port does not forward any traffic, as it is in a discarding state 2 . It only listens to BPDUs and waits for the root port to fail. If the root port fails, the alternate port can immediately transition to a forwarding state and become the new root port 1 .
* C is correct because RSTP alternate ports provide an alternate higher cost path to the root bridge. An alternate port is selected based on the same criteria as the root port, which are the lowest bridge ID, the lowest path cost, the lowest sender port ID, and the lowest receiver port ID 3 . However, an alternate port receives a higher cost BPDU than the root port, otherwise it would be the root port itself 1 .
Therefore, an alternate port provides an alternate higher cost path to the root bridge than the root port.

NEW QUESTION: 43
What does the MAC limiting feature do on EX Series switches?
A. It limits the number of MAC addresses learned on an access port.
B. It limits the number of MAC addresses learned on a trunk port.
C. It limits the acceptable values for a MAC address to a specified range.
D. It limits the time a learned MAC address stays in the MAC routing table.
Answer: (SHOW ANSWER)
MAC limiting protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). This feature is enabled on interfaces (ports). MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access port. The MAC limit value in the EX-series switch's default configuration is five MAC addresses. You can configure the MAC limit on an interface, on an interface in a VLAN, or on all interfaces. You can also specify the action to be taken when the MAC limit is exceeded, such as dropping packets, logging messages, or shutting down the interface.
Additionally, you can configure specific "allowed" MAC addresses for the access port. Any MAC address that is not in the list of configured addresses is not learned. Allowed MAC binds MAC addresses to a VLAN so that the address does not get registered outside the VLAN. References:
[EX] How to control MAC address access through Junos OS for EX Series switches1 Configuring MAC Limiting2 Understanding MAC Limiting and MAC Move Limiting3