JN0-351 Dumps

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

https://www.certleader.com/JN0-351-dumps.html

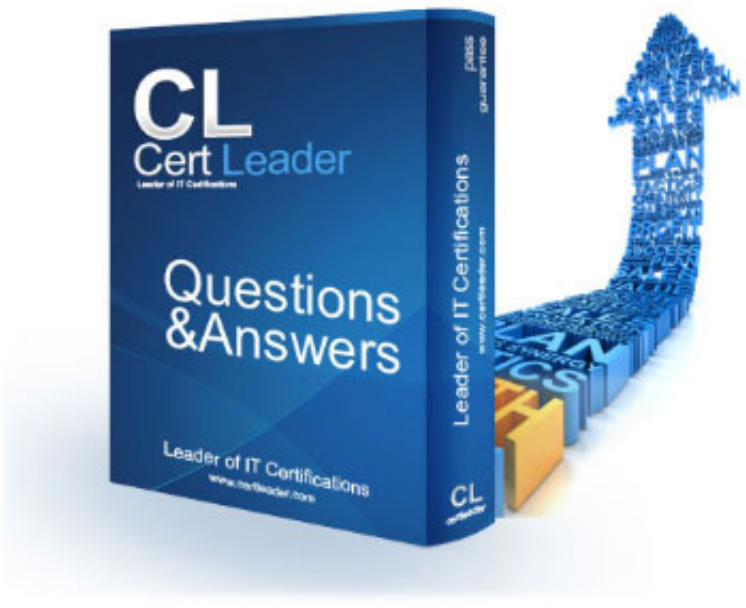

NEW QUESTION 1
Exhibit.

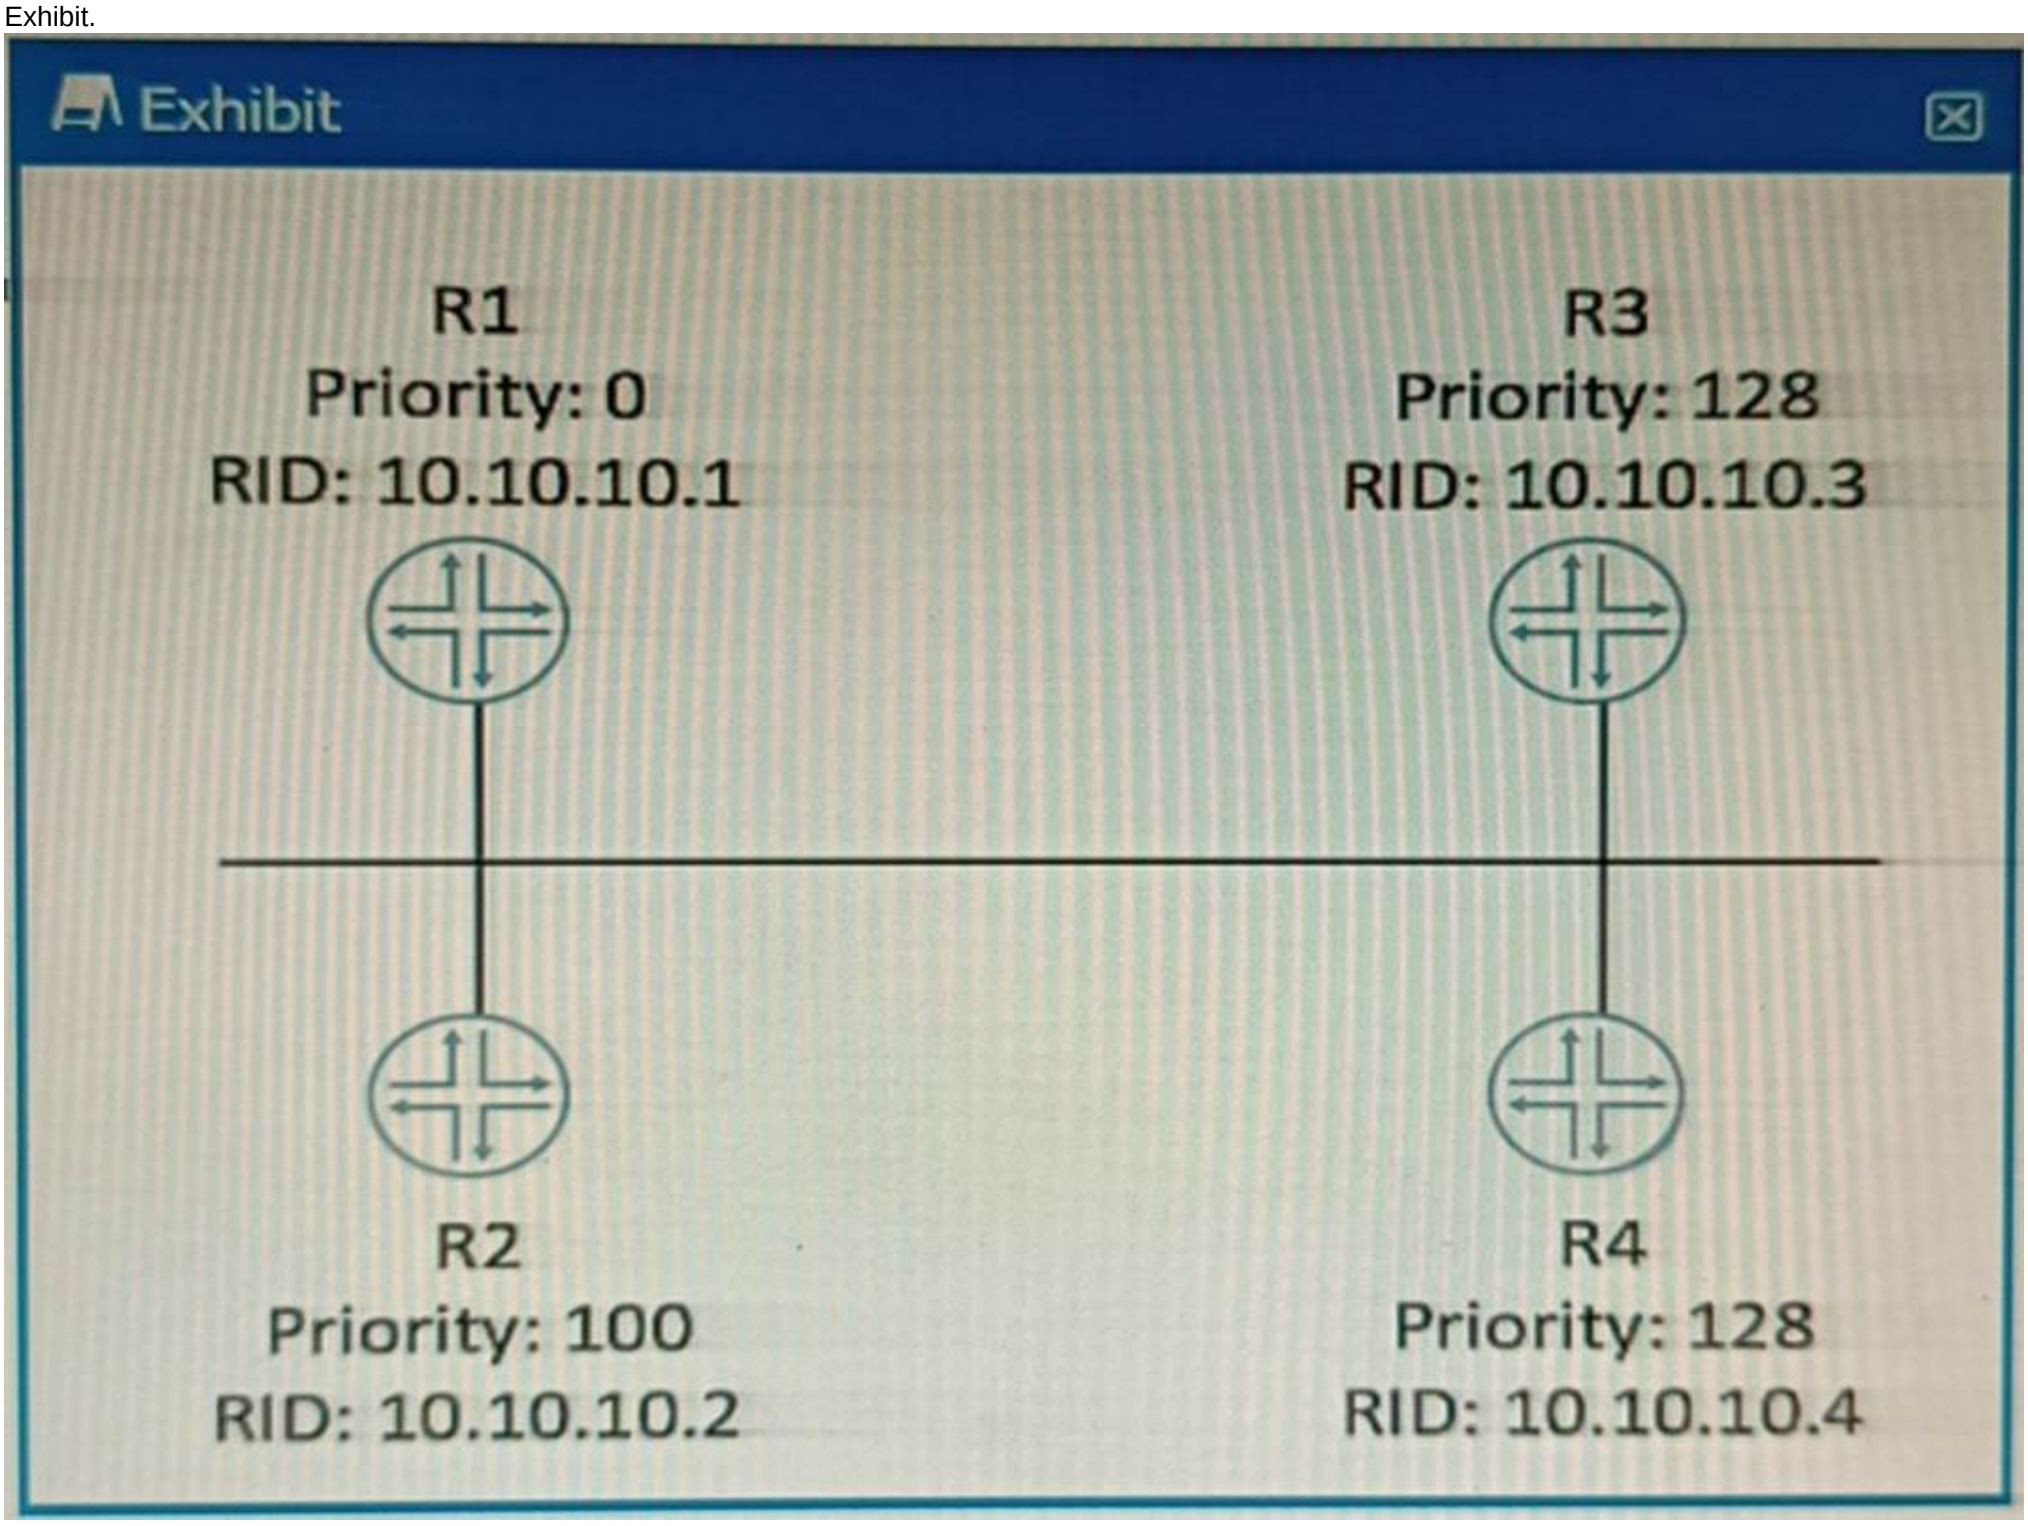

Which router will become the OSPF BDR if all routers are powered on at the same time?

A. R4
B. R1
C. R3
D. R2

Answer: A

Explanation:
OSPF DR/BDR election is a process that occurs on multi-access data links. It is intended to select two OSPF nodes: one to be acting as the Designated Router (DR), and another to be acting as the Backup Designated Router (BDR).The DR and BDR are responsible for generating network LSAs for the multi-access network and synchronizing the LSDB with other routers on the same network1.
The DR/BDR election is based on two criteria: the OSPF priority and the router ID. The OSPF priority is a value between 0 and 255 that can be configured on each interface participating in OSPF. The default priority is 1. A priority of 0 means that the router will not participate in the election and will never become a DR or BDR. The router with the highest priority will become the DR, and the router with the second highest priority will become the BDR. If there is a tie in priority, then the router ID is used as a tie-breaker. The router ID is a 32-bit number that uniquely identifies each router in an OSPF domain.It can be manually configured or automatically derived from the highest IP address on a loopback interface or any active interface2.
In this scenario, all routers have the same priority of 1, so the router ID will determine the outcome of the election. The router IDs are shown in the exhibit as RID values. The highest
RID belongs to R4 (10.10.10.4), so R4 will become the DR. The second highest RID belongs to R3 (10.10.10.3), so R3 will become the BDR.
References:
1:OSPF DR/BDR Election: Process, Configuration, and Tuning2:OSPF Designated Router (DR) and Backup Designated Router (BDR)

NEW QUESTION 2
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: D

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 3
Exhibit.

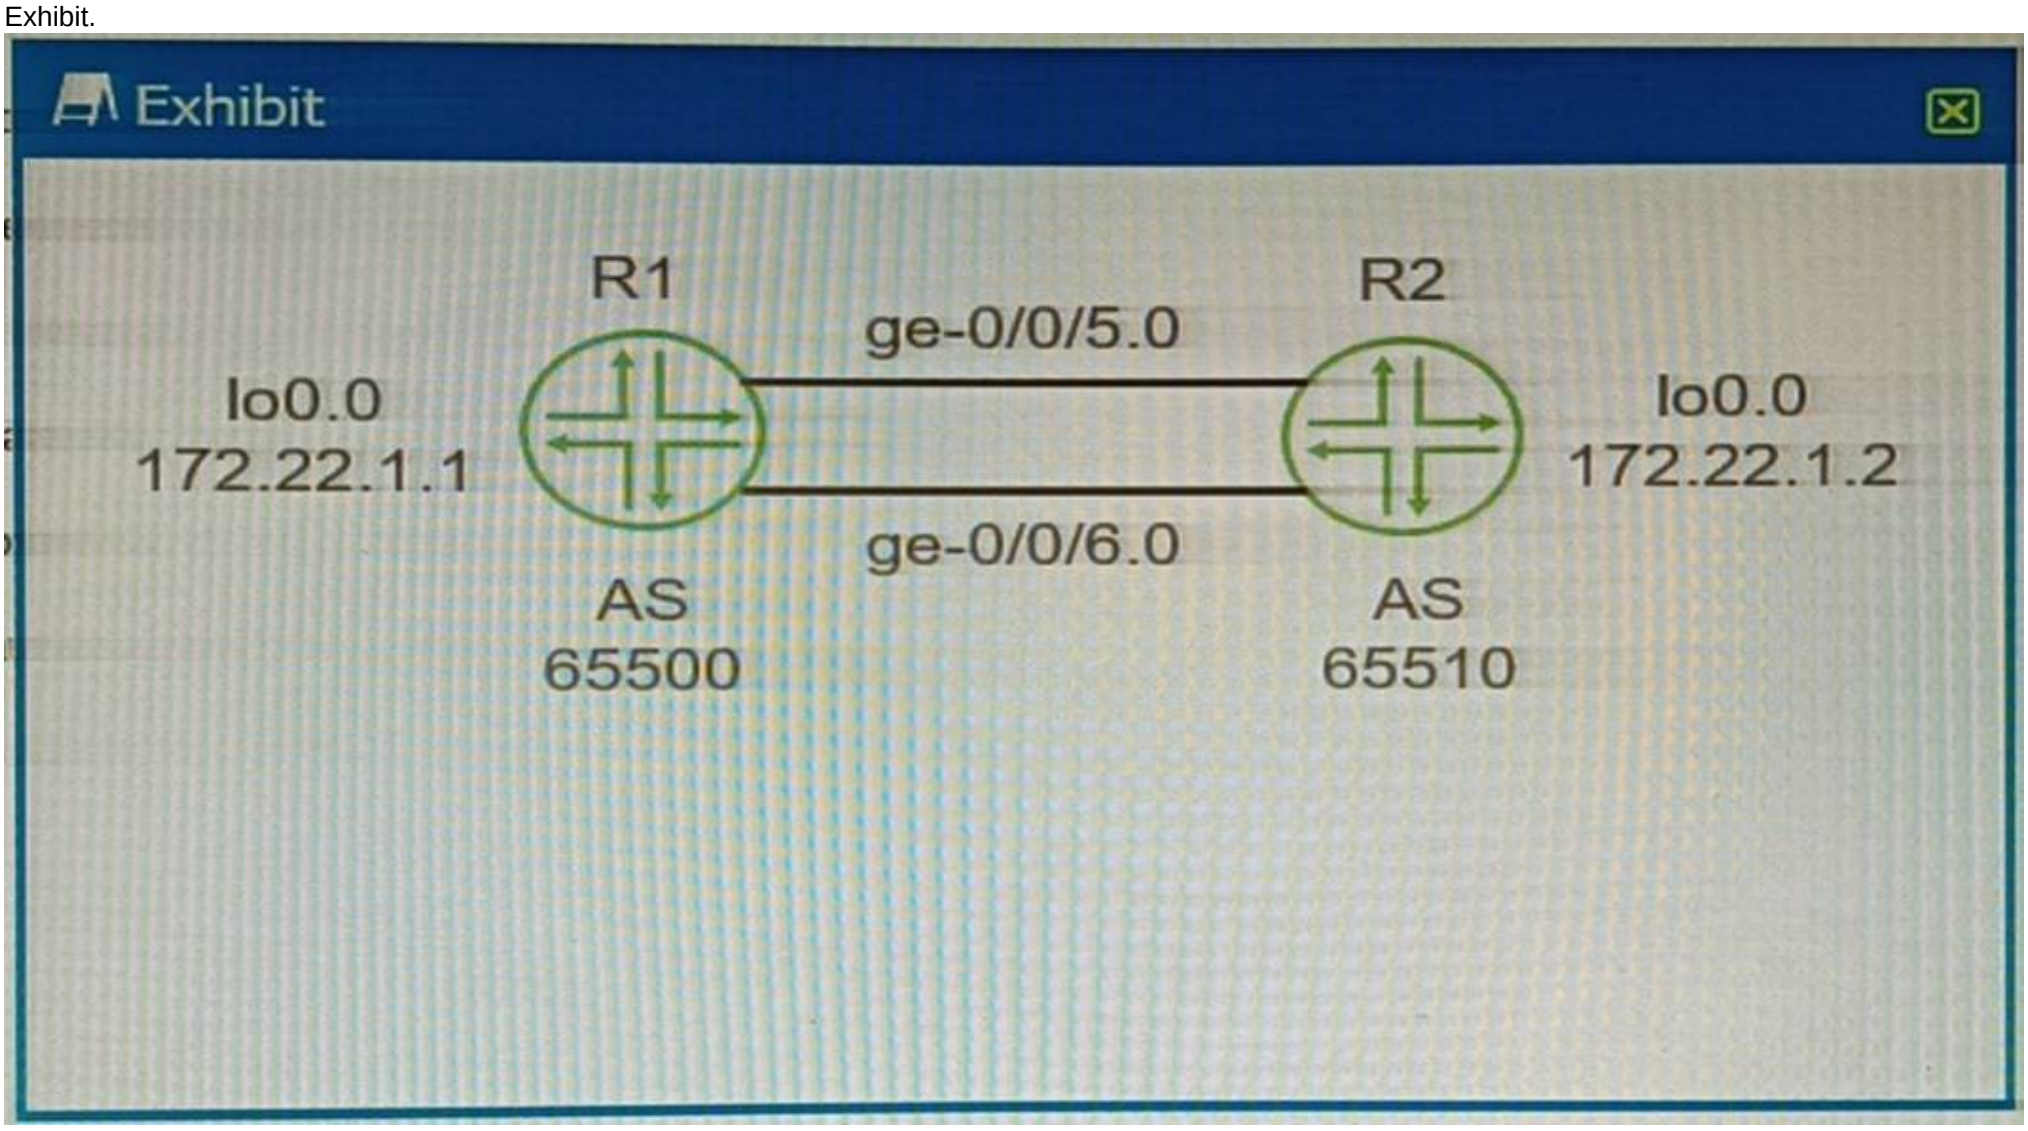

You want to enable redundancy for the EBGP peering between the two routers shown in the exhibit. Which three actions will you perform in this scenario? (Choose three.)

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: ABC

Explanation:
? 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 4
Exhibit

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
  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: BC

Explanation:
Option B suggests that the BGP session is not established. This is correct because in the output, the state of the BGP session is shown as ??Idle??. In BGP, an ??Idle?? state means that the BGP session is not currently established1.
Option C suggests that R1 is configured for AS 65400. This is also correct because in the output, it??s shown that the local AS number is 654001. The local AS number represents the Autonomous System (AS) number of the router on which you??re checking the BGP session1.

NEW QUESTION 5
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: BD

Explanation:
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 endpoint1.
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 power2.
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 6
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

Explanation:
The default MAC age-out timer on an EX Series switch is 300 seconds12. The MAC age-out timer is the maximum time that an entry can remain in the MAC table before it ??ages out,?? or is removed31. This configuration can influence efficiency of network resource use by affecting the amount of traffic that is flooded to all interfaces1. When traffic is received for MAC addresses no longer in the Ethernet routing table, the router floods the traffic to all interfaces1.

NEW QUESTION 7
Which two events cause a router to advertise a connected network to OSPF neighbors? (Choose two.)

A. When an OSPF adjacency is established.
B. When an interface has the OSPF passive option enabled.
C. When a static route to the 224.0.0.6 address is created.
D. When a static route to the 224.0.0.5 address is created.

Answer: AD

Explanation:
? A is correct because when an OSPF adjacency is established, a router will advertise a connected network to OSPF neighbors. An OSPF adjacency is a logical relationship between two routers that agree to exchange routing information using the OSPF protocol1. To establish an OSPF adjacency, the routers must be in the same area, have compatible parameters, and exchange hello packets1. Once an OSPF adjacency is formed, the routers will exchange database description (DBD) packets, which contain summaries of their link-state databases (LSDBs)1. The LSDBs include information about the connected networks and their costs2. Therefore, when an OSPF adjacency is established, a router will advertise
a connected network to OSPF neighbors through DBD packets.
? D is correct because when a static route to the 224.0.0.5 address is created, a router will advertise a connected network to OSPF neighbors. The 224.0.0.5 address is the multicast address for all OSPF routers3. A static route to this address can be used to send OSPF hello packets to all OSPF neighbors on a network segment3. This can be useful when the network segment does not support multicast or when the router does not have an IP address on the segment3. When a static route to the 224.0.0.5 address is created, the router will send hello packets to this address and establish OSPF adjacencies with other routers on the segment3. As explained above, once an OSPF adjacency is formed, the router will advertise a connected network to OSPF neighbors through DBD packets.

NEW QUESTION 8
Which two mechanisms are part of building and maintaining a Layer 2 bridge table? (Choose two.)

A. blocking
B. flooding
C. learning
D. listening

Answer: BC

Explanation:
? Option B is correct. Flooding is a mechanism used in Layer 2 bridging where the switch sends incoming packets to all its ports except for the port where the packet originated1. This is done when the switch doesn??t know the destination MAC address or when the packet is a broadcast or multicast1.
? Option C is correct. Learning is another mechanism used in Layer 2 bridging where the switch learns the source MAC addresses of incoming packets and associates them with the port on which they were received23. This information is stored in a MAC address table, also known as a bridge table23.
? Option A is incorrect. Blocking is a state in Spanning Tree Protocol (STP) used to prevent loops in a network2. It??s not a mechanism used in building and maintaining a Layer 2 bridge table2.
? Option D is incorrect. Listening is also a state in Spanning Tree Protocol (STP) where the switch listens for BPDUs to make sure no loops occur in the network before transitioning to the learning state2. It??s not a mechanism used in building and maintaining a Layer 2 bridge table2.

NEW QUESTION 9
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: CD

Explanation:
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 commands12.
? 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) instance1. 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 traffic1.
? 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_junos1.
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 traffic1.

NEW QUESTION 10
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: A

Explanation:
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 only support encapsulating IP traffic, which means that the payload of the inner packet must be an IP packet. IP-IP tunnels cannot encapsulate non-IP traffic, such as Ethernet frames or MPLS labels1.
Option A is correct, because IP-IP tunnels only support encapsulating IP traffic. Option B is incorrect, because IP-IP tunnels only support encapsulating non-IP traffic. Option C is incorrect, because the TTL in the inner packet is not decremented during transit to the tunnel endpoint. The TTL in the outer packet is decremented by each router along the path, but the TTL in the inner packet is preserved until it reaches the tunnel endpoint2. Option D is incorrect, because there are 20 bytes of overhead with IP-IP encapsulation. The overhead consists of the header of the outer packet, which has a fixed size of 20 bytes for IPv43.
References:
1: IP-IP Tunneling 2: What is tunneling? | Tunneling in networking 3: IPv4 - Header

NEW QUESTION 10
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: C

Explanation:
? The exhibit shows the output of the command show vlans brief, which displays brief information about VLANs and their associated interfaces1.
? 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 traffic1. This can be verified by the command show interfaces terse, which displays the status of the interfaces2.

NEW QUESTION 11
Exhibit.

Exhibit

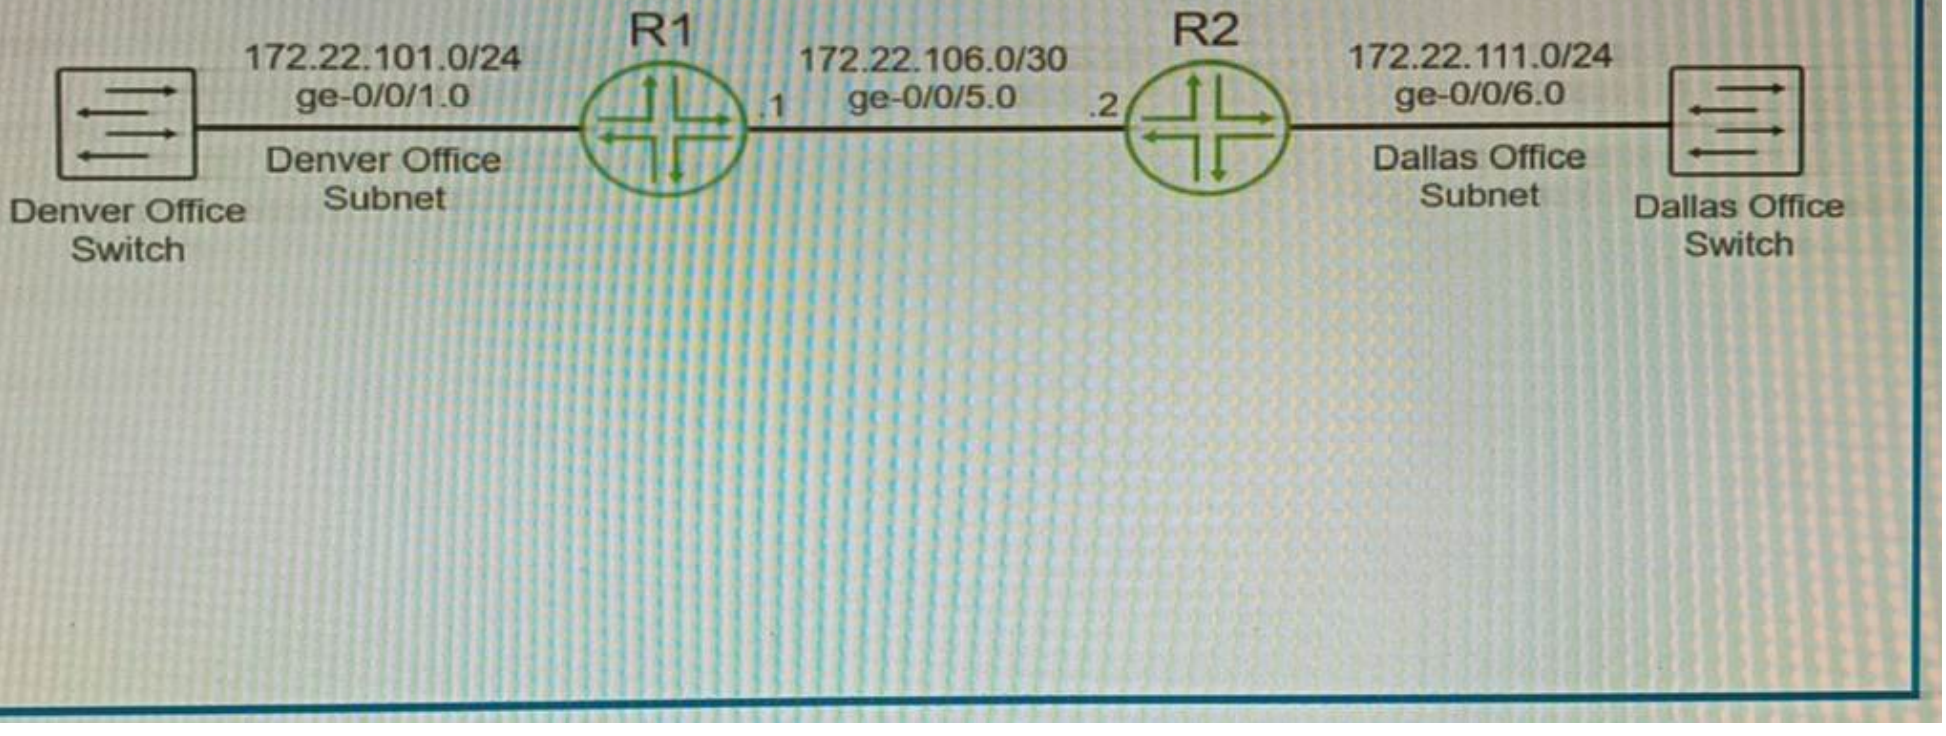

You are using OSPF to advertise the subnets that are used by the Denver and Dallas offices. The routers that are directly connected to the Dallas and Denver subnets are not advertising the connected subnets.
Referring to the exhibit, which two statements are correct? (Choose two.)

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: CD

Explanation:
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 OSPF1.
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 OSPF1.
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 packets1. 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 address1. Therefore, enabling passive interface can help in advertising the Dallas and Denver subnets.

NEW QUESTION 13
Which three protocols support BFD? (Choose three.)

A. RSTP
B. BGP
C. OSPF
D. LACP
E. FTP

Answer: BCD

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 17
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: AD

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 21
Exhibit

Exhibit

```
user@host# show
  protocols {
        oam {
              gre-tunnel {
                    interface gr-1/1/10.1 {
                          keepalive-time 10;
                                hold-time 10;
                    }
              }
        }
        lldp {
              interface all;
        }
  }
```

You have configured a GRE tunnel. To reduce the risk of dropping traffic, you have configured a keepalive OAM probe to monitor the state of the tunnel; however, traffic drops are still occurring.
Referring to the exhibit, what is the problem?

A. For GRE tunnels, the OAM protocol requires that the BFD protocols also be used.
B. The "event link-adjacency-loss" option must be set.
C. LLDP needs to be removed from the gr-1/1/10.1 interface.
D. The hold-time value must be two times the keepalive-time value

Answer: D

Explanation:
A keepalive OAM probe is a mechanism that can be used to monitor the state of a GRE tunnel and detect any failures in the tunnel path. A keepalive OAM probe consists of sending periodic packets from one end of the tunnel to the other and expecting a reply. If no reply is received within a specified time, the tunnel is considered down and the line protocol of the tunnel interface is changed to down1.
To configure a keepalive OAM probe for a GRE tunnel, you need to specify two parameters: the keepalive-time and the hold-time. The keepalive-time is the interval between each keepalive packet sent by the local router. The hold-time is the maximum time that the local router waits for a reply from the remote router before declaring the tunnel down2.
According to the Juniper Networks documentation, the hold-time value must be two times the keepalive-time value for a GRE tunnel2. This is because the hold-time value must account for both the round-trip time of the keepalive packet and the processing time of the remote router. If the hold-time value is too small, it may cause false positives and unnecessary tunnel flaps.
In the exhibit, the configuration shows that the keepalive-time is set to 10 seconds and the hold-time is set to 15 seconds for the gr-1/1/10.1 interface. This means that the local router will send a keepalive packet every 10 seconds and will wait for 15 seconds for a reply from the remote router. However, this hold-time value is not two times the keepalive-time value, which violates the recommended configuration. This may cause traffic drops if the remote router takes longer than 15 seconds to reply.
Therefore, option D is correct, because the hold-time value must be two times the keepalive-time value for a GRE tunnel. Option A is incorrect, because BFD is not required for GRE tunnels; BFD is another protocol that can be used to monitor tunnels, but it is not compatible with GRE keepalives3. Option B is incorrect, because the ??event link- adjacency-loss?? option is not related to GRE tunnels; it is an option that can be used to trigger an action when a link goes down4. Option C is incorrect, because LLDP does not need to be removed from the gr-1/1/10.1 interface; LLDP is a protocol that can be used to discover neighboring devices and their capabilities, but it does not interfere with GRE tunnels5.
References:
1: Configuring Keepalive Time and Hold time for a GRE Tunnel Interface 2: keepalive | Junos OS | Juniper Networks 3: Configuring Bidirectional Forwarding Detection 4: event link-adjacency-loss | Junos OS | Juniper Networks 5: Understanding Link Layer Discovery Protocol

NEW QUESTION 23
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: CD

Explanation:
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 device1. 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-switching2. 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 term3. 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 networks4.
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 protocols5.
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 VLAN6.
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

NEW QUESTION 28
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your JN0-351 Exam with Our Prep Materials Via below:

https://www.certleader.com/JN0-351-dumps.html