Juniper

Exam Questions JN0-351

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

NEW QUESTION 1
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 2
Exhibit

Exhibit

```
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;
}
```

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

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 3
Exhibit.

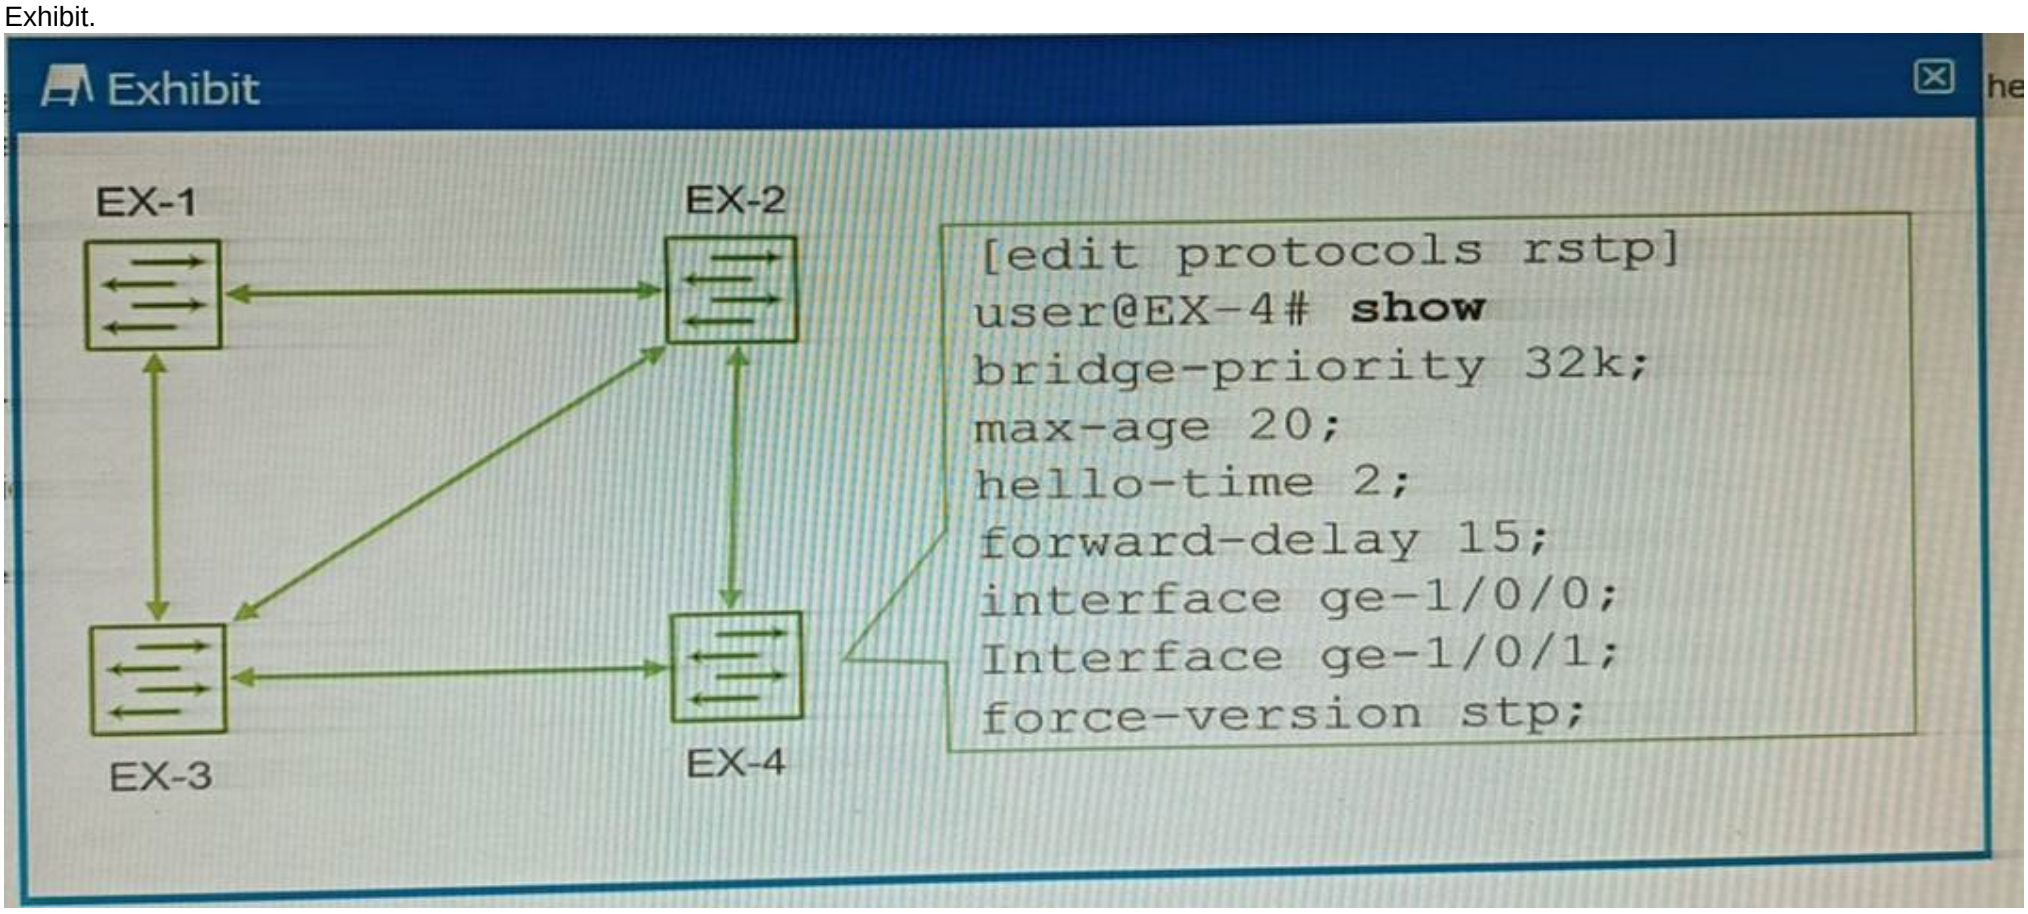

You have configured the four EX Series switches with RSTP, as shown in the exhibit. You discover that whenever a link between switches goes up or down, the switches take longer
than expected for RSTP to converge, using the default settings.
In this scenario, which action would solve the delay in RSTP convergence?

A. The hello-time must be increased.
B. The force-version must be removed.
C. The bridge priority for EX-4 must be set at 4000.
D. The max-age must be increased to 20

Answer: B

Explanation:
? The exhibit shows the configuration of RSTP on EX-4, which has the command force-version stp. This command forces the switch to use the legacy STP protocol instead of RSTP, even though the switch supports RSTP1. This means that EX-4 will not be able to take advantage of the faster convergence and enhanced features of RSTP, such as edge ports, link type, and proposal/agreement sequence2.
? The other switches in the network are likely to be running RSTP, as it is the default protocol for EX Series switches3. Therefore, there will be a compatibility issue between EX-4 and the other switches, which will result in longer convergence times and suboptimal performance. The switch will also generate a warning message that says ??Warning: STP version mismatch with neighbor?? when it receives a BPDU from a RSTP neighbor1.
? To solve this problem, the force-version command must be removed from EX-4, so that it can run RSTP natively and interoperate with the other switches in the network. This will enable faster convergence and better stability for the network topology. To remove the command, you can use the delete protocols rstp force-version command in configuration mode1.

NEW QUESTION 4
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: AB

Explanation:
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 endpoints34. 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 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
After receiving a BGP route, which two conditions are verified by the receiving router to ensure that the received route is valid? (Choose two)

A. The AS-path length is greater than 0.
B. The loops do not exist.
C. The next hop is reachable.
D. The local preference is greater than 0.

Answer: BC

Explanation:
? B is correct because the loops do not exist is one of the conditions that are verified by the receiving router to ensure that the received BGP route is valid. A loop in BGP means that a route has been advertised by the same AS more than once, which can cause routing instability and inefficiency1. To prevent loops, BGP uses the AS-path attribute, which lists the AS numbers that a route has traversed from the origin to the destination2. The receiving router checks the AS-path attribute of the received route and discards it if it finds its own AS number in the list2. This way, BGP avoids accepting routes that contain loops.
? C is correct because the next hop is reachable is one of the conditions that are verified by the receiving router to ensure that the received BGP route is valid. The next hop is the IP address of the next router that is used to forward packets to the destination network3. The receiving router checks the next hop attribute of the received route and verifies that it has a valid route to reach it3. If the next hop is not reachable, the received route is not usable and is rejected by the receiving router3. This way, BGP ensures that only feasible routes are accepted.

NEW QUESTION 7
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 8
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: AB

Explanation:
? 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 9
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 10
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: AD

Explanation:
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 adjacencies1. If the MTU is less than this, IS-IS adjacencies will not be established1.
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-IS1. If the loopback interface is not included in IS-IS, it could prevent IS-IS adjacencies from being established1.
Therefore, options A and D are correct.

NEW QUESTION 10
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: AC

Explanation:
? 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 port1. However, an alternate port does not forward any traffic, as it is in a discarding state2. 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 port1.
? 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 ID3. However, an alternate port receives a higher cost BPDU than the root port, otherwise it would be the root port itself1. Therefore, an alternate port provides an alternate higher cost path to the root bridge than the root port.

NEW QUESTION 14
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: AC

Explanation:
? 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 18
Which statement is correct about graceful Routing Engine switchover (GRES)?

A. The PFE restarts and the kernel and interface information is lost.
B. GRES has a helper mode and a restarting mode.
C. When combined with NSR, routing is preserved and the new master RE does not restart rpd.

D. With no other high availability features enabled, routing is preserved and the new master RE does not restart rpd.

Answer: C

Explanation:
The Graceful Routing Engine Switchover (GRES) feature in Junos OS enables a router with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails1. GRES preserves interface and kernel information, ensuring that traffic is not interrupted1. However, GRES does not preserve the control plane1.
To preserve routing during a switchover, GRES must be combined with either Graceful
Restart protocol extensions or Nonstop Active Routing (NSR)1. When GRES is combined with NSR, nearly 75 percent of line rate worth of traffic per Packet Forwarding Engine remains uninterrupted during GRES1. Any updates to the primary Routing Engine are replicated to the backup Routing Engine as soon as they occur1.
Therefore, when GRES is combined with NSR, routing is preserved and the new master RE does not restart rpd1.

NEW QUESTION 23
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: BC

Explanation:
? 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 27
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 31
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

JN0-351 Practice Exam Features:

* JN0-351 Questions and Answers Updated Frequently

* JN0-351 Practice Questions Verified by Expert Senior Certified Staff

* JN0-351 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* JN0-351 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year