



Juniper

Exam Questions JN0-351

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

You are attempting to configure the initial two aggregated Ethernet interfaces on a router but there are no aggregated Ethernet interfaces available. In this scenario, which configuration will enable these interfaces on this router?

A)

```
user@router# show chassis
aggregated-devices {
    ethernet {
        lacp {
            system-priority 10;
        }
    }
}
```

B)

```
user@router# show chassis
aggregated-devices {
    ethernet {
        device-count 10;
    }
}
```

C)

```
user@router# show chassis
maximum-ecmp 16;
aggregated-devices {
    ethernet {
        device-count 1;
    }
}
```

D)

```

user@router# show chassis
aggregated-devices {
  ethernet {
    device-count 1;
  }
}

```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

The correct answer to your question is C. Option C. Here is why:

? Option C shows the configuration of the chassis statement, which defines the properties of the router chassis, such as the number of aggregated Ethernet interfaces, the number of FPCs, and the number of PICs1.

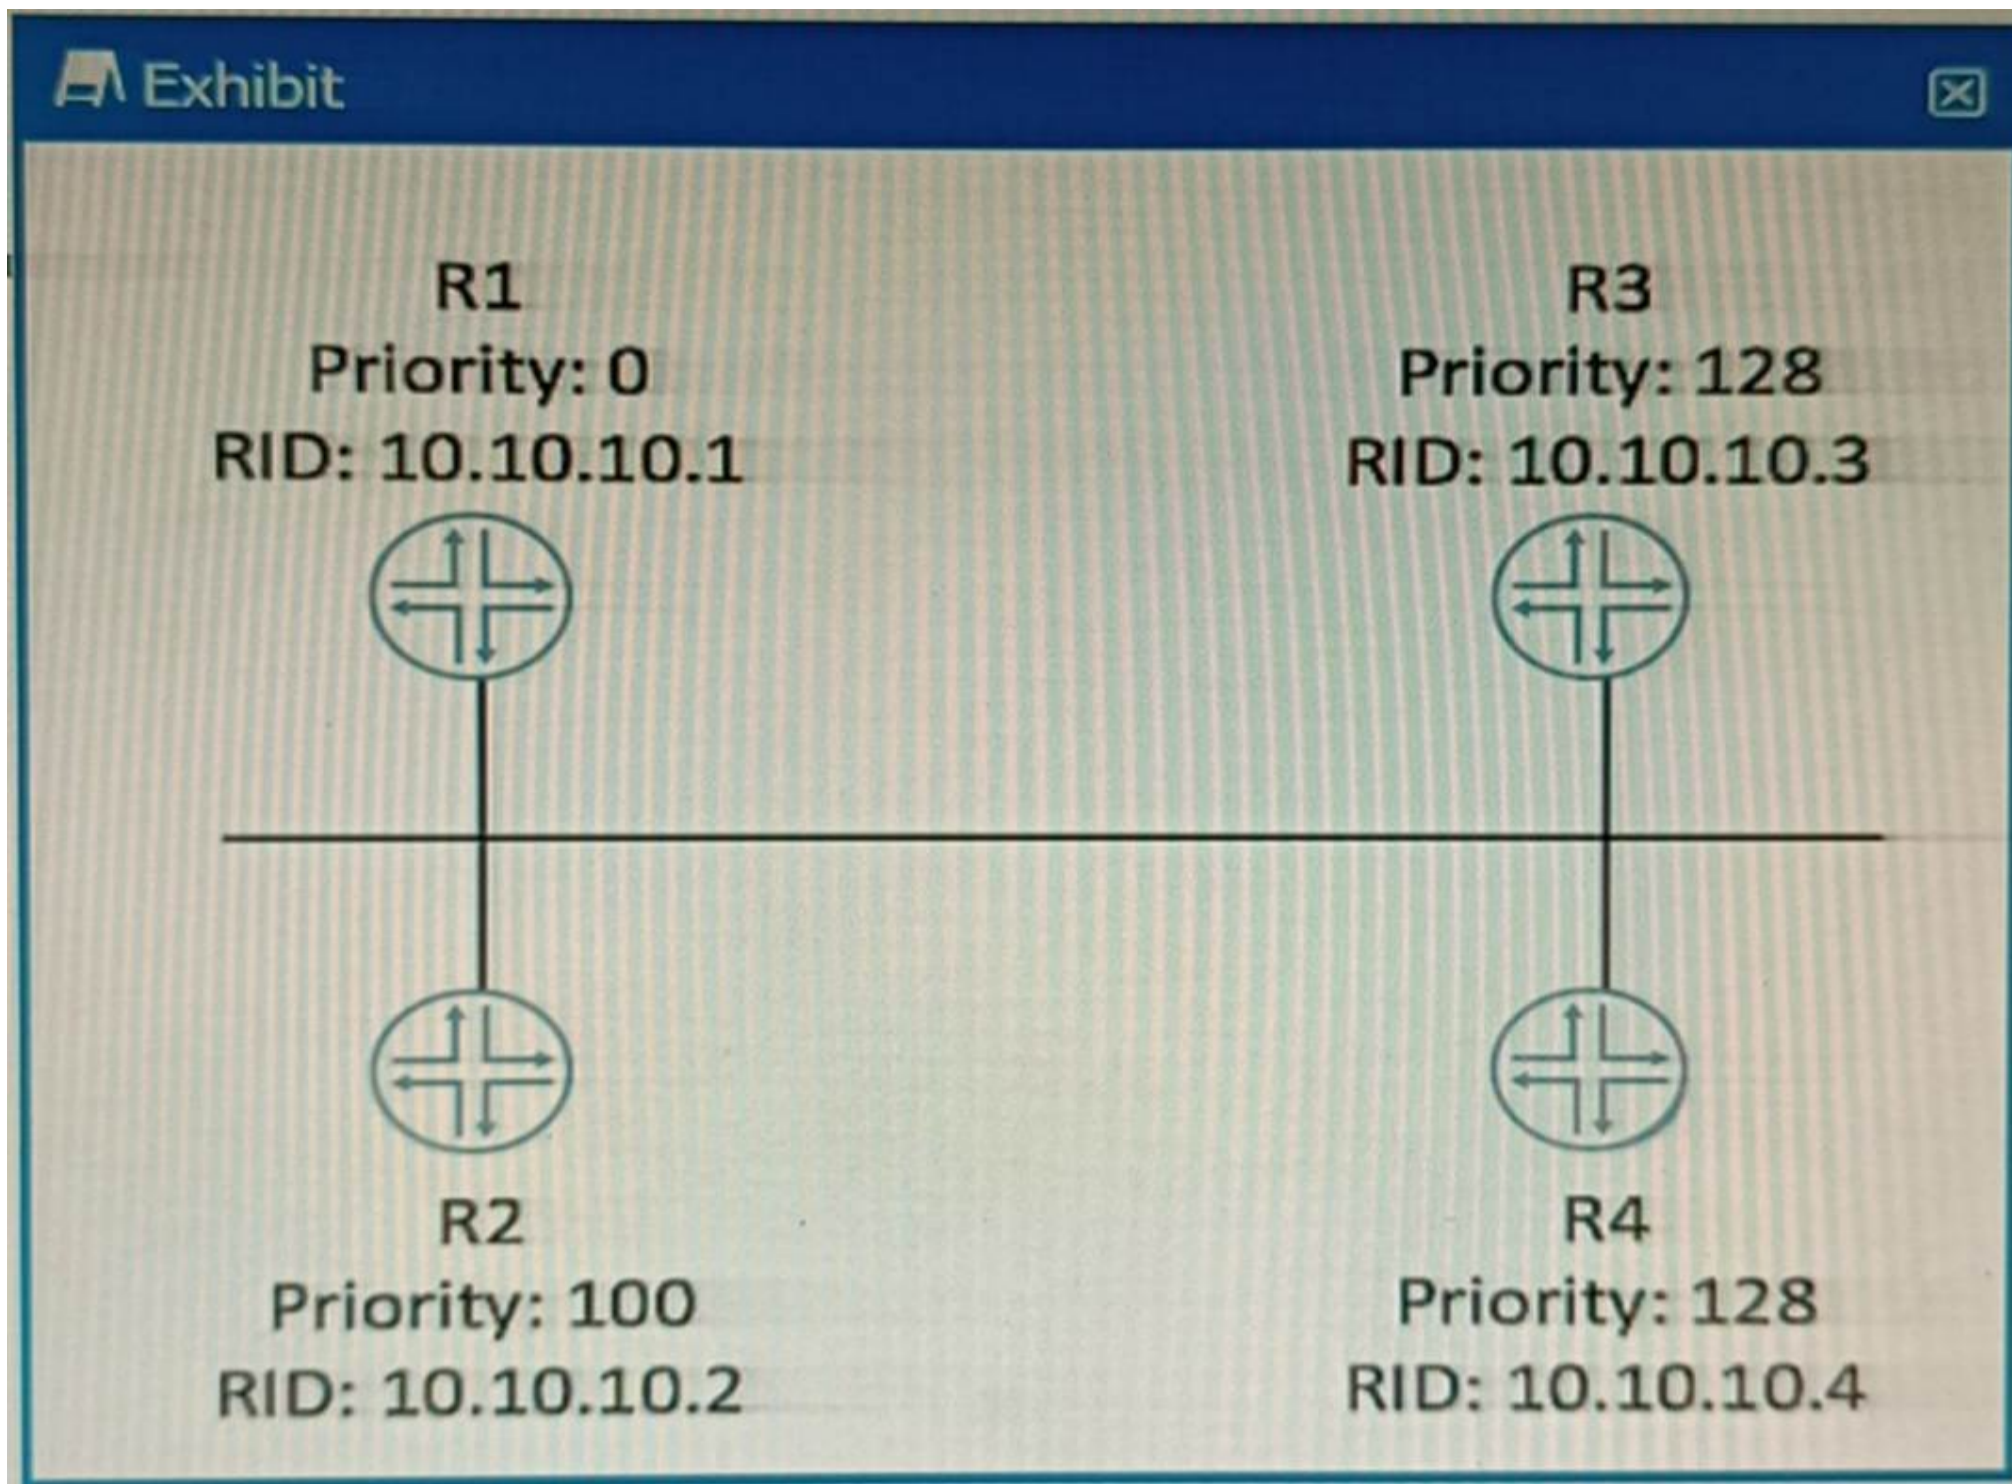
? To enable aggregated Ethernet interfaces on a router, you need to specify the aggregated-devices statement under the chassis statement and set the ethernet parameter to the desired number of interfaces2. For example, to enable two aggregated Ethernet interfaces, you can use the following configuration:
 chassis { aggregated-devices { ethernet { device-count 2; } } }

? Option C shows this configuration with the device-count set to 2, which will enable two aggregated Ethernet interfaces on the router. The other options do not show this configuration and will not enable any aggregated Ethernet interfaces on the router.

? Therefore, option C is the correct answer to your question.

NEW QUESTION 2

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 Tuning
- 2: OSPF Designated Router (DR) and Backup Designated Router (BDR)

NEW QUESTION 3

Exhibit.

Exhibit
✕

```

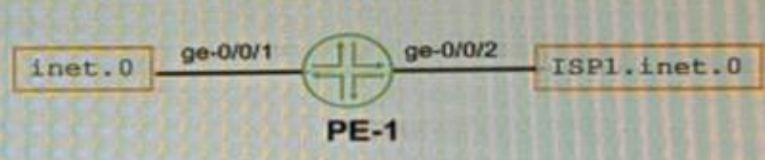
user@PE-1> show route table ISPI.inet.0

user@PE-1> configure

[edit]
user@PE-1# show routing-instances
ISPI {
  instance-type forwarding;
  routing-options {
    static {
      route 0.0.0.0/0 next-hop 203.0.113.2;
    }
    instance-import ISPI-import;
  }
}

[edit]
user@PE-1# show policy-options
policy-statement ISPI-import {
  from instance master;
  then accept;
}

```



The ispi _ inet. 0 route table has currently no routes in it.
 What will happen when you commit the configuration shown on the exhibit?

- A. The ine
- B. 0 route table will be completely overwritten by the ispi . ine
- C. 0 route table.
- D. The ine
- E. 0 route table will be imported into the ispi . ine
- F. 0 route table.
- G. The ISPI . ine
- H. 0 route table will be completely overwritten by the ine
- I. o route table.
- J. The ISPI . ine
- K. 0 route table will be imported into the ine
- L. 0 route table.

Answer: B

Explanation:

The configuration shown in the exhibit is an example of a routing instance of type virtual-router. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters that create a separate routing domain on a Juniper device¹. A virtual-router routing instance allows administrators to divide a device into multiple independent virtual routers, each with its own routing table².

The configuration also includes a rib-group statement, which is used to import routes from one routing table to another. A rib-group consists of an import-rib statement, which specifies the source routing table, and an export-rib statement, which specifies the destination routing table.

In this case, the rib-group name is inet-to-ispi, and the import-rib statement specifies inet.0 as the source routing table. The export-rib statement specifies ispi.inet.0 as the destination routing table. This means that the routes from inet.0 will be imported into ispi.inet.0. Therefore, the correct answer is B. The inet.0 route table will be imported into the ispi.inet.0 route table.

References:

- 1: Routing Instances Overview 2: Virtual Routing Instances : [rib-group (Routing Options)]

NEW QUESTION 4

Which two statements about redundant trunk groups on EX Series switches are correct? (Choose two.)

- A. Redundant trunk groups use spanning tree to provide loop-free redundant uplinks.
- B. Redundant trunk groups load balance traffic across two designated uplink interfaces.
- C. Layer 2 control traffic is permitted on the secondary link.
- D. If the active link fails, then the secondary link automatically takes over.

Answer: CD

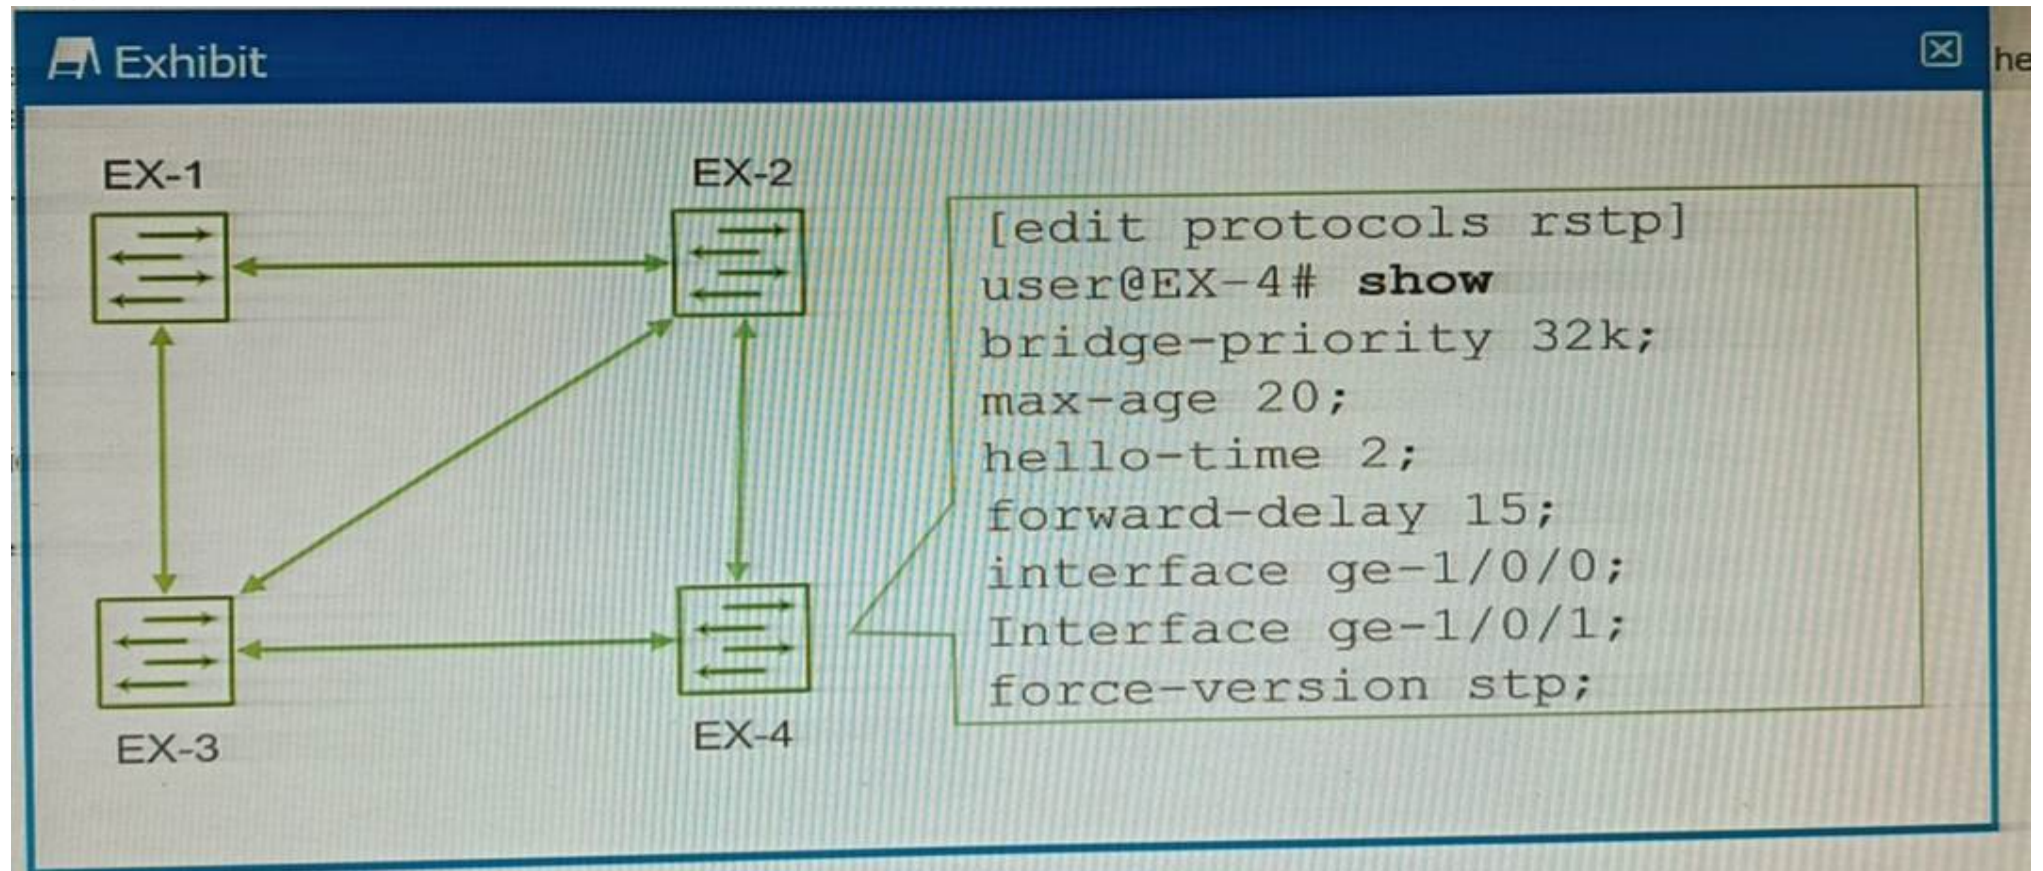
Explanation:

? C is correct because Layer 2 control traffic is permitted on the secondary link of a redundant trunk group (RTG) on EX Series switches. Layer 2 control traffic includes protocols such as LLDP, LACP, and STP, which are used to exchange information and coordinate actions between switches¹. According to the Juniper Networks documentation², Layer 2 control traffic is allowed to pass through both the active and the secondary links of an RTG, but data traffic is only forwarded through the active link. This allows the switches to maintain their Layer 2 adjacencies and monitor the link status on both links.

? D is correct because if the active link fails, then the secondary link automatically takes over in an RTG on EX Series switches. An RTG consists of two trunk links: an active or primary link, and a secondary or backup link². The active link is used to forward data traffic, while the secondary link is in standby mode. If the active link fails or becomes unavailable, the secondary link immediately transitions to a forwarding state and takes over the data traffic without waiting for normal STP convergence². This provides fast recovery and redundancy for the network.

NEW QUESTION 5

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 6

Two routers share the same highest priority and start time.

- A. In this situation, what is evaluated next when determining the designated router? The router with the lowest router ID become the DR.
- B. The router with the highest router ID becomes the DR
- C. The routers perform another DR election.
- D. The router with the highest MAC address become the DR

Answer: B

Explanation:

? According to the OSPF protocol, the designated router (DR) is the router that acts as the focal point for exchanging routing information on a multi-access network segment, such as a LAN1. The DR election process is based on the following criteria, in order of precedence1:

? In your scenario, two routers share the same highest priority and start time. This means that they have equal chances of becoming the DR based on the first and third criteria. Therefore, the second criterion will be used to break the tie, which is the router ID. The router with the highest router ID will become the DR, and the other router will become the backup designated router (BDR), which is ready to take over the role of DR if it fails1.

NEW QUESTION 7

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: BD

Explanation:

Junos devices support various types of tunnels for different purposes12.

? 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 network1. Junos devices support GRE tunnels1.

- ? 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 session¹. Junos devices support IPsec tunnels¹.
- ? 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 network².
- ? Option C is incorrect. While Junos devices do support IP-IP (also known as IP tunneling), it's not supported on all Junos devices¹.

NEW QUESTION 8

You are an operator for a network running IS-IS. Two routers are failing to form an adjacency. What are two reasons for this problem? (Choose two.)

- A. There are mismatched router IDs on the L2 routers.
- B. There is no configured ISO address on any IS-IS interface.
- C. There is a mismatched area ID between the L2 routers.
- D. The family iso configuration is missing from the adjacency interface.

Answer: BD

Explanation:

The two reasons for the failure to form an adjacency in a network running IS-IS could be:

* B. There is no configured ISO address on any IS-IS interface. IS-IS requires each router interface to have an ISO address configured. Without this address, the routers cannot form an adjacency¹.

* D. The family iso configuration is missing from the adjacency interface. The family iso configuration is essential for IS-IS to function correctly. If this configuration is missing from the adjacency interface, it could prevent the formation of an adjacency¹.

These explanations are based on the Enterprise Routing and Switching Specialist (JNCIS-ENT) documents and learning resources available at Juniper Networks^{2,3}.

NEW QUESTION 9

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 10

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: A

Explanation:

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 generated¹. If the switch has been configured with the port-error-disable statement, the disabled interface recovers automatically upon expiration of the specified disable timeout¹. 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 again¹. This explanation is based on the Enterprise Routing and Switching Specialist (JNCIS-ENT) documents and learning resources available at Juniper Networks¹.

NEW QUESTION 10

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: AC

Explanation:

? 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 snooping¹. DAI discards any ARP packets that do not match the database or have invalid formats¹.

? 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 ports². DHCP snooping also builds a database of MAC addresses, IP addresses, lease times, and binding types for each client².

NEW QUESTION 13

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 17

Which two statements about BGP facilitate the prevention of routing loops between two autonomous systems? (Choose two.)

- A. EBGp routers will append their AS number when advertising routes to their neighbors.
- B. EBGp routers will only accept routes that contain their own AS number in the AS_PATH.
- C. EBGp routers will drop routes that contain their own AS number in the AS_PATH
- D. EBGp routers will prepend their AS number when advertising routes to their neighbors

Answer: AC

Explanation:

BGP (Border Gateway Protocol) is a protocol designed to exchange routing and reachability information among autonomous systems (AS) on the internet¹.

? Option A is correct. When an EBGp router advertises routes to its neighbors, it appends its AS number to the AS_PATH attribute¹. This is a key mechanism in BGP to prevent routing loops¹.

? Option C is correct. BGP has a built-in loop prevention mechanism whereby if a BGP router detects its own AS in the AS_PATH attribute, it will drop the prefix and will not continue to advertise it². This helps to prevent routing loops².

? Option B is incorrect. EBGp routers do not accept routes that contain their own AS number in the AS_PATH². Instead, they drop such routes as part of the loop prevention mechanism².

? Option D is incorrect. While it's true that EBGp routers append their AS number when advertising routes, they do not prepend their AS number¹. The term "prepend" in BGP usually refers to a technique used to influence path selection by artificially lengthening the AS_PATH³.

NEW QUESTION 21

.....

Relate Links

100% Pass Your JN0-351 Exam with ExamBible Prep Materials

<https://www.exambible.com/JN0-351-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>