

Exam Questions JN0-637

Security - Professional (JNCIP-SEC)

<https://www.2passeasy.com/dumps/JN0-637/>



NEW QUESTION 1

Exhibit:

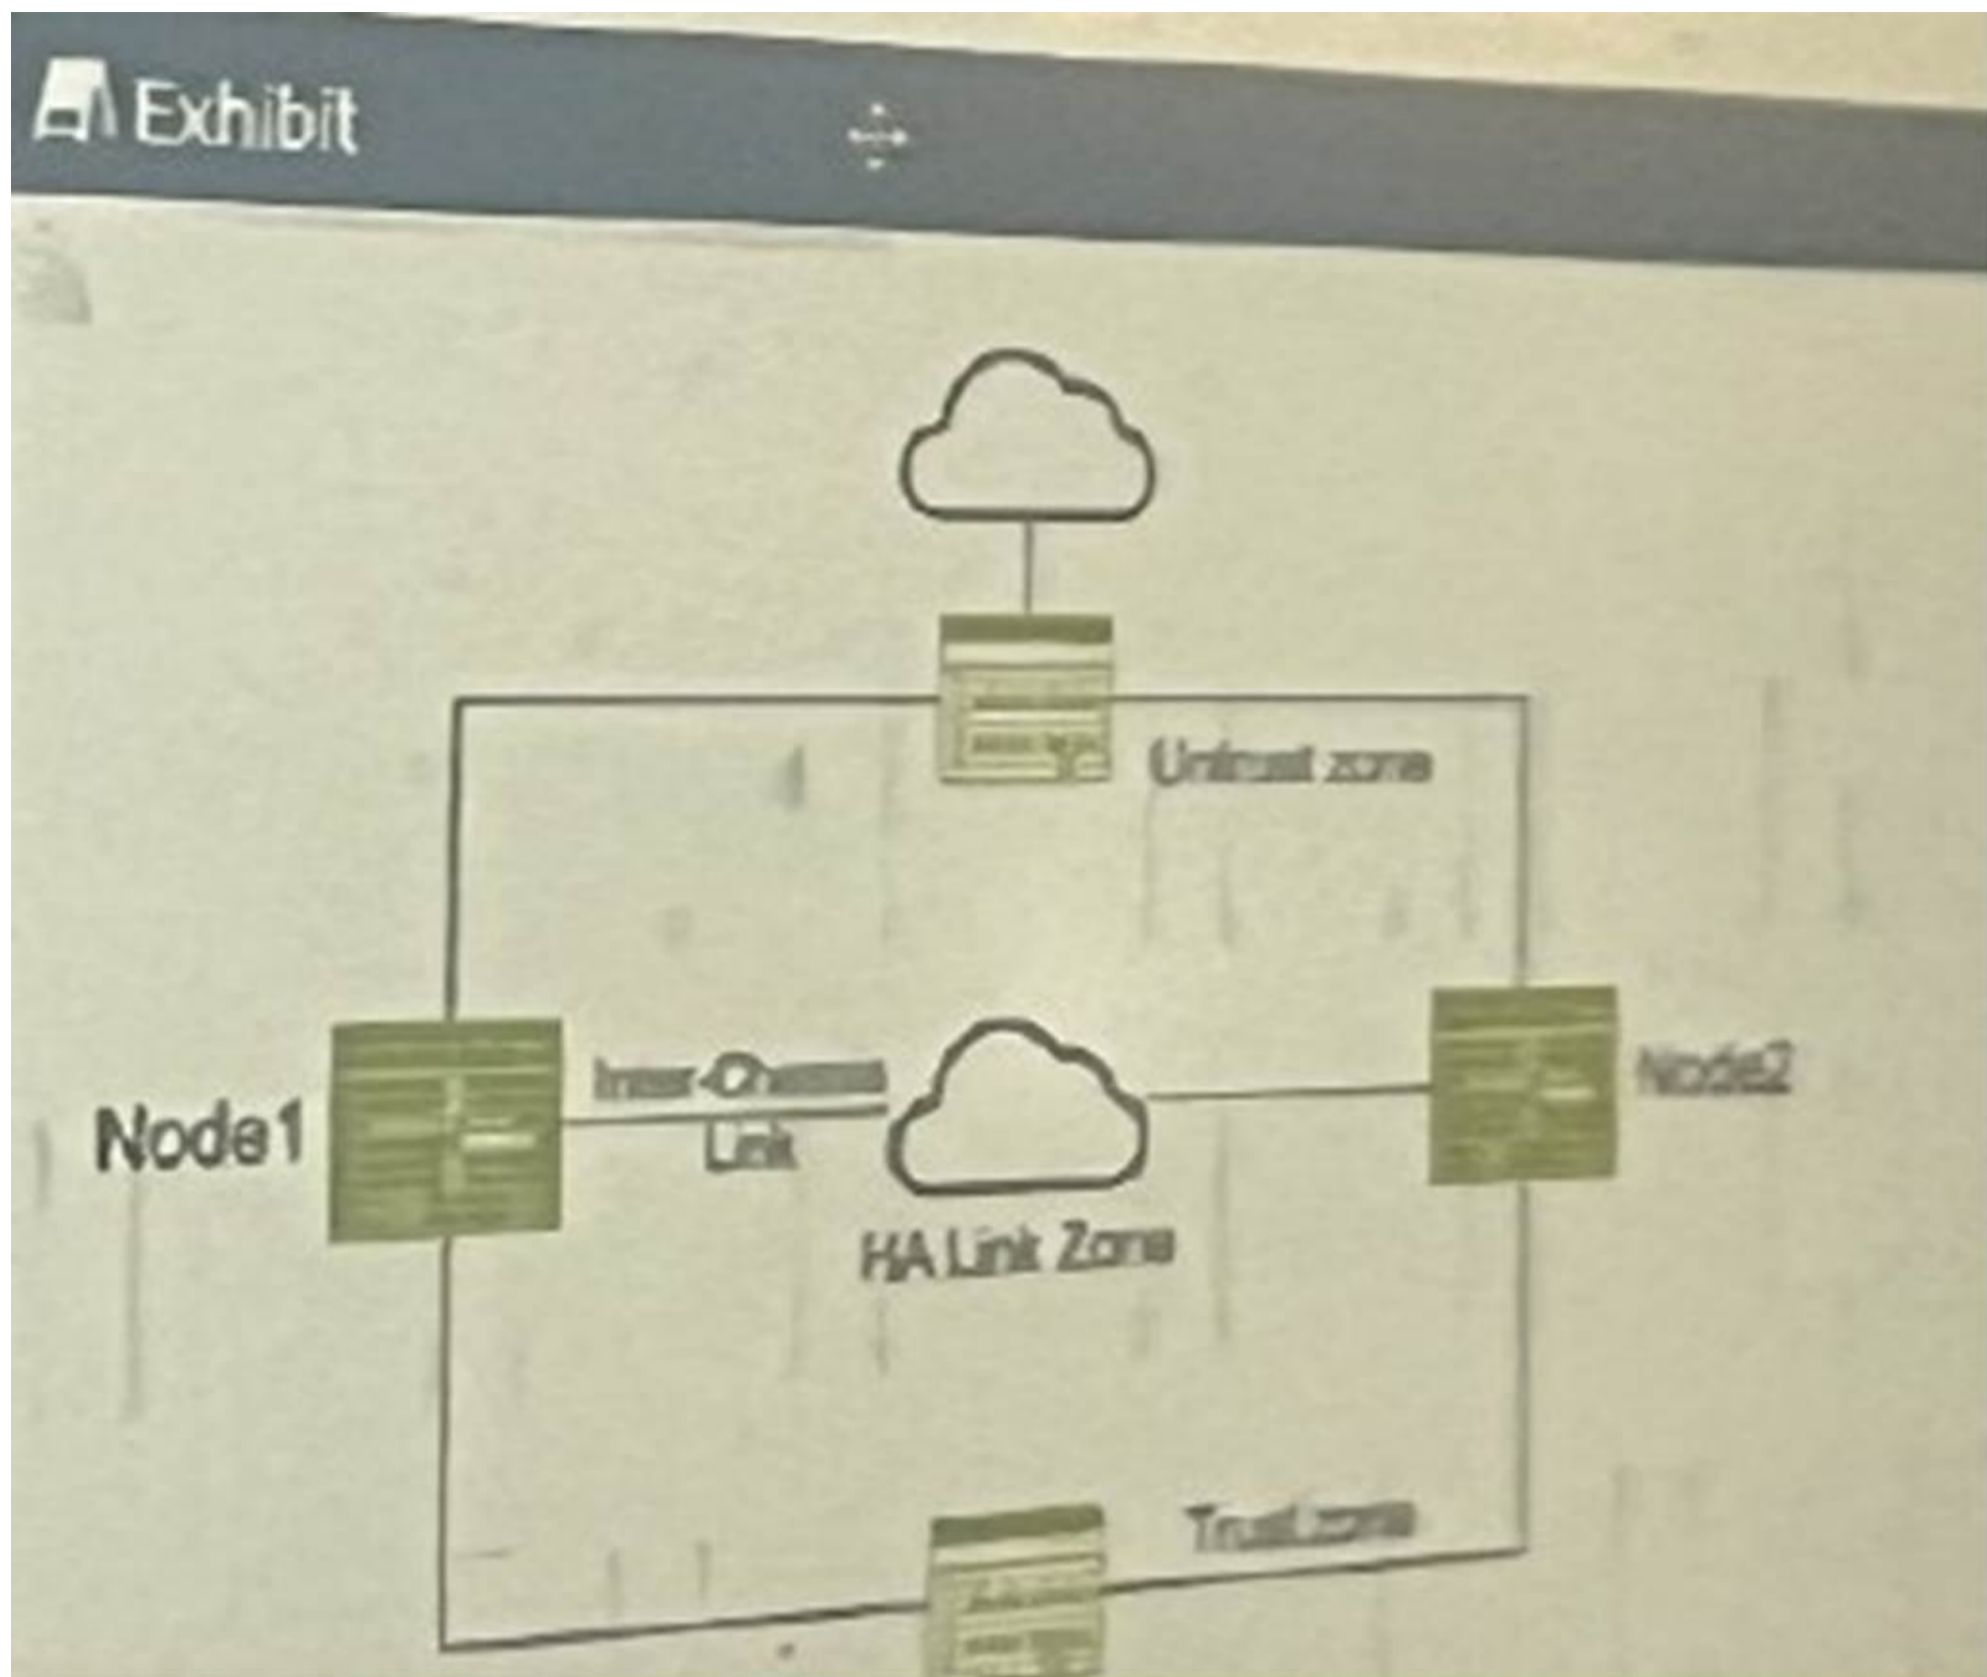
```
[edit]
user@srx# show security nat
source {
    pool ipv4-source-pool {
        address {
            10.10.101.10/32;
        }
    }
}
rule-set ipv6-source {
    from zone trust;
    to zone untrust;
    rule ipv6-host-source {
        match {
            source-address 2001:db8::1/128;
            destination-address 10.10.201.10/32;
        }
        then {
            source-nat {
                pool {
                    ipv4-source-pool;
                }
            }
        }
    }
}
```

You are configuring NAT64 on your SRX Series device. You have committed the configuration shown in the exhibit. Unfortunately, the communication with the 10.10.201.10 server is not working. You have verified that the interfaces, security zones, and security policies are all correctly configured. In this scenario, which action will solve this issue?

- A. Configure source NAT to translate return traffic from IPv4 address to the IPv6 address of your source device.
- B. Configure proxy-ARP on the external IPv4 interface for the 10.10.201.10/32 address.
- C. Configure proxy-NDP on the IPv6 interface for the 2001:db8::1/128 address.
- D. Configure destination NAT to translate return traffic from the IPv4 address to the IPv6 address of your source device.

Answer: D**NEW QUESTION 2**

Exhibit:



You have deployed a pair of SRX series devices in a multimode HA environment. You need to enable IPsec encryption on the interchassis link. Referring to the exhibit, which three steps are required to enable ICL encryption? (Choose three.)

- A. Install the Junos IKE package on both nodes.
- B. Enable OSPF for both interchassis link interfaces and turn on the dynamic-neighbors parameter.
- C. Configure a VPN profile for the HA traffic and apply to both nodes.
- D. Enable HA link encryption in the IPsec profile on both nodes.
- E. Enable HA link encryption in the IKE profile on both nodes.

Answer: ACD

Explanation:

? A. Install the Junos IKE package on both nodes. While I previously stated that IKE is usually included in the base Junos OS image, it's essential to ensure that the necessary IKE package is indeed installed and enabled on both SRX nodes to support ICL encryption.

? C. Configure a VPN profile for the HA traffic and apply it to both nodes. This dedicated VPN profile defines the security parameters (encryption algorithms, authentication, etc.) specifically for the ICL traffic.

? D. Enable HA link encryption in the IPsec profile on both nodes. Within the IPsec profile, you must explicitly enable ICL encryption to ensure that all traffic traversing the interchassis link is protected.

Why E is incorrect:

? E. Enable HA link encryption in the IKE profile on both nodes. While securing IKE negotiations is important, it's typically handled within the IPsec profile itself when configuring ICL encryption on SRX devices.

NEW QUESTION 3

You are enabling advanced policy-based routing. You have configured a static route that has a next hop from the inet.0 routing table. Unfortunately, this static route is not active in your routing instance. In this scenario, which solution is needed to use this next hop?

- A. Use RIB groups.
- B. Use filter-based forwarding.
- C. Use transparent mode.
- D. Use policies.

Answer: A

Explanation:

To enable advanced policy-based routing in Junos OS and activate a static route with a next-hop address in the inet.0 table within your routing instance, you should utilize RIB groups. RIB groups allow you to import routes from one routing table to another. In this scenario, the static route within the routing instance needs access to the inet.0 routes, which is facilitated by configuring a RIB group. Juniper's documentation outlines RIB groups as a necessary component for handling instances where routes need to be shared across routing tables, thereby ensuring seamless traffic flow through specified routes. For more details, refer to the Juniper Networks Documentation on RIB Groups.

In Junos OS for SRX Series devices, when enabling advanced policy-based routing and configuring a static route with a next-hop from the inet.0 routing table, the issue arises because the static route is not being used in the routing instance. This is a common scenario when the next-hop belongs to a different routing table or instance, and the routing instance is not aware of that next-hop.

To resolve this, RIB (Routing Information Base) groups are used. RIB groups allow routes from one routing table (RIB) to be shared or imported into another routing table. This means that the routing instance can import the necessary routes from inet.0 and make them available for the routing instance where the policy-based routing is applied.

Detailed Steps:

? Configure the Static Route: First, configure the static route pointing to the next-hop in inet.0. Here's an example:

```
bash
set routing-options static route 10.1.1.0/24 next-hop 192.168.1.1 This static route will be placed in the inet.0 routing table by default.
```

? Create and Apply a RIB Group: To import routes from inet.0 into the routing instance, create a RIB group configuration. This will allow the static route from inet.0 to be visible within the routing instance.

```
Example configuration for the RIB group: bash
set routing-options rib-groups RIB-GROUP import-rib inet.0
set routing-options rib-groups RIB-GROUP import-rib <routing-instance-name>.inet.0
```

This configuration ensures that routes from inet.0 are imported into the specified routing instance.

? Apply the RIB Group to the Routing Instance: Once the RIB group is configured, apply it to the appropriate routing instance:

```
bash
set routing-instances <routing-instance-name> routing-options rib-group RIB-GROUP
```

? Verify Configuration: Use the following command to verify that the static route has been imported into the routing instance:

```
bash
show route table <routing-instance-name>.inet.0
```

The output should now display the static route imported from inet.0.

Juniper Security Reference:

? RIB Groups Overview: Juniper's documentation provides detailed information on how RIB groups function and how to use them to share routes between different routing tables. This is essential for scenarios involving policy-based routing where routes from one instance (like inet.0) need to be available in another instance.

Reference: Juniper Networks Documentation on RIB Groups.

By using RIB groups, you ensure that the static route from inet.0 is available in the appropriate routing instance for policy-based routing to function correctly. This avoids the need for other methods like filter-based forwarding or transparent mode, which do not address the specific issue of static route visibility across routing instances.

=====

NEW QUESTION 4

You are experiencing problem with your ADVPN tunnels getting established. The tunnel and egress interface are located in different zone. What are two reasons for these problems? (Choose two.)

- A. IKE is not an allowed protocol in the external interfaces' security zone.
- B. IKE is not an allowed protocol in the tunnel endpoints' security zone.
- C. OSPF is not an allowed protocol in the tunnel endpoints' security zone.
- D. BGP is not an allowed protocol in the tunnel endpoints' security zone.

Answer: AB

NEW QUESTION 5

Which two elements are necessary to configure a rule under an APBR profile? (Choose Two)

- A. instance type
- B. match condition
- C. then action
- D. RIB group

Answer: BC

Explanation:

Here's why those elements are necessary for configuring a rule under an APBR profile:

? B. Match condition: This defines the criteria for matching traffic to the APBR rule. It can include:

? C. Then action: This specifies the action to take when traffic matches the rule. The primary action in APBR is:

Why other options are incorrect:

? A. Instance type: While routing instances are used in APBR, the "instance type" itself is not configured within the APBR rule. You define the instance type separately when configuring the routing instance.

? D. RIB group: RIB groups are used for route management and are not directly involved in APBR rule configuration.

NEW QUESTION 6

Which two statements about the differences between chassis cluster and multinode HA on SRX series devices are true? (Choose Two)

- A. Multinode HA member nodes require Layer 2 connectivity.
- B. Multinode HA supports Layer 2 and Layer 3 connectivity between nodes.
- C. Multinode HA requires Layer 3 connectivity between nodes.
- D. Chassis cluster member nodes require Layer 2 connectivity.

Answer: BD

NEW QUESTION 7

Which two statements are correct about DNS doctoring?

- A. The DNS ALG must be disabled.
- B. Proxy ARP is required if your NAT pool for the server is on the same subnet as the uplink interface.
- C. Proxy ARP is required if your NAT pool for the server is on a different subnet as the uplink interface
- D. The DNS ALG must be enabled.

Answer: BD

NEW QUESTION 8

Which two statements are correct about automated threat mitigation with Security Director? (Choose two.)

- A. It works with third-party switches.
- B. It provides endpoint protection by running a Juniper ATP Cloud agent on the servers.
- C. It provides endpoint protection by running a Juniper ATP Cloud agent on EX Series devices.
- D. It works with SRX Series devices.

Answer: AD

NEW QUESTION 9

You want to deploy two vSRX instances in different public cloud providers to provide redundant security services for your network. Layer 2 connectivity between the two vSRX instances is not possible.

What would you configure on the vSRX instances to accomplish this task?

- A. Chassis cluster
- B. Secure wire
- C. Multinode HA
- D. Virtual chassis

Answer: C

NEW QUESTION 10

In a multinode HA environment, which service must be configured to synchronize between nodes?

- A. Advanced policy-based routing
- B. PKI certificates
- C. IPsec VPN
- D. IDP

Answer: B

NEW QUESTION 10

You configured two SRX series devices in an active/passive multimode HA setup. In this scenario, which statement is correct?

- A. Both devices are in the passive state until the activeness determination process is completed.
- B. Both devices start in a hold state until the activeness determination process is completed.
- C. Both devices start in the undiscovered state until the activeness determination process is completed.
- D. Both devices are in the active state until the activeness determine determination process is completed.

Answer: D

NEW QUESTION 12

You want to bypass IDP for traffic destined to social media sites using APBR, but it is not working and IDP is dropping the session.

What are two reasons for this problem? (Choose two.)

- A. IDP disable is not configured on the APBR rule.
- B. The application services bypass is not configured on the APBR rule.
- C. The APBR rule does a match on the first packet.
- D. The session did not properly reclassify midstream to the correct APBR rule.

Answer: AD

Explanation:

Comprehensive Detailed Step-by-Step Explanation with All Juniper Security References

Understanding the Problem:

? The goal is to bypass IDP for traffic destined to social media sites using Application-Based Policy Routing (APBR).

? Despite the configuration, IDP is still dropping the sessions.

? Need to identify two reasons why this is happening.

Key Concepts:

? Application-Based Policy Routing (APBR): Allows routing decisions based on the application identified in the traffic.

? IDP (Intrusion Detection and Prevention): Monitors network traffic for malicious activity and can drop suspicious packets.

? Bypassing IDP: To bypass IDP for certain traffic, specific configurations are required within the APBR rule.

Option A: IDP disable is not configured on the APBR rule.

? Explanation:

Reference:

Juniper Networks Documentation:

"To bypass IDP processing for traffic matching an APBR rule, include the idp-disable statement in the rule configuration."

Source: Juniper TechLibrary - Configuring APBR to Bypass IDP

Option D: The session did not properly reclassify midstream to the correct APBR rule.

* Explanation:

Midstream Reclassification: APBR relies on application identification, which may occur after several packets have been exchanged (not just the first packet).

When the application is identified mid-session, the session should be reclassified according to the correct APBR rule.

If midstream reclassification does not occur properly, the session continues under the initial policy, and IDP continues to inspect and potentially drop the traffic.

Possible Causes:

Session Setup Issues: If the session was established before the application was identified, and reclassification is not enabled or not functioning, the session won't switch to the APBR rule that bypasses IDP.

Configuration Errors: Incorrect or missing configuration for midstream reclassification.

Reference:

Juniper Networks Documentation:

"For APBR to reclassify sessions after the application is identified, ensure that midstream reclassification is enabled."

Source: Juniper TechLibrary - Understanding APBR and Midstream Reclassification

Why Options B and C are Incorrect:

Option B: The application services bypass is not configured on the APBR rule.

* Explanation:

There is no specific application-services bypass option within APBR rules for bypassing IDP.

To bypass IDP, the idp-disable option must be used.

Application services bypass generally refers to bypassing other services like UTM, not specifically IDP within APBR.

Reference:

Juniper Networks Documentation:

"APBR rules can include the idp-disable statement to bypass IDP. There is no application- services bypass statement for APBR."

Option C: The APBR rule does a match on the first packet.

* Explanation:

By default, APBR can match on the first packet, but for applications that require deeper inspection, you can configure the rule to not match on the first packet.

Matching on the first packet is generally beneficial for routing decisions.

In this scenario, matching on the first packet is not the reason why IDP is dropping the session.

Reference:

Juniper Networks Documentation:

"If you configure APBR to match on the first packet, the routing decision is made immediately. If the application is not identified on the first packet, the default routing is used until the application is identified."

Conclusion: Correct Answers:

* A. IDP disable is not configured on the APBR rule.

Without idp-disable, IDP will continue to inspect and possibly drop the traffic matching the APBR rule.

* D. The session did not properly reclassify midstream to the correct APBR rule.

If midstream reclassification fails, the session remains under the initial policy, and IDP processing continues.

Resolution Steps:

Configure idp-disable: Ensure that the APBR rule includes the idp-disable statement to bypass IDP for the specified traffic.

arduino Copy code

```
set security application-path-routing rule <rule-name> then idp-disable
```

Enable Midstream Reclassification: Verify that midstream reclassification is enabled and functioning correctly to reclassify sessions once the application is identified.

Note: Midstream reclassification is enabled by default, but verify that no configuration is preventing it.

Additional References:

Juniper TechLibrary:

"Application-Based Policy Routing Overview" - Provides an overview of APBR features and configurations.

Source: Juniper TechLibrary - APBR Overview

"Configuring IDP Policy Bypass" - Discusses how to bypass IDP for specific traffic. Source: Juniper TechLibrary - Configuring IDP Bypass

Juniper Networks Day One Book:

"Advanced Security Policies" - Offers insights into configuring advanced security policies, including APBR and IDP interactions.

NEW QUESTION 14

A company has acquired a new branch office that has the same address space as one of its local networks, 192.168.100.0/24. The offices need to communicate with each other.

Which two NAT configurations will satisfy this requirement? (Choose two.)

A. [edit security nat source]user@OfficeA# show rule-set OfficeBtoA { from zone OfficeB;to zone OfficeA; rule 1 {match {source-address 192.168.210.0/24; destination-address 192.168.200.0/24;}then {source-nat { interface; }}}}

B. [edit security nat static]user@OfficeA# show rule-set From-Office-B { from interface ge-0/0/0.0;rule 1 { match {destination-address 192.168.200.0/24;}then { static-nat {prefix { 192.168.100.0/24; }}}}}

C. [edit security nat static]user@OfficeB# show rule-set From-Office-A { from interface ge-0/0/0.0;rule 1 { match {destination-address 192.168.210.0/24;}then { static-nat {prefix { 192.168.100.0/24; }}}}}

D. [edit security nat source]user@OfficeB# show rule-set OfficeAtoB { from zone OfficeA;to zone OfficeB; rule 1 {match {source-address 192.168.200.0/24; destination-address 192.168.210.0/24;}then {source-nat { interface; }}}}

Answer: BC

Explanation:

* 1. Static NAT Configuration at Office A (Option B):

? Configuration:

```
[edit security nat static]
```

```
user@OfficeA# show rule-set From-Office-B { from interface ge-0/0/0.0;
```

```
rule 1 { match {
```

```
destination-address 192.168.200.0/24;
```

```
}
```

```
then { static-nat {
```

```
prefix { 192.168.100.0/24; }
```

```
}
```


connections to the internal client using the same NAT translation. It's essential for applications that require consistent NAT mappings across multiple sessions.

Command Example: bash

```
set security nat source persistent-nat permit target-host-port
```

These features ensure that applications with multiple TCP sessions work seamlessly across NAT.

: Juniper NAT and persistent NAT documentation.

=====

NEW QUESTION 24

Referring to the exhibit,

```
security {
  advance-policy-based-routing {
    profile profile1 {
      rule Web-Proxy {
        match {
          dynamic-application [ junos:HTTP junos:HTTPS ];
        }
        then {
          routing-instance R1;
        }
      }
      rule DNS {
        match {
          dynamic-application-group junos:DNS;
        }
        then {
          routing-instance R2;
        }
      }
    }
  }
}
routing-instances {
  R1 {
    instance-type forwarding;
    routing-options {
      static {
        route 192.168.0.0/16 next-hop 10.1.0.1;
      }
    }
  }
}
```

which statement about TLS 1.2 traffic is correct?

- A. TLS 1.2 traffic will be sent to routing instance R1 but not forwarded to the next hop.
- B. TLS 1.2 traffic will be sent to routing instance R1 and forwarded to next hop 10.1.0.1.
- C. TLS 1.2 traffic will be sent to routing instance R2 but not forwarded to the next hop.
- D. TLS 1.2 traffic will be sent to routing instance R2 and forwarded to next hop 10.2.0.1.

Answer: A

NEW QUESTION 28

You are asked to configure tenant systems.

Which two statements are true in this scenario? (Choose two.)

- A. A tenant system can have only one administrator.
- B. After successful configuration, the changes are merged into the primary database for each tenant system.
- C. Tenant systems have their own configuration database.
- D. You can commit multiple tenant systems at a time.

Answer: CD

Explanation:

Each tenant system maintains its own configuration database, isolating configurations from others, enhancing security and operational efficiency. Junos OS supports multiple concurrent commit operations across tenant systems. Further details are covered in the Juniper Tenant System Guide.

When configuring tenant systems on an SRX device, the following principles apply:

? Tenant Systems Have Their Own Configuration Database (Answer C): Each tenant system has its own isolated configuration database, ensuring that changes made in one tenant system do not affect others. This allows for multi-tenant environments where different tenants can have independent configurations.

? Commit Multiple Tenant Systems Simultaneously (Answer D): The system allows for multiple tenant systems to be committed at the same time, simplifying management when working with multiple tenants. This is particularly useful in large environments where multiple logical systems or tenants need updates simultaneously.

: Juniper documentation on tenant systems and configuration databases.

=====

NEW QUESTION 30

Referring to the exhibit, you are assigned the tenantSYS1 user credentials on an SRX series device.

In this scenario, which two statements are correct? (Choose two.)

- A. When you log in to the device, you will be located at the operational mode of the main system hierarchy.
- B. When you log in to the device, you will be located at the operational mode of the Tenant.SY51 logical system hierarchy.
- C. When you log in to the device, you will be permitted to view only the routing tables for the Tenant SYS1 logical system.
- D. When you log in to the device, you will be permitted to view all routing tables available on the on an SYS1 Series device.

Answer: BC

NEW QUESTION 32

You have configured the backup signal route IP for your multinode HA deployment, and the ICL link fails. Which two statements are correct in this scenario? (Choose two.)

- A. The current active node retains the active role.
- B. The active node removes the active signal route.
- C. The backup node changes the routing preference to the other node at its medium priority.
- D. The active node keeps the active signal route.

Answer: AC

NEW QUESTION 36

You are deploying a large-scale VPN spanning six sites. You need to choose a VPN technology that satisfies the following requirements:

- ? All sites must have secure reachability to all other sites.
- ? New spoke sites can be added without explicit configuration on the hub site.
- ? All spoke-to-spoke communication must traverse the hub site. Which VPN technology will satisfy these requirements?

- A. ADVPN
- B. Group VPN
- C. Secure Connect VPN
- D. AutoVPN

Answer: D

Explanation:

AutoVPN simplifies deployment by dynamically establishing tunnels from spokes to the hub. This architecture supports easy scaling with minimal configuration changes, ensuring spoke-to-spoke traffic flows through the hub. For more information, see Juniper AutoVPN Overview.

In this scenario, you need a VPN solution that ensures secure, dynamic connectivity between multiple sites, with the following conditions:

- ? All sites must have secure reachability.
- ? New spoke sites can be added without explicit configuration on the hub site.
- ? Spoke-to-spoke communication must traverse the hub.

The correct technology to meet these requirements is AutoVPN. It simplifies VPN configurations by automating the setup between hub and spoke sites.

Additionally, AutoVPN automatically establishes secure tunnels for new spoke sites without requiring manual configuration at the hub, and all spoke-to-spoke traffic is routed through the hub.

: Juniper AutoVPN technology for dynamic VPN setups.

=====

NEW QUESTION 37

Exhibit:

```
[edit class-of-service]
user@srx# show
classifiers {
  dscp ba-classifier {
    import default;
    forwarding-class best-effort {
      loss-priority high code-points 000000;
    }
    forwarding-class ef-class {
      loss-priority high code-points 000001;
    }
    forwarding-class af-class {
      loss-priority high code-points 001010;
    }
    forwarding-class network-control {
      loss-priority high code-points 000011;
    }
    forwarding-class res-class {
      loss-priority high code-points 000100;
    }
    forwarding-class web-data {
      loss-priority high code-points 000101;
    }
    forwarding-class control-data {
      loss-priority high code-points 000111;
    }
    forwarding-class voip-data {
      loss-priority high code-points 000110;
    }
  }
}
```

You have configured a CoS-based VPN that is not functioning correctly. Referring to the exhibit, which action will solve the problem?

- A. You must delete one forwarding class.
- B. You must change the loss priorities of the forwarding classes to low.
- C. You must use inet precedence instead of DSCP.
- D. You must change the code point for the DB-data forwarding class to 10000.

Answer: A

Explanation:

In the exhibit, the CoS-based VPN configuration is not functioning correctly due to an issue with the number of forwarding classes. The maximum number of forwarding classes supported for CoS-based VPNs with multiple SAs (security associations) is typically four forwarding classes. In this case, more than four forwarding classes are defined.

To solve the issue, one forwarding class must be deleted to ensure that the total number of forwarding classes is reduced to four or fewer.

: Juniper CoS-based VPNs and forwarding class limitations.

=====

NEW QUESTION 41

Which two statements about policy enforcer and the forescout integration are true? (Choose two)

- A. 802.1X authenticated devices are supported.
- B. 802.1X authenticated devices are not supported.
- C. A Forescout CounterACT agent must be installed on third-party devices
- D. A Forescout CounterACT agent is agentless and does not need to be installed on third-party device

Answer: AD

NEW QUESTION 42

Exhibit:

Exhibit

```
[edit routing-instances]
user@vSRX-1# show
APBR-1 {
  routing-options {
    static {
      route 0.0.0.0/0 next-hop 172.16.9.2;
    }
  }
}
[edit routing-options]
user@vSRX-1# show
interface-routes {
  rib-group inet APBR-group;
}
static {
  route 0.0.0.0/0 next-hop 192.168.101.1;
}
rib-groups {
  APBR-group {
    import-rib [ inet.0 APBR-1.inet.0 ];
  }
}
[edit security advance-policy-based-routing]
user@vSRX-1# show
profile APBR-profile {
  rule ssh {
    match {
      dynamic-application junos:SSH;
    }
  }
}
```

Exhibit

```
import-rib [ inet.0 APBR-1.inet.0 ];
}
}
[edit security advance-policy-based-routing]
user@vSRX-1# show
profile APBR-profile {
  rule ssh {
    match {
      dynamic-application junos:SSH;
    }
    then {
      routing-instance APBR-1;
    }
  }
}
from-zone DC9-zone {
  policy move-ssh {
    match {
      source-address any;
      destination-address any;
      application any;
    }
    then {
      application-services {
        advance-policy-based-routing-profile APBR-profile;
      }
    }
  }
}
```

You are having problems configuring advanced policy-based routing. What should you do to solve the problem?

- A. Apply a policy to the APBR RIB group to only allow the exact routes you need.
- B. Change the routing instance to a forwarding instance.
- C. Change the routing instance to a virtual router instance.
- D. Remove the default static route from the main instance configuration.

Answer: B

NEW QUESTION 43

You need to generate a certificate for a PKI-based site-to-site VPN. The peer is expecting to use your domain name vpn.juniper.net.

Which two configuration elements are required when you generate your certificate request? (Choose two.)

- A. ip-address 10.100.0.5
- B. subject CN=vpn.juniper.net
- C. email admin@juniper.net
- D. domain-name vpn.juniper.net

Answer: BD

NEW QUESTION 47

Which three statements about persistent NAT are correct? (Choose Three)

- A. New sessions can only be initiated from a source towards the reflexive address.
- B. New sessions can be initiated from a destination towards the reflexive address.
- C. Persistent NAT only applies to source NAT.
- D. All requests from an internal address are mapped to the same reflexive address.
- E. Persistent NAT applies to both destination and source NAT.

Answer: BCD

NEW QUESTION 51

You are using AutoVPN to deploy a hub-and-spoke VPN to connect your enterprise sites. In this scenario, which two statements are true? (Choose two.)

- A. New spoke sites can be added without explicit configuration on the hub.
- B. Direct spoke-to-spoke tunnels can be established automatically.
- C. All spoke-to-spoke IPsec communication will pass through the hub.

D. AutoVPN requires OSPF over IPsec to discover and add new spokes.

Answer: AC

NEW QUESTION 55

Click the Exhibit button.

```
[edit class-of-service]
user@srx# show
classifiers {
  dscp ba-classifier {
    import default;
    forwarding-class best-effort {
      loss-priority high code-points 000000;
    }
    forwarding-class ef-class {
      loss-priority high code-points 000001;
    }
    forwarding-class af-class {
      loss-priority high code-points 001010;
    }
    forwarding-class network-control {
      loss-priority high code-points 000011;
    }
    forwarding-class res-class {
      loss-priority high code-points 000100;
    }
    forwarding-class web-data {
      loss-priority high code-points 000101;
    }
  }
}
```

You have configured a CoS-based VPN that is not functioning correctly. Referring to the exhibit, which action will solve the problem?

- A. You must change the loss priorities of the forwarding classes to low.
- B. You must change the code point for the DB-data forwarding class to 10000.
- C. You must use inet precedence instead of DSCP.
- D. You must delete one forwarding class.

Answer: D

Explanation:

Comprehensive Detailed Step-by-Step Explanation with All Juniper Security References

Understanding the Problem:

? A CoS-based VPN has been configured but is not functioning correctly.

? The exhibit shows that under the class-of-service configuration, six forwarding classes are defined.

Forwarding Classes in the Exhibit:

? best-effort

? ef-class

? af-class

? network-control

? res-class

? web-data

Juniper CoS-Based VPN Limitations:

? Maximum Number of Forwarding Classes: In CoS-based VPNs (Layer 3 VPNs), there is a limitation on the number of forwarding classes that can be used.

? Supported Forwarding Classes: Only up to four forwarding classes are supported in an L3VPN for CoS purposes.

Reference:

Juniper Networks Documentation:

"For Layer 3 VPNs, the maximum number of forwarding classes supported is four. If you configure more than four forwarding classes, CoS functionality might not work as expected."

Source: Juniper TechLibrary - Class of Service Limitations in VPNs

* Explanation:

Issue Identification:

The VPN is not functioning correctly because it exceeds the maximum number of supported forwarding classes for a CoS-based VPN.

Solution:

Option D: You must delete one forwarding class.

By reducing the number of forwarding classes to four or fewer, the CoS-based VPN will comply with the limitations and function correctly.

Why Other Options Are Incorrect:

Option A: You must change the loss priorities of the forwarding classes to low.

Changing loss priorities does not affect the limitation on the number of forwarding classes.

The issue is not related to loss priority settings but to the number of forwarding classes. Option B: You must change the code point for the DB-data forwarding class to 10000. There is no forwarding class named DB-data in the exhibit.

Changing a code point does not address the issue of exceeding the maximum number of forwarding classes.

Option C: You must use inet precedence instead of DSCP.

Switching from DSCP to IP Precedence does not resolve the issue of having too many forwarding classes.

The limitation on the number of forwarding classes remains the same regardless of the classification method used.

Conclusion:

To resolve the issue with the CoS-based VPN not functioning correctly due to exceeding the maximum number of forwarding classes, you must delete forwarding classes to reduce the total number to four or fewer.

* Answer: D. You must delete one forwarding class.

Additional References: Juniper TechLibrary:

"Configuring Class of Service for MPLS VPNs" - Discusses CoS considerations and limitations in MPLS L3VPN deployments.

Source: Juniper TechLibrary - CoS for VPNs

Juniper Networks Day One Book:

"Deploying MPLS Layer 3 VPNs" - Provides insights into CoS limitations and best practices for VPN deployments.

NEW QUESTION 56

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual JN0-637 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the JN0-637 Product From:

<https://www.2passeasy.com/dumps/JN0-637/>

Money Back Guarantee

JN0-637 Practice Exam Features:

- * JN0-637 Questions and Answers Updated Frequently
- * JN0-637 Practice Questions Verified by Expert Senior Certified Staff
- * JN0-637 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * JN0-637 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year