ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

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.

Question: 2

Exhibit:

```
Auj 3 02:10:28 02:10:26.045090;CiD-O;THREN>_ID-Ol;RT; <10.10.101.W/60858->10.10.102.10/22:6,0*0 matched filter filter-1

Aug 3 02:10:28 02:10:2 8.04510 0: CID-O: THREADED-01: RT: no session found, start first path, in_tunnel - 0x3, from_cp_flag - 0
Aug s 52:10:28 02 :1 1:26.045154 :OZE— .■ :THREAC_ID-Q1 :RI: flow firstcraatesessiar.

Aug 3 02:10:28 02:10:28.045143:CID-0:THREAD_IE-01:RT: routed (x_dst ip 10.10.102.10) from trust (g*-0/0/4.0 in 0) to ge- u/0/5.0, Next-hop: 10.10*102.10
Aug 3 02:10:2# 02:10:28.G45158:CiD-0:TERSAr ~-- .'“ I •st_cclicy_search: policy search from zone trust-> zone dmz
< 0 xC, jxe it a. 616,0x1€)

Aug  3  02:10:28  02:10:28.045191:CID-0:THREAD_ID-01:RI:  packer  dropped, denied by policy
Aug  3  02:10:28  02:10;28.045192:CID-0:THREAD_ID-01:RT:  denied  by policy default-palicy-logical-systea-OG€2), dropping pkt
Aug  3  02:10:26  02:10:28.045192:CID-0;IHP^AD_ID-01:RT:  packet  dropped, policy deny.
Aug  3  0 2:10:2 8  02:10:28.045195:CID-0:THREAD_ID-Q1:RT:  flowJMtiat^firstjpatb: first pak  no  session
```

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

A. The packet is dropped by the default security policy.
B. The packet is dropped by a configured security policy.
C. The data shown requires a traceoptions flag of host-traffic.
D. The data shown requires a traceoptions flag of basic-datapath.

Answer: AD

Explanation:

Understanding the Flow Log Output:

From the flow logs in the exhibit, we can observe the following key events:

The session creation was initiated (flow_first_create_session), but the policy search failed (flow_first_policy_search), which implies that no matching policy was found between the zones involved (zone trust-> zone dmz).

The packet was dropped with the reason "denied by policy." This shows that the packet was dropped either due to no matching security policy or because the default policy denies the traffic (packet dropped, denied by policy).

The line denied by policy default-policy-logical-system-00(2) indicates that the default security policy is responsible for denying the traffic, confirming that no explicit security policy was configured to allow this traffic.

Explanation of Answer A (Dropped by the default security policy):

The log message clearly states that the packet was dropped by the default security policy (defaultpolicy-logical-system-00). In Junos, when a session is attempted between two zones and no explicit policy exists to allow the traffic, the default policy is to deny the traffic. This is a common behavior in Junos OS when a security policy does not explicitly allow traffic between zones.

Explanation of Answer D (Requires traceoptions flag of basic-datapath):

The information displayed in the log involves session creation, flow policy search, and packet dropping due to policy violations, which are all part of basic packet processing in the data path. This type of information is logged when the traceoptions flag is set to basic-datapath. The basic-datapath traceoption provides detailed information about the forwarding process, including policy lookups and packet drops, which is precisely what we see in the exhibit.

The traceoptions flag host-traffic (Answer C) is incorrect because host-traffic is typically used for traffic destined to or generated from the Junos device itself (e.g., SSH or SNMP traffic to the SRX device), not for traffic passing through the device.

To capture flow processing details like those shown, you need the basic-datapath traceoptions flag, which provides details about packet forwarding and policy evaluation.

Step-by-Step Configuration for Tracing (Basic-Datapath):

Enable flow traceoptions:

To capture detailed information about how traffic is being processed, including policy lookups and flow session creation, enable traceoptions for the flow.

bash

```
set security flow traceoptions file flow-log
set security flow traceoptions flag basic-datapath
```

Apply the configuration and commit:

bash

```
commit
```

View the logs:

Once enabled, you can check the trace logs for packet flows, policy lookups, and session creation details:

bash

```
show log flow-log
```

This log will contain information similar to the exhibit, including session creation attempts and packet drops due to security policy.

Juniper Security Reference:

Default Security Policies: Juniper SRX devices have a default security policy to deny all traffic that is not explicitly allowed by user-defined policies. This is essential for security best practices. Reference: Juniper Networks Documentation on Security Policies.

Traceoptions for Debugging Flows: Using traceoptions is crucial for debugging and understanding how traffic is handled by the SRX, particularly when issues arise from policy misconfigurations or routing. Reference: Juniper Traceoptions.

By using the basic-datapath traceoptions, you can gain insights into how the device processes traffic, including

policy lookups, route lookups, and packet drops, as demonstrated in the exhibit.

Question: 3

Exhibit:

```
[edit]
user?srx< show security rat source (
    pool ipv4—source—pool ( address {
        10.10.101.10/32;

    rule-set ipv€-source ■. from zone trust; to zone untrust; rule ipv€-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

Explanation:

Question: 4

What are three core components for enabling advanced policy-based routing? (Choose three.)

A. Filter-based forwarding
B. Routing options
C. Routing instance
D. APBR profile
E. Policies

Answer: ACD

Explanation:

To enable Advanced Policy-Based Routing (APBR) on SRX Series devices, three key components are necessary: filter-based forwarding, routing instances, and APBR profiles. Filter-based forwarding is utilized to direct specific traffic flows to a routing instance based on criteria set by a policy. Routing instances allow the traffic to be managed independently of the main routing table, and APBR profiles define how and when traffic should be forwarded. These elements ensure that APBR is flexible and tailored to the network's requirements. Refer to Juniper's APBR Documentation for more details.

Advanced policy-based routing (APBR) in Juniper's SRX devices allows the selection of different paths for traffic based on policies, rather than relying purely on routing tables. To enable APBR, the following core components are required:

Filter-based Forwarding (Answer A): Filter-based forwarding (FBF) is a technique used to forward traffic based on policies rather than the default routing table. It is essential for enabling APBR, as it helps match traffic based on filters and directs it to specific routes.

Configuration Example: bash

```
set firewall family inet filter FBF match-term source-address 192.168.1.0/24
set firewall family inet filter FBF then routing-instance custom-routing-instance
```

Routing Instance (Answer C): A routing instance is required to define the separate routing table used by APBR. You can create multiple routing instances and assign traffic to these instances based on policies. The traffic will then use the routes defined within the specific routing instance.

Configuration Example: bash

```
set routing-instances custom-routing-instance instance-type forwarding
set routing-instances custom-routing-instance routing-options static route 0.0.0.0/0 next-hop 10.10.10.1
```

APBR Profile (Answer D): The APBR profile defines the rules and policies for advanced policy-based routing. It allows you to set up conditions such as traffic type, source/destination address, and port, and then assign actions such as redirecting traffic to specific routing instances.

Configuration Example: bash

```
set security forwarding-options advanced-policy-based-routing profile apbr-profile match application http
set security forwarding-options advanced-policy-based-routing profile apbr-profile then routing-instance custom-routing-instance
```

Other Components:

Routing Options (Answer B) are not a core component of APBR, as routing options define the general behavior of

the routing table and protocols. However, APBR works by overriding these default routing behaviors using policies.
Policies (Answer E) are crucial in many network configurations but are not a core component of enabling APBR. APBR specifically relies on profiles rather than standard security policies.
Juniper Security Reference:
Advanced Policy-Based Routing (APBR): Juniper's APBR is a powerful tool that allows routing based on specific traffic characteristics rather than relying on static routing tables. APBR ensures that specific types of traffic can take alternate paths based on business or network needs. Reference: Juniper Networks APBR Documentation.

Question: 5

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

Answer: AC

Explanation:

Explanation of Answer A (Session Reclassification):
APBR (Advanced Policy-Based Routing) requires the session to be classified based on the specified rule, which can change midstream as additional packets are processed. If the session was already established before the APBR rule took effect, the traffic may not be correctly reclassified to match the new APBR rule, leading to IDP (Intrusion Detection and Prevention) processing instead of being bypassed. This can occur especially when the session was already established before the rule change. Explanation of Answer C (Application Services Bypass):
For APBR to work and bypass the IDP service, the application services bypass must be explicitly configured. Without this configuration, the APBR rule may redirect the traffic, but the IDP service will still inspect and potentially drop the traffic. This is especially important for traffic destined for specific sites like social media platforms where bypassing IDP is desired.
Example configuration for bypassing IDP services: bash

set security forwarding-options advanced-policy-based-routing profile <profile-name> application-services-bypass
Step-by-Step Resolution:
Reclassify the Session Midstream:
If the traffic was already being processed before the APBR rule was applied, ensure that the session is reclassified by terminating the current session or ensuring the APBR rule is applied from the start. Command to clear the session:
bash

clear security flow session destination-prefix <ip-address>
Configure Application Services Bypass:
Ensure that the APBR rule includes the application services bypass configuration to properly bypass IDP or any other security services for traffic that should not be inspected.
Example configuration:
bash

set security forwarding-options advanced-policy-based-routing profile <profile-name> application- services-bypass
Juniper Security Reference:
Session Reclassification in APBR: APBR requires reclassification of sessions in real-time to ensure midstream packets are processed by the correct rule. This is crucial when policies change dynamically or new rules are added.
Application Services Bypass in APBR: This feature ensures that security services such as IDP are bypassed for traffic that matches specific APBR rules. This is essential for applications where performance is a priority and security inspection is not necessary.

Question: 6

Which two statements are correct about mixed mode? (Choose two.)

A. Layer 2 and Layer 3 interfaces can use the same security zone.
B. IRB interfaces can be used to route traffic.
C. Layer 2 and Layer 3 interfaces can use separate security zones.
D. IRB interfaces cannot be used to route traffic.

Answer: BC

Explanation:

Question: 7

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

Explanation:

Question: 8

Exhibit:

```
user3srx> show ethernet-
switching globa--information
Global Configuration:
KA: aging interval                : 300 Enabled
MAC learning                      : Disabled €5536
MAC statistics                    : Disabled
MAC limit Count                   : Disabled
MAC limit hit                     : IPv4 - 1200 seconds
MAC packet action drop:           IPv6 - 1200 seconds 65536
MAC+IP aging interval :           NO
                                  1200
MAC+IP limit Count : MAC+IP       1200
limit reached : LE aging time     : Iransparent bridge
IE VLAN aging time : Global       Master
Mede                              : Disabled
RE state                          : Disabled
VXLAN Overlay load bal:           Disabled 0
VXLAN ECMP                        :
Fast Update                       :
Host Pkts GBP src tag : [edit
interfaces] usertsrx# show ge-
0/0/0 { unit 0 (
                family ethernet vlan {
                        members vl. 1;
                switching (
```

MAC+IP limit count : MAC-IPIPv6 - 1200 seconds
limit reached : IE aging time : 65536
LE VLAN aging time : Global
Kode : 1200
RE state : 1200
VXWi Overlay lead cal: Transparent bridge
VXLAN ECMP : Master
Fast update : Disabled
Host Pkts GBP sre tag : [edit Disabled
interfaces] userSsrx# snow ge- Disabled
0/0/0 (

family ethernet vlan •
members

switching (

vlOO;

ge-0/0/1 {

family met {
address
172.16.0.1/24;

In which mode is the SRX Series device?

A. Packet
B. Ethernet switching
C. Mixed
D. Transparent

Answer: C

Explanation:

Question: 9

You configure two Ethernet interfaces on your SRX Series device as Layer 2 interfaces and add them to the same VLAN. The SRX is using the default L2-learning setting. You do not add the interfaces to a security zone.
Which two statements are true in this scenario? (Choose two.)

A. You are unable to apply stateful security features to traffic that is switched between the two interfaces.
B. You are able to apply stateful security features to traffic that enters and exits the VLAN.
C. The interfaces will not forward traffic by default.

D. You cannot add Layer 2 interfaces to a security zone.

Answer: AC

Explanation:

When Ethernet interfaces are configured as Layer 2 and added to the same VLAN without being assigned to a security zone, they will not forward traffic by default. Additionally, because they are operating in a pure Layer 2 switching mode, they lack the capability to enforce stateful security policies. For further details, refer to Juniper Ethernet Switching Layer 2 Documentation.

Explanation of Answer A (Unable to Apply Stateful Security Features):
When two interfaces are configured as Layer 2 interfaces and belong to the same VLAN but are not assigned to any security zone, traffic switched between them is handled purely at Layer 2. Stateful security features, such as firewall policies, are applied at Layer 3, so traffic between these interfaces will not undergo any stateful inspection or firewalling by default.
Explanation of Answer C (Interfaces Will Not Forward Traffic):
In Junos, Layer 2 interfaces must be added to a security zone to allow traffic forwarding. Since the interfaces in this scenario are not part of a security zone, they will not forward traffic by default until assigned to a zone. This is a security measure to prevent unintended forwarding of traffic.
Juniper Security Reference:
Layer 2 Interface Configuration: Layer 2 interfaces must be properly assigned to security zones to enable traffic forwarding and apply security policies. Reference: Juniper Networks Layer 2 Interface Documentation.

Question: 10

Which two statements are true about the procedures the Junos security device uses when handling traffic destined for the device itself? (Choose two.)

A. If the received packet is addressed to the ingress interface, then the device first performs a security policy evaluation for the junos-host zone.
B. If the received packet is destined for an interface other than the ingress interface, then the device performs a security policy evaluation for the junos-host zone.
C. If the received packet is addressed to the ingress interface, then the device first examines the host-inbound-traffic configuration for the ingress interface and zone.
D. If the received packet is destined for an interface other than the ingress interface, then the device performs a security policy evaluation based on the ingress and egress zone.

Answer: BC

Explanation:

When handling traffic that is destined for itself, the SRX examines the host-inbound-traffic configuration for the ingress interface and the associated security zone. It evaluates whether the traffic should be allowed based on this configuration. Traffic not addressed to the ingress interface is handled based on security policies within the junos-host zone, which applies to traffic directed to the SRX itself. For more details, refer to Juniper Host Inbound Traffic Documentation.

When handling traffic that is destined for the SRX device itself (also known as host-bound traffic), the SRX follows a specific process to evaluate the traffic and apply the appropriate security policies. The junos-host zone is a special security zone used for managing traffic destined for the device itself, such as management traffic (SSH, SNMP, etc.).

Explanation of Answer B (Packet to a Different Interface):

If the packet is destined for an interface other than the ingress interface, the SRX performs a security policy evaluation specifically for the junos-host zone. This ensures that management or host-bound traffic is evaluated according to the security policies defined for that zone.

Explanation of Answer C (Packet to the Ingress Interface):

If the packet is addressed to the ingress interface, the device first checks the host-inbound-traffic configuration for the ingress interface and zone. This configuration determines whether certain types of traffic (such as SSH, HTTP, etc.) are allowed to reach the device on that specific interface.

Step-by-Step Handling of Host-Bound Traffic:

Host-Inbound Traffic: Define which services are allowed to the SRX device itself: bash

set security zones security-zone <zone-name> host-inbound-traffic system-services ssh

Security Policy for junos-host: Ensure policies are defined for managing traffic destined for the SRX device: bash

set security policies from-zone <zone-name> to-zone junos-host policy allow-ssh match sourceaddress any

set security policies from-zone <zone-name> to-zone junos-host policy allow-ssh match destinationaddress any

Juniper Security Reference:

Junos-Host Zone: This special zone handles traffic destined for the SRX device, including management traffic. Security policies must be configured to allow this traffic. Reference: Juniper Networks Host-Inbound Traffic Documentation.

Question: 11

Exhibit:

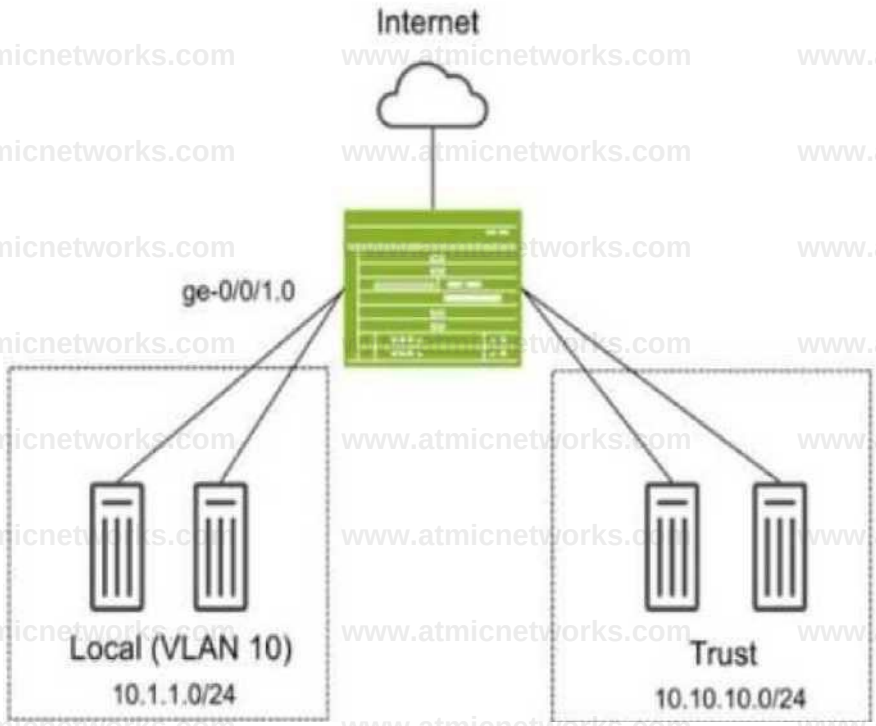

You have deployed an SRX Series device as shown in the exhibit. The devices in the Local zone have recently been added, but their SRX interfaces have not been configured. You must configure the SRX to meet the following requirements:

Devices in the 10.1.1.0/24 network can communicate with other devices in the same network but not with other networks or the SRX.

You must be able to apply security policies to traffic flows between devices in the Local zone.

Which three configuration elements will be required as part of your configuration? (Choose three.)

A. set security zones security-zone Local interfaces ge-0/0/1.0
B. set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan-members 10
C. set protocols l2-learning global-mode switching
D. set protocols l2-learning global-mode transparent-bridge
E. set security zones security-zone Local interfaces irb.10

Answer: ABD

Explanation:

In this scenario, we need to configure the SRX Series device so that devices in the Local zone (VLAN 10, 10.1.1.0/24 network) can communicate with each other but not with other networks or the SRX itself. Additionally, you must be able to apply security policies to traffic flows between the devices in the Local zone.

Explanation of Answer A (Assigning Interface to Security Zone):

You need to assign the interface ge-0/0/1.0 to the Local security zone. This is crucial because the SRX only applies security policies to interfaces assigned to security zones. Without this, traffic between devices in the Local zone won't be processed by security policies.

Configuration:

set security zones security-zone Local interfaces ge-0/0/1.0

Explanation of Answer B (Configuring Ethernet-Switching for VLAN 10):

Since we are using Layer 2 switching between devices in VLAN 10, we need to configure the interface to operate in Ethernet switching mode and assign it to VLAN 10.

Configuration:

set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan-members 10

Explanation of Answer D (Transparent Bridging Mode for Layer 2):

The global mode for Layer 2 switching on the SRX device must be set to transparent-bridge. This ensures that the SRX operates in Layer 2 mode and can switch traffic between devices without routing.

Configuration:

set protocols l2-learning global-mode transparent-bridge

Summary:

Interface Assignment: Interface ge-0/0/1.0 is assigned to the Local zone to allow policy enforcement.

Ethernet-Switching: The interface is configured for Layer 2 Ethernet switching in VLAN 10.

Transparent Bridging: The SRX is configured in Layer 2 transparent-bridge mode for switching between devices.

Juniper Security Reference:

Layer 2 Bridging and Switching Overview: This mode allows the SRX to act as a Layer 2 switch for forwarding traffic between VLAN members without routing. Reference: Juniper Transparent Bridging Documentation.

Question: 12

Exhibit:

user?peerl> show chassis high-availability information Neda failure codes: HW Hardware monitoring IE Looptack monitoring MB Mbuf monitoring s? SPU monitoring CS Cold Sync monitoring SV Software Upgrade Node Status: ONLINE Local-id: 1 Local-iP: 10.10.1.1 HA Peer Information:

Beer Id: 2 IP address: 10.10.1.2 Interface: ge-0/0/1.0 Routing Instance: default Encrypted: NO conn state: UP Cold Sync Status: COMPLETE Services Redundancy Group: . Current State: ONLINE Feer information: Peer Id: 2 SRG failure event codes: BF BFD monitoring IF I? monitoring IF Interface monitoring CP Control Plane monitoring Services Redundancy Group: - Deployment Type: SWITCHINS status: ACTIVE Activeness Priority: ICO Preemption: ENABLED Precess Packet In Backup State: NO Control Plane State: READY System integrity check: N/A Failure Events: NONE Peer Information: Feer id: 2 status : BACKUP Health Status: HEALTHY Failover Readiness: PEADY

Referring to the exhibit, which statement is true?

A. SRG1 is configured in hybrid mode.
B. The ICL is encrypted.
C. If SRG1 moves to peer 2, peer 1 will drop packets sent to the SRG1 interfaces.
D. If SRG1 moves to peer 2, peer 1 will forward packets sent to the SRG1 interfaces.

Answer: D

Explanation:

The exhibit describes a Chassis Cluster configuration with high availability (HA) settings. The key information is related to Service Redundancy Group 1 (SRG1) and its failover behavior between the

two peers.

Explanation of Answer D (Packet Forwarding after Failover):

In a typical SRX HA setup with active/backup configuration, if the SRG1 group moves to peer 2 (the backup), peer 1 (previously the active node) will forward packets to peer 2 instead of dropping them. This ensures smooth failover and seamless continuation of services without packet loss.

This behavior is part of the active/backup failover process in SRX chassis clusters, where the standby peer takes over traffic processing without disruption.

Juniper Security Reference:

Chassis Cluster Failover Behavior: When a service redundancy group fails over to the backup peer, the previously active peer forwards traffic to the new active node. Reference: Juniper Chassis Cluster Documentation.

Question: 13

You are asked to create multiple virtual routers using a single SRX Series device. You must ensure that each virtual router maintains a unique copy of the routing protocol daemon (RPD) process.

Which solution will accomplish this task?

A. Secure wire
B. Tenant system
C. Transparent mode
D. Logical system

Answer: D

Explanation:

Logical systems on SRX Series devices allow the creation of separate virtual routers, each with its unique RPD process. This segmentation ensures that routing and security policies are isolated across different logical systems, effectively acting like independent routers within a single SRX device. For further information, see Juniper Logical Systems Documentation.

To create multiple virtual routers on a single SRX Series device, each with its own unique copy of the routing protocol daemon (RPD) process, you need to use logical systems. Logical systems allow for the segmentation of an SRX device into multiple virtual routers, each with independent configurations, including routing instances, policies, and protocol daemons.

Explanation of Answer D (Logical System):

A logical system on an SRX device enables you to create multiple virtual instances of the SRX, each operating independently with its own control plane and routing processes. Each logical system gets a separate copy of the RPD process, ensuring complete isolation between virtual routers.

This is the correct solution when you need separate routing instances with their own RPD processes on the same physical device.
Configuration Example:
bash
set logical-systems <logical-system-name> interfaces ge-0/0/0 unit 0
set logical-systems <logical-system-name> routing-options static route 0.0.0.0/0 next-hop 192.168.1.1
Juniper Security Reference:
Logical Systems Overview: Logical systems allow for the creation of multiple virtual instances within a single SRX device, each with its own configuration and control plane. Reference: Juniper Logical Systems Documentation.

Question: 14

Click the Exhibit button.

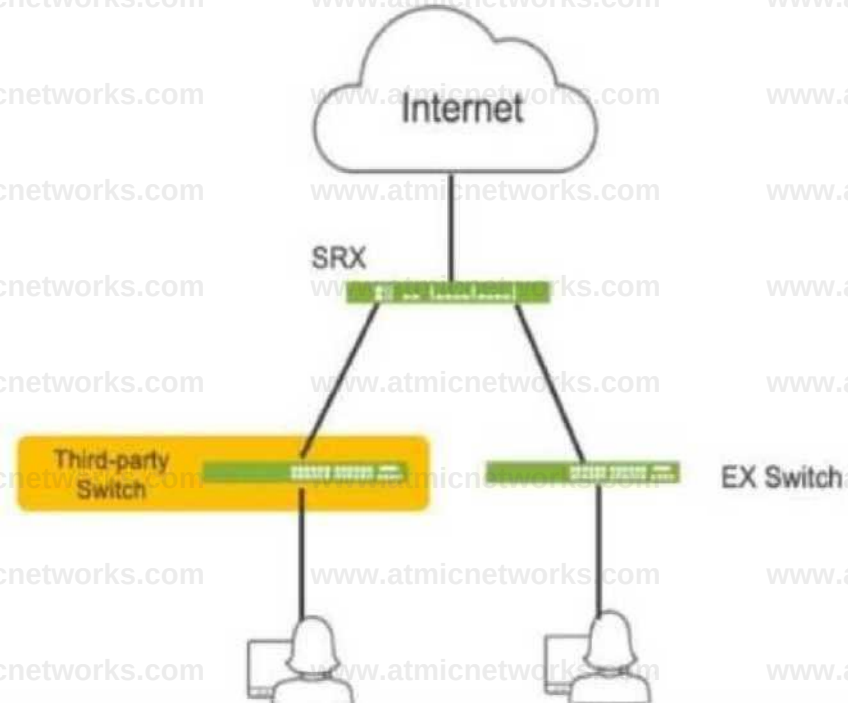

Referring to the exhibit, which three actions do you need to take to isolate the hosts at the switch port level if they become infected with malware? (Choose three.)

A. Enroll the SRX Series device with Juniper ATP Cloud.
B. Use a third-party connector.
C. Deploy Security Director with Policy Enforcer.
D. Configure AppTrack on the SRX Series device.
E. Deploy Juniper Secure Analytics.

Answer: ABC

Explanation:

A. Enroll the SRX Series device with Juniper ATP Cloud. This is essential for the SRX to receive threat intelligence from ATP Cloud, enabling it to identify infected hosts and take action.
B. Use a third-party connector. In this specific scenario, a third-party connector is required to

integrate the SRX with the third-party switch. While Juniper has native integration for its EX switches, a connector is necessary to communicate with and manage the third-party switch.
C. Deploy Security Director with Policy Enforcer. Security Director orchestrates the automated response, and Policy Enforcer translates the policies into device-specific commands for the SRX and the third-party switch (via the connector).

Question: 15

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

Explanation:

Question: 16

You are asked to connect two hosts that are directly connected to an SRX Series device. The traffic should flow unchanged as it passes through the SRX, and routing or switch lookups should not be performed. However, the traffic should still be subjected to security policy checks.
What will provide this functionality?

A. MACsec
B. Mixed mode
C. Secure wire
D. Transparent mode

Answer: C

Explanation:

Secure wire mode on SRX devices allows traffic to flow transparently through the firewall without being routed or switched, while still applying security policies. This is ideal for scenarios where traffic inspection is required without altering the traffic path or performing additional routing decisions. For further details on Secure Wire, refer to Juniper Secure Wire Documentation.

In this scenario, you want traffic to pass through the SRX unchanged (without routing or switching lookups) but still be subject to security policy checks. The best solution for this requirement is Secure

Wire.

Explanation of Answer C (Secure Wire):

Secure Wire allows traffic to flow through the SRX without any Layer 3 routing or Layer 2 switching decisions. It effectively bridges two interfaces at Layer 2 while still applying security policies. This ensures that traffic remains unchanged, while security policies (such as firewall rules) can still be enforced.

This is an ideal solution when you need the SRX to act as a "bump in the wire" for security enforcement without changing the traffic or performing complex network lookups.

Juniper Security Reference:

Secure Wire Functionality: Provides transparent Layer 2 forwarding with security policy enforcement, making it perfect for scenarios where traffic needs to pass through unchanged. Reference: Juniper Secure Wire Documentation.

Question: 17

Which two statements are true when setting up an SRX Series device to operate in mixed mode? (Choose two.)

A. A physical interface can be configured to be both a Layer 2 and a Layer 3 interface at the same time.
B. User logical systems support Layer 2 traffic processing.
C. The SRX must be rebooted after configuring at least one Layer 3 and one Layer 2 interface.
D. Packets from Layer 2 interfaces are switched within the same bridge domain.

Answer: CD

Explanation:

In mixed mode, SRX devices can simultaneously handle Layer 2 switching and Layer 3 routing, but a reboot is required when configuring Layer 2 and Layer 3 interfaces to ensure the configuration takes effect. Layer 2 packets are switched within the defined bridge domain. Further guidance on SRX mixed mode can be found at Juniper Mixed Mode Documentation.

When an SRX Series device is configured in mixed mode, both Layer 2 switching and Layer 3 routing functionalities can be used on the same device. This enables the SRX to act as both a router and a switch for different interfaces. However, there are certain considerations:

Explanation of Answer C (Reboot Requirement):

After configuring the SRX to operate with at least one Layer 2 interface and one Layer 3 interface, the device needs to be rebooted. This is required to properly initialize the mixed mode configuration, as the SRX needs to switch between Layer 2 and Layer 3 processing modes.

Explanation of Answer D (Layer 2 Traffic Handling):

In mixed mode, traffic from Layer 2 interfaces is switched within the same bridge domain. A bridge domain

defines a Layer 2 broadcast domain, and packets from Layer 2 interfaces are forwarded based on MAC addresses within that domain.

Juniper Security Reference:

Mixed Mode Overview: Juniper SRX devices can operate in mixed mode to handle both Layer 2 and Layer 3 traffic simultaneously. Reference: Juniper Mixed Mode Documentation.

Question: 18

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

Explanation:

Question: 19

Exhibit:

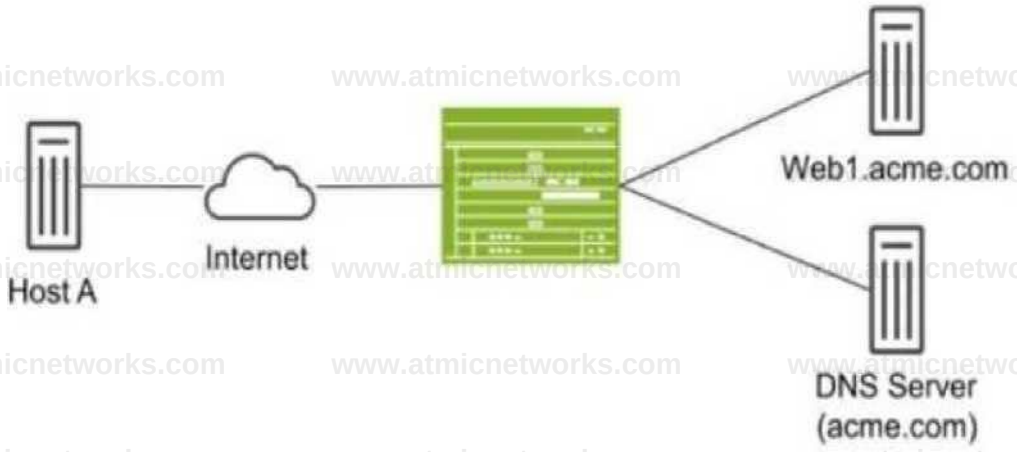

Host A shown in the exhibit is attempting to reach the Web1 webserver, but the connection is failing. Troubleshooting reveals that when Host A attempts to resolve the domain name of the server (web.acme.com), the request is resolved to the private address of the server rather than its public IP. Which feature would you configure on the SRX Series device to solve this issue?

A. Persistent NAT

B. Double NAT
C. DNS doctoring
D. STUN protocol

Answer: C

Explanation:

DNS doctoring modifies DNS responses for hosts behind NAT devices, allowing them to receive the correct public IP address for internal resources when queried from the public network. This prevents issues where private IPs are returned and are not reachable externally. For details, visit Juniper DNS Doctoring Documentation.

In this scenario, Host A is trying to resolve the domain name web.acme.com, but the DNS resolution returns the private IP address of the web server instead of its public IP. This is a common issue in networks where private addresses are used internally, but public addresses are required for external clients.

Explanation of Answer C (DNS Doctoring):

DNS doctoring is a feature that modifies DNS replies as they pass through the SRX device. In this case, DNS doctoring can be used to replace the private IP address returned in the DNS response with the correct public IP address for Host A. This allows external clients to reach internal resources without being aware of their private IP addresses.

Configuration Example: bash

```
set security nat dns-doctoring from-zone untrust to-zone trust
```

Juniper Security Reference:

DNS Doctoring Overview: DNS doctoring is used to modify DNS responses so that external clients can access internal resources using public IP addresses. Reference: Juniper DNS Doctoring Documentation.

Question: 20

Exhibit:

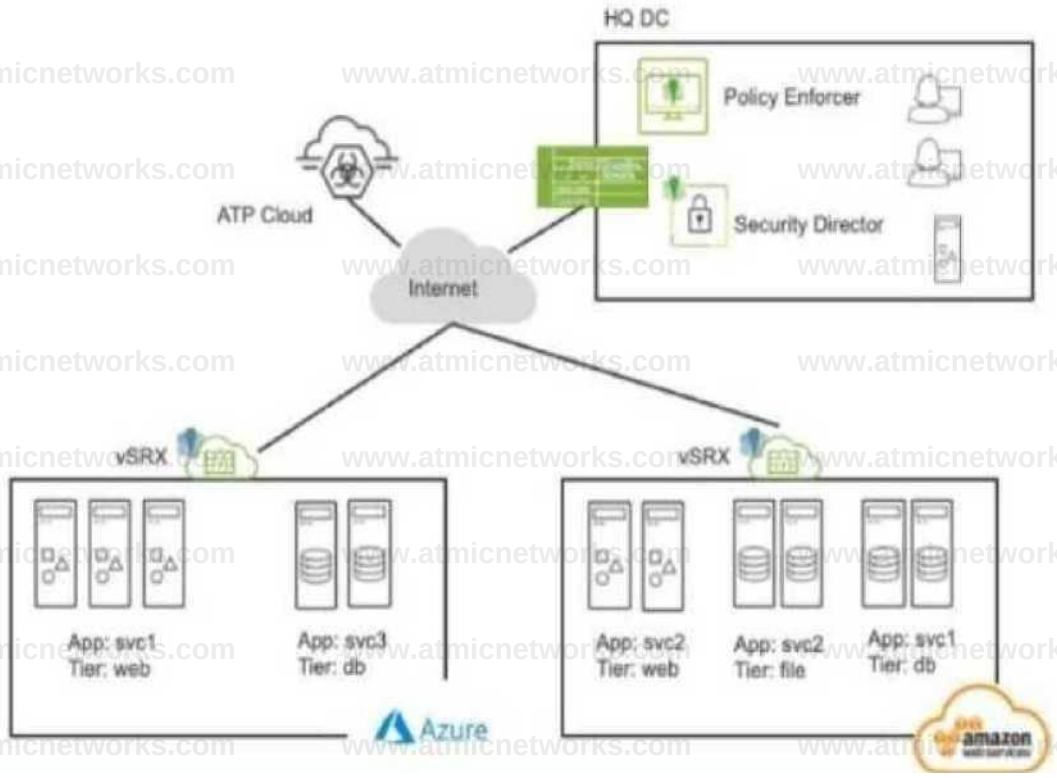

Referring to the exhibit, what do you use to dynamically secure traffic between the Azure and AWS clouds?

A. You can dynamically secure traffic between the clouds by using user identities in the security policies.
B. You can dynamically secure traffic between the clouds by using advanced connection tracking in the security policies.
C. You can dynamically secure traffic between the clouds by using security tags in the security policies.
D. You can dynamically secure traffic between the clouds by using URL filtering in the security policies.

Answer: C

Explanation:

Security tags facilitate dynamic traffic management between cloud environments like Azure and AWS. Tags allow flexible policies that respond to cloud-native events or resource changes, ensuring secure inter-cloud communication. For more information, see Juniper Cloud Security Tags.

In the scenario depicted in the exhibit, where traffic needs to be dynamically secured between Azure and AWS clouds, the best method to achieve dynamic security is by using security tags in the security policies.
Explanation of Answer C (Security Tags in Security Policies):
Security tags allow dynamic enforcement of security policies based on metadata rather than static IP addresses or zones. This is crucial in cloud environments, where resources and IP addresses can change dynamically.

Using security tags in the security policies, you can associate traffic flows with specific applications, services, or virtual machines, regardless of their underlying IP addresses or network locations. This ensures that security policies are automatically updated as cloud resources change.

Juniper Security Reference:

Dynamic Security with Security Tags: This feature allows you to dynamically secure cloud-based traffic using metadata and tags, ensuring that security policies remain effective even in dynamic environments. Reference: Juniper Security Tags Documentation.

Question: 21

Exhibit:

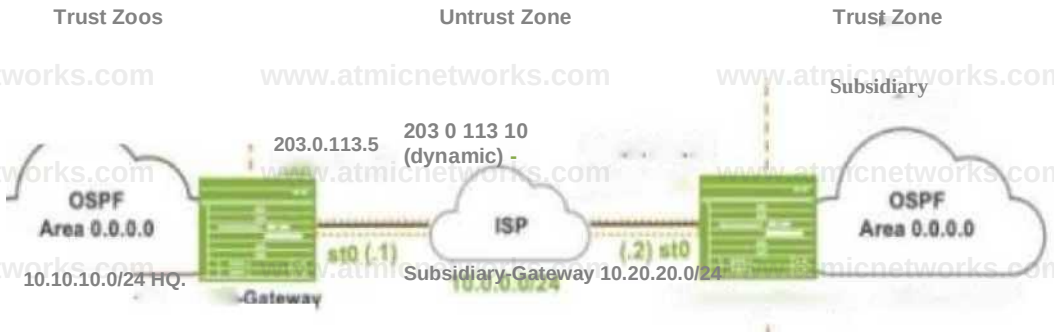

Referring to the exhibit, which IKE mode will be configured on the HQ-Gateway and SubsidiaryGateway?

A. Main mode on both the gateways
B. Aggressive mode on both the gateways
C. Main mode on the HQ-Gateway and aggressive mode on the Subsidiary-Gateway
D. Aggressive mode on the HQ-Gateway and main mode on the Subsidiary-Gateway

Answer: B

Explanation:

Question: 22

You are deploying threat remediation to endpoints connected through third-party devices.
In this scenario, which three statements are correct? (Choose three.)

A. All third-party switches must support AAA/RADIUS and Dynamic Authorization Extensions to the RADIUS protocol.
B. The connector uses an API to gather endpoint MAC address information from the RADIUS server.
C. All third-party switches in the specified network are automatically mapped and registered with the RADIUS server.
D. The connector queries the RADIUS server for the infected host endpoint details and initiates a change of authorization (CoA) for the infected host.
E. The RADIUS server sends Status-Server messages to update infected host information to the connector.

Answer: ABD

Explanation:

For threat remediation in a third-party network, the RADIUS protocol is necessary to communicate with the RADIUS server for details about infected hosts. CoA enables security measures to be enforced based on endpoint information provided by the RADIUS server. Details on this setup can be found in Juniper RADIUS and AAA Documentation.

When deploying threat remediation to endpoints connected through third-party devices, such as switches, the following conditions must be met for proper integration and functioning: Explanation of Answer A (Support for AAA/RADIUS and Dynamic Authorization Extensions): Third-party switches must support AAA (Authentication, Authorization, and Accounting) and RADIUS with Dynamic Authorization Extensions. These extensions allow dynamic updates to be made to a session's authorization parameters, which are essential for enforcing access control based on threat detection.
Explanation of Answer B (Connector Gathers MAC Information via API):
The connector uses an API to gather MAC address information from the RADIUS server. This MAC address data is necessary to identify and take action on infected hosts or endpoints.
Explanation of Answer D (Connector Initiates CoA):
The connector queries the RADIUS server for infected host details and triggers a Change of Authorization (CoA) for the infected host. The CoA allows the connector to dynamically alter the host's access permissions or isolate the infected host based on its threat status.
Juniper Security Reference:
Threat Remediation via RADIUS: Dynamic remediation actions, such as CoA, can be taken based on information received from the RADIUS server regarding infected hosts. Reference: Juniper RADIUS and CoA Documentation.

Question: 23

Exhibit:

```
user8srx> show ethernet-switching global-information
```

```
Global Configuration.:
MAC aging interval          : 300
MAC learning                : Enabled
MAC statistics              : Disabled
MAC limit Count             : €5536
MAC limit hit               : Disabled
MAC packet action drop: Disabled
MAC-IP agi: |        -: val  : IPv4 - 1200 seconds
                              IPv6 - 1200 seconds
MAC*IP limit Count          : 65536
MAC-IF limit reached        : NO
LE aging time               : 1200
IE VLAN aging time          : 1200
Global Mode                 : Transparent bridge
EE state                    : Master
```

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

A. You cannot secure intra-VLAN traffic with a security policy on this device.
B. You can secure inter-VLAN traffic with a security policy on this device.
C. The device can pass Layer 2 and Layer 3 traffic at the same time.
D. The device cannot pass Layer 2 and Layer 3 traffic at the same time.

Answer: BC

Explanation:

The exhibit provides information about an SRX Series device operating in transparent mode (Layer 2) and Layer 3 routing at the same time. Let's break down the correct answers: Explanation of Answer B (Secure Inter-VLAN Traffic with a Security Policy):
The SRX device can secure inter-VLAN traffic because it supports security policies for Layer 3 traffic between different VLANs. In this case, traffic moving between different VLANs (i.e., Layer 3 traffic) can be processed and controlled using security policies.
Explanation of Answer C (Pass Layer 2 and Layer 3 Traffic Simultaneously):
The SRX device can handle both Layer 2 and Layer 3 traffic simultaneously. In mixed mode, the device is capable of switching traffic at Layer 2 (intra-VLAN) while also routing traffic at Layer 3 (inter-VLAN). This is evident from the global configuration showing transparent bridge mode and Layer 3 interfaces.
Juniper Security Reference:
Mixed Mode Overview: Juniper SRX devices in mixed mode can operate as both a Layer 2 switch and a Layer 3 router, allowing it to pass traffic at both layers simultaneously. Reference: Juniper Mixed Mode Documentation.

Question: 24

You want to test how the device handles a theoretical session without generating traffic on the Junos

security device.
Which command is used in this scenario?

A. request security policies check
B. show security flow session
C. show security match-policies
D. show security policies

Answer: A

Explanation:

The request security policies check command allows you to simulate a session through the SRX device, checking the security policy action that would apply without needing to send real traffic. This helps in validating configurations before actual deployment. For more details, see Juniper Security Policies Testing.

The command request security policies check is used to test how a Junos security device handles a theoretical session without generating actual traffic. This command is useful for validating how security policies would be applied to a session based on various parameters like source and destination addresses, application type, and more.
Explanation of Answer A (request security policies check):
This command allows you to simulate a session and verify which security policies would be applied to the session. It's a proactive method to test security policy configurations without the need to generate real traffic.
Example usage: bash

request security policies check from-zone trust to-zone untrust source 10.1.1.1 destination 192.168.1.1 protocol tcp application junos-https
Juniper Security Reference:
Security Policies Check: This command provides a way to simulate and verify security policy behavior without actual traffic. Reference: Juniper Security Policy Documentation.

Question: 25

Exhibit:

user8srxl> show chassis hig.-.-availatility services-reclundancy-group 1 SRG failure event codes: BF EFC monitoring CP IF monitoring IF Interface monitoring CP Control Elate monitoring
Services Redundancy Group: 1
Deployment Type: SWITCHING Status: ACTIVF Activeness Priority: 200 Preemption: ENABLED
Process Packet in Backup state: NO Centre! Plane State: READY System Integrity Check: N/A Failure Events: NONE Peer Information:
Feer Id: 2
Status : BACKUP
Health Status: HEALTHY
Failover Readiness: READY
Virt-al IF Info: Index: 2
IP! 138.51.1...110/24

VMAC: N/A interface: ge-'/5/3.' Status: INSTALLED Index: 1
IP: 10.10.101.1/24

Feer Informal ion: Feer Id: 2 Status : BACKUP Health Status: HEALTHY Failover Readiness: READY
Virtual IF info: Index: 2
IF: 19e.51.lfj.ljj/24 VMAC: N/A
Interface: ge-0/0/3.0 Status: INSTALLED Index: 1
IF: 10.10.101.1/24
VMAC: N/A

Interface: ge-0/0/4.0

Status: INSTALLED
Split-brain Prevention Probe Info: DST-IP: 10.10.101.1
Routing instance: default status: NOT RUNNING
Result: N/A Reason: N/A
Interface Monitoring:
Status: UP
IF Name: ge-0/0/4 State: Up
IF Name: ge-0/0/3 State: Up
SRGID Iable:

SRGID	IF Prefix	Routing Table
1	198.51.100.100/32	default
1	10.10.101.1/32	default

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

A. The ge-0/0/3.0 and ge-0/0/4.0 interfaces are not active and will not respond to ARP requests to the virtual IP MAC address.
B. This device is the backup node for SRG1.
C. The ge-0/0/3.0 and ge-0/0/4.0 interfaces are active and will respond to ARP requests to the virtual IP MAC address.
D. This device is the active node for SRG1.

Answer: AB

Explanation:

The interfaces are active and respond to ARP for virtual IP as long as the node is the primary or active node in the SRG group. This ensures high availability and proper traffic forwarding. For information, refer to Juniper SRX HA Documentation.

The exhibit shows information about a chassis cluster and its services redundancy group (SRG1). Let's analyze the relevant details:
Explanation of Answer B (Backup Node for SRG1):
The exhibit indicates that this SRX device is in the backup role for SRG1. The status: BACKUP field confirms that this device is currently in a standby role and is not the active node for the services redundancy group.
Explanation of Answer A (Interfaces Not Active):
Since the device is in the backup role, the interfaces ge-0/0/3.0 and ge-0/0/4.0 will not respond to ARP requests for the virtual IP's MAC address. Only the active node's interfaces respond to ARP requests in a chassis cluster configuration.
Juniper Security Reference:
Chassis Cluster Redundancy Overview: In a chassis cluster, the backup node does not respond to ARP requests for

the virtual IP. Only the active node handles such requests to ensure seamless traffic forwarding. Reference: Juniper Chassis Cluster Documentation.

Question: 26

Which role does an SRX Series device play in a DS-Lite deployment?

A. Softwire concentrator
B. STUN server
C. STUN client
D. Softwire initiator

Answer: A

Explanation:

Question: 27

Which two statements are correct about the ICL in an active/active mode multinode HA environment? (Choose two.)

A. The ICL is strictly a Layer 2 interface.
B. The ICL uses a separate routing instance to communicate with remote multinode HA peers.
C. The ICL traffic can be encrypted.
D. The ICL is the local device management interface in a multinode HA environment.

Answer: BC

Explanation:

Question: 28

Exhibit:

Your company uses SRX Series devices to establish an IPsec VPN that connects Site-1 and the HQ networks. You want VoIP traffic to receive priority over data traffic when it is forwarded across the VPN.
Which three actions should you perform in this scenario? (Choose three.)

A. Enable next-hop tunnel binding.
B. Create a firewall filter that identifies VoIP traffic and associates it with the correct forwarding class.
C. Configure CoS forwarding classes and scheduling parameters.
D. Enable the copy-outer-dscp parameter so that DSCP header values are copied to the tunneled packets.
E. Enable the multi-sa parameter to enable two separate IPsec SAs for the VoIP and data traffic.

Answer: BCE

Explanation:

Question: 29

Your IPsec tunnel is configured with multiple security associations (SAs). Your SRX Series device supports the CoS-based IPsec VPNs with multiple IPsec SAs feature. You are asked to configure CoS for this tunnel.
Which two statements are true in this scenario? (Choose two.)

A. The local and remote gateways do not need the forwarding classes to be defined in the same order.
B. A maximum of four forwarding classes can be configured for a VPN with the multi-sa forwardingclasses statement.
C. The local and remote gateways must have the forwarding classes defined in the same order.
D. A maximum of eight forwarding classes can be configured for a VPN with the multi-sa forwardingclasses statement.

Answer: AD

Explanation:

Question: 30

The exhibit shows part of the flow session logs.

Mar 1 -31:26:23 01:28:23.434801;CID-0:THREAD_ID-d:RT;<172.2G.2ul-10/590Q9-H0.0.1.129/22; 6,0x0 Batched filter MatchTraffic:
Mar 1 01:28:23 01:28:23.434817:CID-0:THREADED-01:RT: ge-O/G/4.0:172.20.lD1.10/59D09->10.0.1.129/22, top, flag 2 3yn
Mar 7 01:28:23 01:28:23,434619:ciD-0:THREAD_ID-01:RT: find flow: table 0x2G€a€0a0, hash 43106(Oxffff), sa 172.20.101.10, da 10.0.1.129, sp 59009, dp 22, proto 6, tok 9, conn-tag 0x00000000
Mar ‘ .1:28:11 CL:2: :2i .434c£2;C11--:1HPZAL’ ID-Qi:RT: no session found, start first path- in tunnel - 0x1, from cp flag -
Mar 7 71:28:23 Cl:28:23.43482c :CID-0:THREAD_ID-0'. :FT: flow_first_create_session
Mar 7 ?!:28:23 01:28:23.434834:CID-0:THPEAD^ID-01:R": flowflrst_in_dsc_r.at: in <ge-0/0/3.0>, out <N/A> dst_adr 10.0.1.129, sp 59009, dp 22
Mar 7 01:28:23 01:28:23.434835:CIt-2:THREAD_ID-O1:RI: chose interface ge-G/0/4.0 as incoming nar if.
Mar 7 71:26:23 Ci:2~:23.434?36ICID-*:THREAD ID-II:FT:flow first rule dst xlate: DST no-xlate: :.:.2,*C to 10.0.1.129(221

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

A. The existing session is found in the table, and the fast path process begins.
B. This packet arrives on interface ge-0/0/4.0.
C. Junos captures a TCP packet from source address 172.20.101.10 destined to 10.0.1.129.
D. Destination NAT occurs.

Answer: BD

Explanation:

Question: 31

You have deployed automated threat mitigation using Security Director with Policy Enforcer, Juniper ATP Cloud, SRX Series devices, Forescout, and third-party switches.

In this scenario, which device is responsible for communicating directly to the third-party switches when infected hosts need to be blocked?

A. Forescout
B. Policy Enforcer
C. Juniper ATP Cloud
D. SRX Series device

Answer: B

Explanation:

Policy Enforcer receives these policies and translates them into device-specific commands. It then communicates with the third-party switches (using protocols like SNMP, RADIUS, or vendor-specific APIs) to enforce those commands, such as blocking the infected hosts' MAC addresses or port access. Why Policy Enforcer is the Right Choice:

Centralized Enforcement: Policy Enforcer acts as the central point of enforcement for Security Director policies, ensuring consistent security across the network.

Multi-Vendor Support: It can interact with a wide range of network devices, including switches from different vendors.

Automation: Policy Enforcer automates the policy enforcement process, enabling rapid response to threats.

Reference: Forescout and Juniper integration for network access control.

Question: 32

Referring to the exhibit,

```
from zone trust;
to zone untrust;
rule allow {
natch {
        source-address 172.16.1.0/24;
        destination-address 0.0.0.0/0;
then ( source-r.at { interface { persistent-nat ■.
        permit target-hest-p
```

which two statements are correct about the NAT configuration? (Choose two.)

A. Both the internal and the external host can initiate a session after the initial translation.
B. Only a specific host can initiate a session to the reflexive address after the initial session.
C. Any external host will be able to initiate a session to the reflexive address.
D. The original destination port is used for the source port for the session.

Answer: BD

Explanation:

Persistent NAT with target-host restricts session initiation to specific addresses, enhancing security. Reflexive NAT supports multiple connections by preserving the original port. Refer to Juniper NAT Configuration Documentation.

Referring to the NAT configuration shown in the exhibit:
Specific Host Can Initiate a Session (Answer B): The configuration uses persistent NAT with the permit target-host-port statement. This allows a specific external host (based on the target host and port used in the initial session) to initiate a session back to the internal host after the initial session has been established.
Persistent NAT ensures that the translation state is maintained, allowing external hosts to connect back only under specific conditions (e.g., the same target host and port as used in the original connection).
Original Destination Port (Answer D): The original destination port used by the internal host is retained as the source port when the session is established from outside to inside. This behavior is a result of how persistent NAT binds the internal and external sessions, ensuring that communication occurs over the same port used for the initial session.
Reference: Juniper NAT and Persistent NAT configuration documentation.

Question: 33

You are using ADVPN to deploy a hub-and-spoke VPN to connect your enterprise sites.
Which two statements are true in this scenario? (Choose two.)

A. ADVPN creates a full-mesh topology.
B. IBGP routing is required.
C. OSPF routing is required.
D. Certificate-based authentication is required.

Answer: CD

Explanation:

Question: 34

You want to create a connection for communication between tenant systems without using physical revenue ports on the SRX Series device.
What are two ways to accomplish this task? (Choose two.)

A. Use an external router.
B. Use an interconnect VPLS switch.
C. Use a secure wire.
D. Use a point-to-point logical tunnel.

Answer: BD

Explanation:

Question: 35

An ADVPN configuration has been verified on both the hub and spoke devices and it seems fine. However, OSPF is not functioning as expected.

```
[edit protocols cspf) user3ATVBN-HUB# show area 0.0.0.0 {
    interface stO.O { derar.d-circuit;

    interface ge- 0/jz 3. C passive;
```

Referring to the exhibit, which two statements under interface st0.0 on both the hub and spoke devices would

solve this problem? (Choose two.)

A. interface-type p2mp
B. dynamic-neighbors
C. passive
D. interface-type p2p

Answer: AB

Explanation:

For ADVPN with OSPF, using a point-to-multipoint (p2mp) interface type and enabling dynamic-

neighbors are crucial. This configuration allows dynamic discovery of neighbors and the establishment of tunnels. For more information, refer to Juniper ADVPN Configuration Guide.

In the ADVPN configuration, OSPF isn't functioning as expected due to the interface configuration on st0.0. Here are the adjustments needed:
Interface Type p2mp (Answer A): OSPF requires that the tunnel interface be set to p2mp (point-to- multipoint) to allow OSPF to communicate with multiple dynamic neighbors over the ADVPN tunnels.
Command Example: bash

set interfaces st0.0 family inet ospf interface-type p2mp
Dynamic Neighbors (Answer B): The dynamic neighbors statement allows OSPF to discover and communicate with dynamically established spokes in an ADVPN environment. This is essential for ADVPN to function properly since the tunnel endpoints are not static.
Command Example: bash

set protocols ospf area 0.0.0.0 interface st0.0 dynamic-neighbors
These settings ensure OSPF properly functions over dynamically created ADVPN tunnels.
Reference: Juniper ADVPN and OSPF configuration.

Question: 36

You have deployed an SRX Series device at your network edge to secure Internet-bound sessions for your local hosts using source NAT. You want to ensure that your users are able to interact with applications on the Internet that require more than one TCP session for the same application session. Which two features would satisfy this requirement? (Choose two.)

A. address persistence
B. STUN
C. persistent NAT
D. double NAT

Answer: AC

Explanation:

Address persistence ensures that the same NAT IP address is used for all sessions originating from a single source IP. Persistent NAT maintains connections for applications needing multiple sessions, like VoIP. Additional details are available in Juniper NAT Documentation.

For applications that require multiple TCP sessions for the same application session (such as VoIP or certain online games), the SRX device needs to handle NAT properly to maintain session continuity. Here's what helps:
Address Persistence (Answer A): Address persistence ensures that multiple sessions initiated by the same internal host are mapped to the same external IP address. This is crucial for applications that use multiple TCP sessions to maintain a stateful connection with the external server.

Command Example:

bash

set security nat source persistent-nat address-persistence
Persistent NAT (Answer C): This feature allows the external server to initiate new 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.
Reference: Juniper NAT and persistent NAT documentation.

Question: 37

Referring to the exhibit, security (advance - policy-baaed-routing ■ profile profilel { rule Web-Proxy { match I dynamic-application [junos:HTTP junos:HTTM J; then { routing-instance RI;

|

rule DNS { match (dynantrc-application-gxoup junosiDNS; then { routing-instance R2;

|

J routing-instances { RI (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 torouting instance R1 but not forwarded to the next hop.
B. TLS 1.2 traffic will be sent torouting instance R1 and forwarded to next hop 10.1.0.1.
C. TLS 1.2traffic will be sent to routing instance R2 but not forwarded to the next hop.
D. TLS 1.2 traffic will be sent torouting instance R2 and forwarded to next hop 10.2.0.1.

Answer: A

Explanation:

Question: 38

You have an initial setup of ADVPN with two spokes and a hub. A host at partner Spoke-1 is sending traffic to a host at partner Spoke-2.
In this scenario, which statement is true?

A. Spoke-1 will establish a VPN to Spoke-2 when this is first deployed, so traffic will be sent immediately to Spoke-2.
B. Spoke-1 will send the traffic through the hub and not use a direct VPN to Spoke-2.
C. Spoke-1 will establish the tunnel to Spoke-2 before sending any of the host traffic.
D. Spoke-1 will send the traffic destined to Spoke-2 through the hub until the VPN is established between the spokes.

Answer: A

Explanation:

Question: 39

Referring to the exhibit,

```
SRX(ttypO)
login: Userl
Password:
— JUNOS 22.4R1.9 built 2023-03-24 12:52:33
Userl8SRX:LSYS-l>
```

which two statements about User1 are true? (Choose two.)

A. User1 has access to the configuration specific to their assigned logical system.
B. User1 is logged in to logical system LSYS-1.
C. User1 can add logical units to an interface that a primary administrator has not previously assigned.

D. User1 can view outputs from other user logical systems.

Answer: AB

Explanation:

In this configuration, User1 is logged into logical system LSYS-1, which restricts access and visibility to that particular system. This ensures isolation between logical systems on the same physical device. Only a system administrator can assign additional permissions. For more details, see Juniper Logical Systems Guide.

From the exhibit, we see that User1 is logged into logical system LSYS-1:
Access to Assigned Logical System (Answer A): User1, being logged into the logical system LSYS-1, only has access to the configuration and interfaces within that logical system. This is a key feature of logical systems in Junos, ensuring users are restricted to their respective environments.
Logged into LSYS-1 (Answer B): The prompt shows that User1 is currently operating in LSYS-1, as indicated by the User1@SRX:LSYS-1> command line.
Reference: Juniper logical systems configuration and user permissions.

Question: 40

Exhibit:

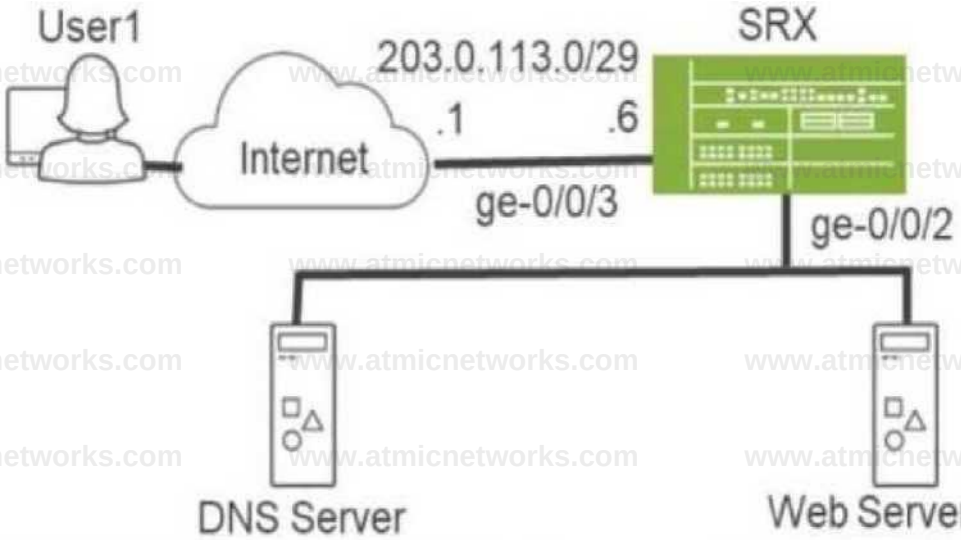

You are asked to ensure that Internet users can access the company's internal webserver using its

FQDN. However, the internal DNS server's A record only points to the webserver's private address. Referring to the exhibit, which two actions are required to complete this task? (Choose two.)

A. Disable the DNS ALG.
B. Configure static NAT for both the DNS server and the webserver.
C. Configure destination NAT for both the DNS server and the webserver.
D. Configure proxy ARP on ge-0/0/3.

Answer: BD

Explanation:

In the scenario where internal users are trying to access the company's web server via its FQDN but the DNS server resolves to a private IP, two key actions are needed:
Static NAT (Answer B): Since the internal DNS server resolves the web server to its private IP address (10.10.10.4/24), you need to configure static NAT for both the DNS server and the webserver. This will ensure that requests coming from the internet will be translated to the web server's public IP (203.0.113.4) and the DNS server's public IP (203.0.113.2).
Example Command: bash

set security nat static rule-set public-to-private from zone untrust
set security nat static rule-set public-to-private rule dns-server match destination-address 203.0.113.2/32
set security nat static rule-set public-to-private rule dns-server then static-nat-prefix 10.10.10.2/32 set security nat static rule-set public-to-private rule web-server match destination-address 203.0.113.4/32
set security nat static rule-set public-to-private rule web-server then static-nat-prefix 10.10.10.4/32 Proxy ARP (Answer D): The SRX needs to respond to ARP requests for the public IP addresses of both the DNS and webserver on the interface facing the internet (ge-0/0/3). This allows the SRX to handle requests directed at the public IPs.
Example Command:

set interfaces ge-0/0/3 unit 0 family inet proxy-arp interface-address 203.0.113.2/32
set interfaces ge-0/0/3 unit 0 family inet proxy-arp interface-address 203.0.113.4/32
These two configurations allow external users to access the internal web server via its public IP, as resolved by the DNS server.
Reference: Juniper NAT and proxy ARP documentation.

Question: 41

How does an SRX Series device examine exception traffic?

A. The device examines the host-inbound traffic for the ingress interface and zone.
B. The device examines the host-outbound traffic for the ingress interface and zone.
C. The device examines the host-inbound traffic for the egress interface and zone.
D. The device examines the host-outbound traffic for the egress interface and zone.

Answer: A

Explanation:

Exception traffic, including management and control plane traffic, is handled by examining hostinbound traffic configurations at the ingress interface and zone. It ensures traffic reaches necessary services like SSH and IKE securely. See Juniper Host Inbound Traffic Documentation for more.

SRX Series devices handle exception traffic (such as management traffic like SSH, Telnet, DNS queries, etc.) differently than regular transit traffic. Exception traffic is examined based on host-inbound traffic for the ingress interface and zone. If traffic is destined for the device itself (e.g., management traffic or routing protocol messages), it must be allowed as host-inbound traffic on both the ingress interface and zone.

Example Command: bash

set security zones security-zone trust host-inbound-traffic system-services ssh

This ensures that traffic destined to the SRX device is inspected based on the ingress interface and zone.

Reference: Juniper documentation on host-inbound traffic and exception handling.

Question: 42

Exhibit:

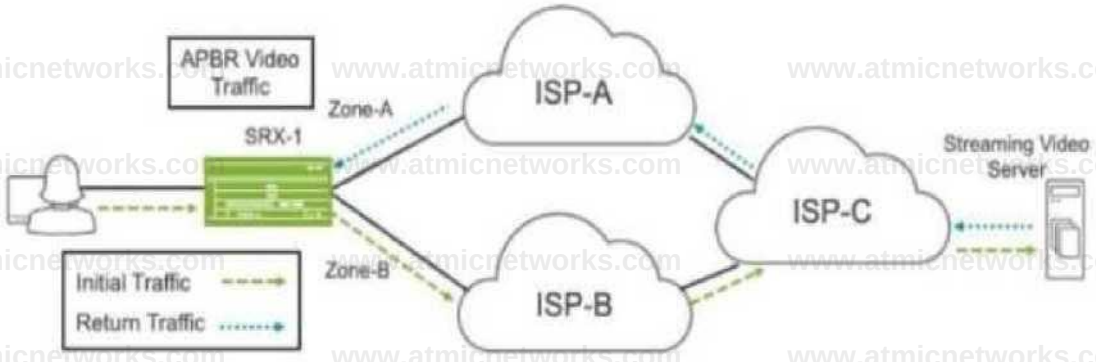

Referring to the exhibit, a default static route on SRX-1 sends all traffic to ISP-

A. You have configured APBR to send all requests for streaming video traffic to ISP-B. However, the return traffic from the streaming video server is coming through ISP-A, and the traffic is being dropped by SRX-1. You can only make changes on SRX-1.

How do you solve this problem?

A. Place both ISP-facing interfaces in the same zone.
B. Change the APBR routing instance from a forwarding instance to a virtual router instance.
C. Enable AppTrack to keep track of the sessions and zones for the streaming video traffic.
D. Configure BGP to control the return path of the streaming video traffic.

Answer: D

Explanation:

Question: 43

You are configuring an interconnect logical system that is configured as a VPLS switch to allow two logical systems to communicate.
Which two parameters are required when configuring the logical tunnel interfaces? (Choose two.)

A. Encapsulation ethernet must be used.
B. The virtual tunnel interfaces should only be configured with two logical unit pairs per logical system interconnect.
C. The logical tunnel interfaces should be configured with two logical unit pairs per logical system interconnect.
D. Encapsulation ethernet-vpls must be used.

Answer: CD

Explanation:

Question: 44

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 nigh code-points JOGCOI;
        forwarding-class af-class { loss-priority high code-points 001010;

        forwarding-class network-control loss-priority high code-points 00CC11;

        £:iwardmg-class res-class • loss-priority nigh code-points JOCISO;

        forwaxdir.g-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.

Reference: Juniper CoS-based VPNs and forwarding class limitations.

Question: 45

Exhibit:

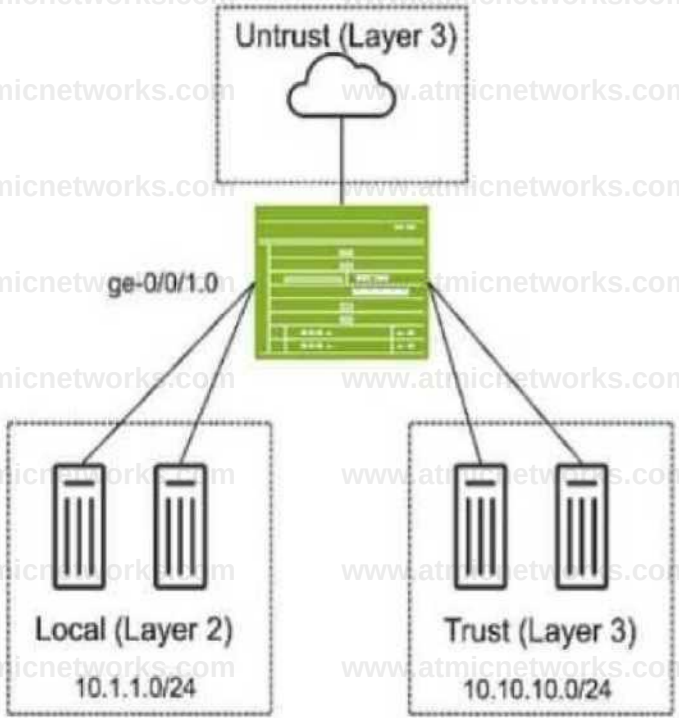


```
user@srx> show ethernet-switching
global-information
Global Configuration:

MAC aging interval     : 300
MAC learning           : Enabled
MAC statistics         : Disabled
MAC limit Count        : 65536
MAC limit hit          : Disabled
MAC packet action drop: Disabled
MAC+IP aging interval  : IPv4 - 1200 seconds
                         IPv6 - 1200 seconds
MAC+IP limit Count     : 65536
MAC+IP limit reached   : No
LE  aging time         : 1200
LE  VLAN aging time    : 1200
Global Mode            : Switching
FE state               : Master
Host Pkts GBP src tag  : 0
```

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

A. Hosts in the Local zone can be enabled for control plane access to the SRX.
B. An IRB interface is required to enable communication between the Trust and the Untrust zones.
C. You can configure security policies for traffic flows between hosts in the Local zone.
D. Hosts in the Local zone can communicate with hosts in the Trust zone with a security policy.

Answer: AD

Explanation:

Question: 46

Your customer needs embedded security in an EVPN-VXLAN solution.
What are two benefits of adding an SRX Series device in this scenario? (Choose two.)

A. It enhances tunnel inspection for VXLAN encapsulated traffic with Layer 4-7 security services. B. It adds extra security with the capabilities of an enterprise-grade firewall in the EVPN-VXLAN underlay.
C. It adds extra security with the capabilities of an enterprise-grade firewall in the EVPN-VXLAN overlay.
D. It enhances tunnel inspection for VXLAN encapsulated traffic with only Layer 4 security services.

Answer: AC

Explanation:

The SRX Series can inspect traffic within VXLAN tunnels, providing in-depth security services across multiple layers. Adding SRX in the overlay network allows comprehensive control, leveraging advanced firewall capabilities. For more details, see Juniper EVPN-VXLAN Security.

When integrating an SRX Series device into an EVPN-VXLAN solution, it offers several security benefits:
Layer 4-7 Security Services (Answer A): The SRX can provide deep packet inspection for VXLAN encapsulated traffic, enhancing security by offering services such as intrusion prevention, application layer filtering, and antivirus scanning. This allows security monitoring of the encapsulated traffic at higher layers of the OSI model (Layers 4-7), which is essential for advanced threat detection.
Security in the Overlay Network (Answer C): The SRX adds security by functioning as an enterprisegrade firewall within the EVPN-VXLAN overlay. This means that traffic flowing between virtualized segments or networks can be inspected and filtered using SRX firewall rules, ensuring that the VXLAN overlay remains secure.
These features make the SRX a powerful addition for securing EVPN-VXLAN environments, providing comprehensive security for encapsulated traffic and ensuring that both the underlay and overlay networks are protected.
Reference: Juniper documentation on SRX integration in EVPN-VXLAN solutions.

Question: 47

You want to use a security profile to limit the system resources allocated to user logical systems.
In this scenario, which two statements are true? (Choose two.)

A. If nothing is specified for a resource, a default reserved resource is set for a specific logical system. B. If you do not specify anything for a resource, no resource is reserved for a specific logical system, but the entire system can compete for resources up to the maximum available.
C. One security profile can only be applied to one logical system.
D. One security profile can be applied to multiple logical systems.

Answer: BD

Explanation:

When using security profiles to limit system resources in Juniper logical systems:
No Resource Specification (Answer B): If a resource limit is not specified for a logical system, no specific amount of system resources is reserved for it. Instead, the logical system competes for resources along with others in the system, up to the maximum available. This allows flexible resource

allocation, where logical systems can scale based on actual demand rather than predefined limits. Multiple Logical Systems per Security Profile (Answer D): A single security profile can be applied to multiple logical systems. This allows administrators to define resource limits once in a profile and apply it across several logical systems, simplifying management and ensuring consistency across different environments.
These principles ensure efficient and flexible use of system resources within a multi-tenant or multi- logical-

system environment.

Reference: Juniper security profiles and logical system documentation.

Question: 48

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.

Reference: Juniper documentation on tenant systems and configuration databases.

Question: 49

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.
Reference: Juniper AutoVPN technology for dynamic VPN setups.

Question: 50

You need to set up source NAT so that external hosts can initiate connections to an internal device, but only if a connection to the device was first initiated by the internal device.
Which type of NAT solution provides this functionality?

A. Address persistence
B. Persistent NAT with any remote host
C. Persistent NAT with target host
D. Static NAT

Answer: C

Explanation:

Persistent NAT with target host allows external hosts to establish connections only when the internal device initiates a session first, ideal for specific interactive applications. Refer to Juniper Persistent NAT Documentation.

The scenario requires that external hosts be able to initiate a connection only if the internal device has already initiated a connection. The correct solution is Persistent NAT with target host, which ensures that a specific external host can initiate new connections back to the internal device, but

only after the internal device has established a session first.
Persistent NAT with Target Host (Answer C): This allows the internal device to initiate a connection, and once established, the specified external host can also initiate new connections to the internal device on the same NAT mapping.
Example Configuration:
bash

set security nat source persistent-nat permit target-host-port
This solution is appropriate when controlled bidirectional communication is required based on an internal-initiated connection.

Reference: Juniper persistent NAT documentation.

Question: 51

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

A. Infected hosts are tracked by their IP address.
B. Infected hosts are tracked by their chassis serial number.
C. Infected hosts are tracked by their MAC address.
D. Infected hosts are tracked by their user identity.

Answer: AC

Explanation:

Question: 52

You have deployed two SRX Series devices in an active/passive multimode HA scenario.
In this scenario, which two statements are correct? (Choose two.)

A. Services redundancy group 1 (SRG1)is used for services that do not have a control plane state.
B. Services redundancy group 0 (SRG0)is used for services that have a control plane state.
C. Services redundancy group 0 (SRG0)is used for services that do not have a control plane state.
D. Services redundancy group 1 (SRG1)is used for services that have a control plane state.

Answer: CD

Explanation:

Question: 53

Which two statements are true regarding NAT64? (Choose two.)

A. An SRX Seriesdevice should bein packet-based forwarding mode for IPv4.
B. An SRX Seriesdevice should bein packet-based forwarding mode for IPv6.
C. An SRX Seriesdevice should bein flow-based forwarding mode for IPv4.
D. An SRX Seriesdevice should bein flow-based forwarding mode for IPv6.

Answer: BC

Explanation:

Question: 54

What is the advantage of using separate st0 logical units for each spoke connection?

A. It is easy to configure even when managing many st0 units.
B. It facilitates scalability.
C. Junos devices can exchange NHTB data automatically using this method.
D. It enables assignments of different settings to each logical unit.

Answer: D

Explanation:

Question: 55

You are asked to select a product offered by Juniper Networks that can collect and assimilate data from all probes and determine the optimal links for different applications to maximize the full potential of AppQoE.
Which product provides this capability?

A. Security Director
B. Network Director
C. Mist
D. Security Director Insights

Answer: C

Explanation:

Question: 56

You are asked to establish IBGP between two nodes, but the session is not established. To troubleshoot this problem, you configured trace options to monitor BGP protocol message exchanges.

Kar 7 02:39:15 02:33;15.353921 iCTD-O:TBREM.ID-OURT: <192.160.2.1/54882-MK .160.1.1/179,'6,0X0 > marched filter tbgp- tzaffio:

Mar 7 02:38:15 02:38 :15.353933:CID-O:THFEAD_TD-O1:RT: ge-0/0/3.0:192.158 .2 .l/54882->192.L«e. 1.1/179, tep, flag 2 sy-
Kar 7 02:38:15 02:30:15.353935:CID-0:TBREAD_ID-0i:RT-. find flow: table 0x206a60a0, hash 6149(0xffff), sa 192.168.2.1, da 192.168.1.1, sp 54382, dp 179, proto £, tok 8, conn-tag jxOOOOOOOu
Mar 7 02:38:15 02:38:15.353938:CIS>-0: THREA»_ID-01:M: no session found, start first path. in_nunnel - 0x0, fzom_cp_flag - 0
Mar 7 02:30:15 02:38:15.353941:CID-0:TBREAD_ID-01:RT: flow_firsP_createsessiOn

Mar 7 02:39:15 02:39:15.353964:CID-O:WR£AD_1D-01:M: Doing DESTIMATION adit zoute-looXup
Mar 7 02:30:15 02: 38 :15.353971: CID-0 : THREAD_ID-01: KT : flow ipv4_rt_lkup success 192.183.1.1, iifl 0x47, oifl 0x0
Hat 7 02:30:15 02:38:15.353975:CIW:TBREAD_ID-01:RT: Changing C«-ifp from .local..0 SC loO.O for dst: 192,160.1.1 in vz id : 0
Mar 7 02:38:15 02:38:15.353976: CID-0: iaRE6D_ID-ei:RT: routed !x_dst_ip 192.168.1.1) iron uncrust (ge-0/0/3.0 ir. 0) to IcO.O, Next-hop: 192.160.-l-X
Mar 7 02:38:15 02:38:15.35S978-:CiD-0UHREAD_iD-l 1:": fltw_first_policy_seazch: policy search from zone untrust-> rone tnist (0*0, 0xd66200b3, 0^:3)
Mar 7 02: 33: 15 02: 38 :15.155966:CID-u:THREAD_1D-O1:RT: Policy Ikup: ways 0 rone 15:global! -> sone(5:global) scope:0

Mar 7 02:39:15 02:30:15.354000:CID-0:THREADJO-01:RT) permitted by policy allow-bgp(6i
Kar 7 02:33:15 02:38:15.354048:CID-j:iHREM_ID-01:RT: flow_flrst_fir.al_ch.eck: Ln 0/3.05, out
Mar 7 02:30:15 02:30:15.354050:CID-0:IBREAD_ID-01:RT: In £-c-.jtcrsz_ccztLete_'etsiac
Kar 7 02:38:15 02:38:15.354051tciu-j:THaBAp_lD-01:RT: ilow_first_camplete_sessibn, pakjtr: ixicSfcdiJ nap: ax2al40340, in_tunnel: 0x0

Mar 7 02:30:15 12:36:15.353976:ClD“t:THREAD,ID-01:87: flow_first_polccy_’earch: policy search from rone utttr«t-> zone trust (0x0.0x466200b3.0xb3)
Mar 7 02:30:15 62:38:IS.353978 :CID-0:mS£AB_ID-'31:RT! flaw_first_poliey_search: pcliey search fieri zone entrust ■ > zone trust (OxO,Gxd662COb3,0xb3)
Mar 7 02:30:15 0-2:38;15.353906;CTD-0:THREAD,lE-OL:RT: Policy Ikup; vsys 0 zone(5:global) -> zone(5:global) scope;O

Max 7 02:30:10 M :3R: AS. 354006 .-CiL'-O: THREAD iz>-fll:.RT: permitted by policy allow-bgp(C)
Mar 7 02:30:15 02;30:15.354046:CI&-C:TaR6M_n)-01:RI: fl9w_fiz«_finai_checb; in 0/3.Oa, cut
Mar 7 82:38:18 :2:33:15.344056:CID-0:Tl®EAD_lD-0i:RT: In flaw_fizst_eMpleta_se53ian
Mat 7 02:38:15 . 2 ;38;1S.354051 :CID—0: THREAD, ID-01: RI; :1:^J i rat,Cvt^leto, sea slot, p*k_ptr: ;x2cSIcdL, nap: 0*24140311, ia_tanael :■ 0x0

Mar 7 02:30:15 52:38:15.353978:ciD-0:THFBAD_rE-81:XT: flow_fitst_policy_seazch: policy search from zone untrust-> zona trust (Oxi,CadiSlO&bSySxfeS)
Mar 7 02:30:15 :2:39 115.353956 :ClC-C:THRW>_lE-01:Jt71 Policy Ikup: vsys 0 sanelSiglsball -> zone 151 global) aecte:1

Mar 7 02:30:15 O2;3s;15.354000;CIC^O:THFEAD_ID-O1:RT: permitted by policy aliow-bgc(£)
Mar 7 02:38:15 02:38:15.354M0 :CID-0:THREM>_rD-01:»T: flw_flist final checi: in 0/3.■>, out
Max 7 02:30:15 02:38:15.154050; CID-0 :THBEAD_rs-01: RI: In rizw_rirst_ccnipiece_sessicn
Mar 7 02:39:15 02: 30:15.354051 :cic-i:ms£AO_tB-Cl:RI: tlow_:irst_coEplete_3es3ion, par_pzr: ix2c5fcd4?, nip: SzIaKOaiu, in_maelt 0x0

Mar 7 02:38:15 0-2:38:13.354055: CID-0 :IHRMD_IE-ai:Jffi: Bessies (id:2C.3?5) czMted for first rah 2
Mat 7 62:38:15 82:30:15.354073:CIE-0:TMiRD_ID-:<:XT: flaw_fitet_in_ifM_nM: it. , out A> Mt_adz 132.168.1.1, sp 54802, dp 179
Mar 7 72:38:1: 02:38:15.354075:CII>-v:THI®Ei_n3’-01 :KT; chose interface IQQ.C as incoming rar if.
Mar 7 62:38:15 02:38:18.3 54675:CID-0:THREAD_IB-61:RT: packer dropped: for self but not interested
Mar 7 02:30:15 02:30:IS.3540781CID-0! THREAD^ID-01:RT: packer dropped, packer dropped: for self bur norinterested.
Mar 7 £2:38:15 G2:3B: 15.354079 :ClD-O:THREAD_ID-'jl: RT: flaw_f iEMjtA*£alIsesalon: Loopback session processing aborted
Mar 7 02:39:15 82:38:15.3 5 4 6 8 0 :C!t>0:THRtAr=_TT-01: RT: first path session Installation failed
Mar 7 02130:13 0 2:3=:15.35 4081:CID-O:THREAD rD-01:RT: flow find session returns error.

Referring to the exhibit, which action would solve the problem?

A. Add the junos-host zone policy to permit the BGP packets.

B. Add a firewall filter to lo0 that permits the BGP packets.
C. Modify the security policy to permit the BGP packets.
D. Add BGP to the lo0 host-inbound-traffic configuration.

Answer: D

Explanation:

Question: 57

You are using trace options to troubleshoot a security policy on your SRX Series device.

user?SRX> show log flow-log l find "policy search"
Jan 5 14:15:3.7 14 :13: 37.325231 :CIi--j:THPXAD_XE-JliLS^i^il- LC.;RT:flov_firsz_poXicy_5earch: policy search front zone Linux-3 zone-> zone juros-host (2x3,Qx94cB3jl6,1X16F, result: ux5ed4h4€9, pending: 07, is_http_cached = '
Jan 9 14:13:37 14:19:37,520232:CID-0:THREAD_ID-01:LSYS_ID“00:RT:flow_fir$tj?olicy_search: dynapp_none_policy: TRUE, ue nonepclxcy; TROE, is_final: CxO, is_explicit: OxC, pdicyjneta^data: 0x0
Jan 9 14:19:37 14:19:37,520233 :CID-0: THREAD,. ID-jl:LSYS_XD-C0:RT: app 22, timeout 1600$, cuxx agecut 20s
Jan 9 14:19:37 14:19:37,S20234:ClD-u!THR£AD_lD-3i:LSYS.lD-vG:RT: packer dropped, denied by policy
4: J^7 (^&a?4A0^*GH>r&4m denied by policy deny-ssh(7), dropping pkt
Jar. 9 14:19:37 14:19:37.520235 :CIM:THREAD_ID-$1;I£YS_I1)H^ packet dropped, policy deny.

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

A. The SSH traffic matches an existing session.
B. No entries are created in the SRX session table.
C. The traffic is not destined for the root logical system.
D. The security policy controls traffic destined to the SRX device.

Answer: AD

Explanation:

Question: 58

You have deployed automated threat mitigation using Security Director with Policy Enforcer, Juniper ATP Cloud, SRX Series devices, and EX Series switches.
In this scenario, which device is responsible for blocking the infected hosts?

A. Policy Enforcer
B. Security Director
C. Juniper ATP Cloud
D. EX Series switch

Answer: A

Explanation:

Policy Enforcer interacts with other network elements like EX switches to enforce blocking of infected hosts based on threat intelligence from ATP Cloud and other sources. For more information, refer to Juniper Policy Enforcer Documentation.

In a Juniper automated threat mitigation setup involving Security Director, Policy Enforcer, Juniper ATP Cloud, SRX Series, and EX Series switches, the Policy Enforcer is the component responsible for blocking infected hosts. The role of each component is as follows:
Policy Enforcer (Correct: Option A):
Policy Enforcer receives threat intelligence from Juniper ATP Cloud and instructs SRX devices and EX Series switches to block or quarantine infected hosts. Policy Enforcer pushes policies to these devices to enforce the mitigation actions.
Security Director (Incorrect):
Security Director provides centralized management and visibility but does not directly enforce policies.
Juniper ATP Cloud (Incorrect):
Juniper ATP Cloud is responsible for analyzing threats and providing intelligence but does not take direct mitigation actions.
EX Series Switch (Incorrect):
EX Series switches can enforce the policy pushed by Policy Enforcer but are not responsible for deciding which hosts to block.
Juniper Reference:

Juniper ATP Cloud and Policy Enforcer Documentation: Details the roles of each component in the automated threat mitigation architecture.

Question: 59

Referring to the exhibit,

user8srx> show chassis high-availability information Nose failure coses: HW Hardware monitoring LB Loopback monitoring
MB Mbuf monitoring SF SPU monitoring
IS Cold Sync monitoring SU Software Upgrade
Node Status: ONLINE Local-id: 1
Local-IP: 10.10.1.1 HA Peer Information:
Peer Id: 2 IP address: 10.10.1.2 Interface: ge-0/0/1.0
Routing Instance: default Encrypted: NO Conn State: UP Cold Sync Status: COMPLETE Services Redundancy Group: 1 Current State: ONLINE Feer information:
Peer Id: 2
SRG failure event codes: BE EFE monitoring IF IF monitoring IF interface monitoring CP Control Plane monitoring
Services Redundancy Group: _ Deployment Type: SWITCHING Status: ACTIVE Activeness Priority: 200 Preemption: ENABLED Process Packet in Backup State: NO

which three statements about the multinode HA environment are true? (Choose three.)

A. Two services redundancy groups are available.
B. IP monitoring has failed for the services redundancy group.
C. Node 1 will host services redundancy group 1 unless it is unavailable.
D. Session state is synchronized on both nodes.
E. Node 2 will process transit traffic that it receives for services redundancy group 1.

Answer: ACD

Explanation:

Referring to the exhibit for a multinode HA environment, we can conclude the following about the HA setup:

Two Services Redundancy Groups (Correct: Option A):

The output shows the status of SRG 0 and SRG 1, confirming that there are two services redundancy groups in the HA configuration.

Node 1 Hosting SRG 1 (Correct: Option C):

The exhibit indicates that Node 1 is currently active for SRG 1. According to the configuration, Node 1 will continue to host SRG 1 unless it becomes unavailable.

Session State Synchronization (Correct: Option D):

In this HA setup, session state synchronization is enabled between the two nodes. This ensures that sessions remain active and seamless failover can occur if one node fails.

Juniper Reference:

Juniper HA Documentation: Provides details on multinode HA setups, SRG configurations, and session synchronization.

Question: 60

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

Explanation:

Question: 61

A company has acquired a new branch office that has the same address space of one of its local networks, 192.168.100/24. 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: AD

Explanation:

The problem describes two offices needing to communicate, but both share the same IP address space, 192.168.100.0/24. To resolve this, NAT must be configured to translate the conflicting address spaces on each side. Here's how each of the configurations works:

Option A (Correct):

This source NAT rule translates the source address of traffic from Office B to Office A. By configuring source NAT, the source IP addresses from Office B (192.168.210.0/24) will be translated when communicating with Office A (192.168.200.0/24). This method ensures that there is no overlap in address space when packets are transmitted between the two offices.

Option D (Correct):

This is a source NAT rule configured on Office B, which translates the source addresses from Office A to prevent address conflicts. It ensures that when traffic is initiated from Office A to Office B, the overlapping address range (192.168.100.0/24) is translated.

Options B and C (Incorrect):

These options involve static NAT rules that map address ranges between the two offices, but they do not resolve the overlapping IP address space issue effectively. Static NAT is not the optimal solution in this scenario since the problem involves address space conflict, which requires translation of source addresses during communication.

Juniper Reference:

Juniper NAT Configuration Guide: Detailed instructions on how to configure source NAT and resolve address conflicts between networks.

Question: 62

Referring to the exhibit,

[edit security r.at] -serSstjc# stow source [interface { pert-overloading *ff;
} rule-set rule'. (from zone trust; to zone untrust; rule allow (match { source-address 172.16.1.0/24; destination-address 0.0.0.0/0

then (source-r.at { interface (persistent-nat { permit taxget-hos

1
I]

which two statements are correct about the NAT configuration? (Choose two.)

A. Both the internal and the external host can initiate a session after the initial translation.
B. Only a specific host can initiate a session to the reflexive address after the initial session.
C. Any external host will be able to initiate a session to the reflexive address.
D. The original destination port is used for the source port for the session.

Answer: AB

Explanation:

Question: 63

You are asked to establish a hub-and-spoke IPsec VPN using an SRX Series device as the hub. All of the spoke devices are third-party devices.

Which statement is correct in this scenario?

A. You must ensure that you are using aggressive mode when incorporating third-party devices as your spokes.
B. You must statically configure the next-hop tunnel binding table entries for each of the third-party spoke

devices.

C. You must create a policy-based VPN on the hub device when peering with third-party devices.

D. You must always peer using loopback addresses when using non-Junos devices as your spokes.

Answer: B

Explanation:

Question: 64

Exhibit:

```
[edit] userSSemoteSitel# snow interfaces ge-0/0/2 (
        unit 0 { family met { dhcp;
                I
        1

stO ( unit 0 { family inet {
                        address 10.0.0.2/30;

} [edit security zones] userSRatcteSitel# show security-zone entrust interfaces (
        ge-Q/0/2.0 ( host-inbound-traffic ( system-services ( ike; dhep;
                        J

J
```

[edit seeurity ike] userSRemoteSiteif show policy ike-policy-1 (node main; proposal-set standard; pre-shared-key ascii-text "S9S€stcCpohsex7viR7vwYzalAB"; ft SECRET-DATA

gateway gateway-1 { ike policy ike policy 1; address 203.0.113.5;
local-identity hostname "Renotesitel0srx.juniper.net"; external-interface ge-0/0/2;
I

[edit security ike] userScorporatef show policy ike-policy-sitel . mode main;
proposal-set standard;
pre-shared-key ascii-text ”S9$63t6CpOhSeX7VlR7VwYZGlAB”; #♦ SECRET-DATA

gateway gateway-site. (
ike-policy ike-policy-sitel;
dynamic hostname "Remotesitel0srx.juniper.net"; external-interface ge-0/0/1;

You are troubleshooting a new IPsec VPN that is configured between your corporate office and the RemoteSite1 SRX Series device. The VPN is not currently establishing. The RemoteSite1 device is being assigned an IP address on its gateway interface using DHCP.
Which action will solve this problem?

A. On the RemoteSite1 device, change the IKE gateway external interface to st0.0.
B. On both devices, change the IKE version to use version 2 only.
C. On both devices, change the IKE policy proposal set to basic.
D. On both devices, change the IKE policy mode to aggressive.

Answer: D

Explanation:

Aggressive mode is required when an IP address is dynamically assigned, such as through DHCP, as it allows for faster establishment with less identity verification. More details are available in Juniper IKE and IPsec Configuration Guide.

The configuration shown in the exhibit highlights that the RemoteSite1 SRX Series device is using DHCP to obtain an IP address for its external interface (ge-0/0/2). This introduces a challenge in IPsec VPN configurations when the public IP address of the remote site is not static, as is the case here. Aggressive mode in IKE (Internet Key Exchange) is designed for situations where one or both peers have dynamically assigned IP addresses. In this scenario, aggressive mode allows the devices to exchange identifying information, such as hostnames, rather than relying on static IP addresses, which is necessary when the remote peer (RemoteSite1) has a dynamic IP from DHCP.
Correct Action (D): Changing the IKE policy mode to aggressive will resolve the issue by allowing the two devices to establish the VPN even though one of them is using DHCP. In aggressive mode, the initiator can present its identity (hostname) during the initial handshake, enabling the VPN to be established successfully.
Incorrect Options:
Option A: Changing the external interface to st0.0 is incorrect because the st0 interface is used for the tunnel interface, not for the IKE negotiation.
Option B: Changing to IKE version 2 would not resolve the dynamic IP issue directly, and IKEv1 works in this scenario.
Option C: Changing the IKE proposal set to basic doesn't address the dynamic IP challenge in this scenario.
Juniper Reference:
Juniper IKE and VPN Documentation: Provides details on when to use aggressive mode, especially when a

dynamic IP address is involved.

Question: 65

You are asked to see if your persistent NAT binding table is exhausted.
Which show command would you use to accomplish this task?

A. show security nat source persistent-nat-table summary
B. show security nat source summary
C. show security nat source pool all
D. show security nat source persistent-nat-table all

Answer: D

Explanation:

The command show security nat source persistent-nat-table all provides a comprehensive view of all entries in the persistent NAT table, enabling administrators to monitor and manage resource exhaustion. Refer to Juniper NAT Monitoring Guide for more.

In Junos OS, when persistent NAT is configured, a binding table is created to keep track of NAT sessions and ensure that specific hosts are allowed to initiate sessions back to internal hosts. To check if the persistent NAT binding table is full or exhausted, the correct command must display the entire table.
Correct Command (D):
The command show security nat source persistent-nat-table all will display the entire persistent NAT binding table. This allows you to check whether the table is exhausted or if there is space available for new persistent NAT sessions.
Incorrect Options:
Option A: The command show security nat source persistent-nat-table summary provides a summary view but does not give detailed insights into whether the table is exhausted.
Option B and Option C: These commands deal with general NAT source summaries or pools, which are not related specifically to persistent NAT bindings.
Juniper Reference:
Juniper Persistent NAT Documentation: Describes the persistent NAT binding table and the commands used to monitor its status.

Question: 66

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: B, C

Explanation:

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

When two networks with overlapping IP address spaces need to communicate, Network Address Translation (NAT) is required to translate the IP addresses so that they become unique across the combined network. In this scenario, both the local network and the new branch office use the same subnet: 192.168.100.0/24. To enable communication without IP conflicts, we need to translate the overlapping addresses to unique ones.

Understanding the Problem:

Local Network (Office A): 192.168.100.0/24

Branch Office (Office B): 192.168.100.0/24

Objective: Allow communication between Office A and Office B despite overlapping IP ranges. Solution Overview:

To resolve the overlapping IP addresses, we can use Static NAT to create a one-to-one mapping between the overlapping IP addresses and a unique IP range. This way, when packets traverse the network boundary, their IP addresses are translated to a non-overlapping range, avoiding conflicts. Option B and Option C implement Static NAT to resolve the issue: Option B (At Office A):

Translates destination addresses from 192.168.200.0/24 to 192.168.100.0/24.

This allows Office B to reach Office A's overlapping network by targeting a unique IP range (192.168.200.0/24).

Option C (At Office B):

Translates destination addresses from 192.168.210.0/24 to 192.168.100.0/24.

This allows Office A to reach Office B's overlapping network by targeting a unique IP range (192.168.210.0/24).

Detailed

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; }
            }
        }
    }
}
```

from interface ge-0/0/0.0;: Specifies the interface through which the traffic is received.

Matching Traffic:

destination-address 192.168.200.0/24;: Matches packets destined for 192.168.200.0/24.

Action:

static-nat { prefix { 192.168.100.0/24; } }: Translates the destination address to 192.168.100.0/24.

Result:

Office B sends packets to 192.168.200.0/24, which are translated to 192.168.100.0/24 upon arrival at Office A.

Reference:

Juniper Networks Documentation: "Configuring Static NAT"

2. Static NAT Configuration at Office B (Option C):

Configuration:

```
[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; }
            }
        }
    }
}
```

from interface ge-0/0/0.0;: Specifies the interface through which the traffic is received.

Matching Traffic:

destination-address 192.168.210.0/24;: Matches packets destined for 192.168.210.0/24.

Action:

static-nat { prefix { 192.168.100.0/24; } }: Translates the destination address to 192.168.100.0/24. Result:

Office A sends packets to 192.168.210.0/24, which are translated to 192.168.100.0/24 upon arrival at Office B.

Reference:

Juniper Networks Documentation: "Configuring Static NAT"

Why Options A and D are Incorrect:

Option A and Option D use Source NAT, which is typically used for translating the source IP address of outgoing traffic.

Source NAT with interface-based translation may not resolve overlapping IP issues effectively because it doesn't provide a one-to-one mapping of the overlapping addresses.

In scenarios with overlapping networks, Static NAT is preferred as it allows for consistent and predictable address translation, essential for two-way communication.

Key Juniper Concepts:

Static NAT:

Provides a one-to-one mapping between local and global addresses.

Useful for scenarios where bidirectional communication is required.

Reference: Juniper Networks Day One Book "Advanced NAT Concepts"

Source NAT:

Typically used for translating private IP addresses to public IP addresses for outbound traffic. Interface-based Source NAT translates the source IP to the IP address of the egress interface. Not ideal for resolving overlapping IP spaces in bidirectional communication.

Additional Reference:

Juniper TechLibrary:
"Understanding NAT in SRX Series Devices"
"Configuring NAT for Overlapping Networks"
Juniper Forums and Knowledge Base Articles:
Discussions on resolving overlapping IP address spaces using Static NAT.

Conclusion:

By implementing Static NAT configurations as shown in Options B and C, both offices can effectively communicate despite having overlapping IP address spaces. Static NAT ensures that IP addresses are uniquely translated, avoiding conflicts and enabling seamless connectivity between the two networks.

Question: 67

Click the Exhibit button.

```
[edit class-of-service]
user@srx# show
        cp ba-classifler (
            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 Reference

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

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.

Explanation:

Additional Reference:

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.

Question: 68

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: A, D

Explanation:

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

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.

To bypass IDP for specific traffic using APBR, you must explicitly configure the idp-disable option within the APBR rule.

Without this configuration, even if APBR redirects the traffic, IDP will still inspect and potentially drop the traffic.

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.

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.

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.

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 Reference:

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.

Question: 69

Which two statements are true regarding NAT64? (Choose two.)

A. An SRX Series device should be in flow-based forwarding mode for IPv4.
B. An SRX Series device should be in packet-based forwarding mode for IPv4.
C. An SRX Series device should be in packet-based forwarding mode for IPv6.
D. An SRX Series device should be in flow-based forwarding mode for IPv6.

Answer: A, D

Explanation:

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

Understanding NAT64:

NAT64 allows IPv6-only clients to communicate with IPv4 servers by translating IPv6 addresses to IPv4 addresses and vice versa.

It is essential in environments where IPv6 clients need access to IPv4 resources.

Flow-Based vs. Packet-Based Forwarding Modes:

Flow-Based Forwarding Mode:

The SRX device processes packets based on the session state.

Supports advanced services like NAT, IDP, and ALG.

Packet-Based Forwarding Mode:

The SRX device processes each packet individually without maintaining session state.

Limited support for advanced services.

Option A: An SRX Series device should be in flow-based forwarding mode for IPv4.

True.

NAT64 requires flow-based mode for IPv4 traffic to properly translate and maintain session states.

Option B: An SRX Series device should be in packet-based forwarding mode for IPv4.

False.

Packet-based mode does not support NAT features.

Option C: An SRX Series device should be in packet-based forwarding mode for IPv6.

False.

Similar to IPv4, NAT64 requires flow-based mode for IPv6 traffic.

Option D: An SRX Series device should be in flow-based forwarding mode for IPv6.

True.

Flow-based mode is necessary for NAT64 to handle IPv6 traffic correctly.

Key Points:

NAT64 Requires Flow-Based Mode:

Both IPv4 and IPv6 interfaces involved in NAT64 must be configured in flow-based mode.

This is because NAT64 relies on session information and stateful packet inspection.

Packet-Based Mode Limitations:

Does not support NAT, as it lacks session awareness.

Not suitable for NAT64 operations.

Juniper Security Reference:

Juniper Networks Documentation:

"NAT64 is supported only in flow-based processing mode."

Source: Configuring NAT64

Understanding Flow-Based and Packet-Based Modes:

"Flow-based mode is required for stateful services such as NAT."

Source: Flow-Based and Packet-Based Processing

Conclusion:

To implement NAT64 on an SRX Series device, both IPv4 and IPv6 traffic must be processed in flowbased forwarding mode.

Therefore, Options A and D are the correct statements.

Question: 70

Click the Exhibit button.

```
user0srx> show ethernet   -switching gl obal-information
Global Configuration:

MAC aging interval      :
MAC learning            : 300 Enabled

MAC statistics          : Disabled
MAC limit Count         : 65536

MAC limit hit           : Disabled
MAC packet action  drop:  Disabled
```

```
MA '+IP aging interval :                    seconds

                                            seconds
MAC*IP   limit   Count   :IPv4 - 1200
MA'+IP limit reached : LEIPv6 - 1200
aging time : LE VLAN aging65536
time :                    No 1200 1200
Global Mode              : Transparent   bridge
RE state                 : Master
```

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

A. You cannot secure intra-VLAN traffic with a security policy on this device.
B. You can secure inter-VLAN traffic with a security policy on this device.
C. The device can pass Layer 2 and Layer 3 traffic at the same time.
D. The device cannot pass Layer 2 and Layer 3 traffic at the same time.

Answer: A, D

Explanation:

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

Understanding the Exhibit:

The SRX device is operating in Transparent Mode, as indicated by:
Global Mode : Transparent bridge
Transparent Mode on SRX Devices:
Transparent Mode (Layer 2 Mode):
The SRX device acts as a Layer 2 switch.
Does not perform routing functions.
Security policies can be applied to inter-VLAN (Layer 2) traffic but not intra-VLAN traffic.
Cannot handle Layer 3 traffic simultaneously.

Option A: You cannot secure intra-VLAN traffic with a security policy on this device.
True.
In Transparent Mode, intra-VLAN traffic is switched within the VLAN and does not pass through the SRX firewall processing engine.
Therefore, security policies cannot be applied to intra-VLAN traffic.
Option B: You can secure inter-VLAN traffic with a security policy on this device.
False.
In Transparent Mode, all interfaces are in the same VLAN (unless VLAN tagging is configured).
Inter-VLAN routing is not possible as the device does not perform Layer 3 functions.
Option C: The device can pass Layer 2 and Layer 3 traffic at the same time.
False.
In Transparent Mode, the SRX device operates exclusively at Layer 2.
It cannot process Layer 3 traffic simultaneously.
Option D: The device cannot pass Layer 2 and Layer 3 traffic at the same time.

True.

The SRX device in Transparent Mode cannot handle both Layer 2 and Layer 3 traffic concurrently.

Key Points:

Intra-VLAN Traffic:

Traffic within the same VLAN.

In Transparent Mode, this traffic is switched and does not go through the firewall's security policies. Inter-VLAN Traffic:

Traffic between different VLANs.

Requires routing capabilities (Layer 3).

In Transparent Mode, the SRX cannot perform routing functions.

Juniper Security Reference:

Juniper Networks Documentation:

"In transparent mode, the SRX Series device acts like a Layer 2 switch or bridge. Security policies cannot control intra-VLAN traffic because such traffic does not pass through the firewall." Source: Understanding Transparent Mode

"The device cannot perform both Layer 2 switching and Layer 3 routing simultaneously in transparent mode." Source: Transparent Mode Limitations

Conclusion:

Option A is correct because intra-VLAN traffic cannot be secured with security policies in Transparent Mode.

Option D is correct because the device cannot pass both Layer 2 and Layer 3 traffic at the same time when operating in Transparent Mode.

Question: 71

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: A, D

Explanation:

Question: 72

You are deploying OSPF over IPsec with an SRX Series device and third-party device using GRE. Which two statements are correct? (Choose two.)

A. The GRE interface should use lo0 as endpoints.
B. The OSPF protocol must be enabled under the VPN zone.
C. Overlapping addresses are allowed between remote networks.
D. The GRE interface must be configured under the OSPF protocol.

Answer: A, D

Explanation:

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

Understanding the Scenario:

Objective: Deploy OSPF over IPsec between an SRX Series device and a third-party device using GRE tunnels.

Components Involved:

GRE (Generic Routing Encapsulation): Encapsulates packets to allow routing protocols like OSPF to run over IPsec tunnels.

IPsec: Provides security for the GRE tunnels.

OSPF: Dynamic routing protocol used over the GRE tunnel.

Option A: The GRE interface should use lo0 as endpoints.

Using the loopback interface (lo0) as the source and destination endpoints for GRE tunnels is a common best practice.

Advantages:

Stability: Loopback interfaces are always up, ensuring the GRE tunnel remains operational even if physical interfaces fail.

Reachability: Provides consistent endpoint IP addresses for GRE tunnels.

Configuration:

Assign IP addresses to lo0 interfaces on both devices.

Configure GRE tunnels to use these lo0 IP addresses as source and destination.

Reference:

Juniper Networks Documentation:

"Using loopback interfaces as GRE tunnel endpoints ensures stability and consistent reachability for routing protocols over GRE tunnels."

Source: Configuring GRE Tunnels

Option D: The GRE interface must be configured under the OSPF protocol.

To run OSPF over the GRE tunnel, the GRE interface must be included in the OSPF configuration.

Configuration Steps:

Create GRE Interface:

Example: set interfaces gr-0/0/0 unit 0 tunnel source <source-ip> tunnel destination <destination-ip>

Assign IP Address to GRE Interface:

Example: set interfaces gr-0/0/0 unit 0 family inet address <ip-address>

Include GRE Interface in OSPF:

Example: set protocols ospf area <area-id> interface gr-0/0/0.0

Result:

OSPF will establish adjacencies over the GRE interface and exchange routing information.

Reference:

Juniper Networks Documentation:

"To enable OSPF over GRE tunnels, you must include the GRE interfaces in the OSPF configuration."

Source: OSPF over GRE Configuration

Why Options B and C are Incorrect:

Option B: The OSPF protocol must be enabled under the VPN zone.

Since OSPF is running over the GRE tunnel, which is encapsulated over IPsec, the OSPF packets are encapsulated within GRE and IPsec.

The SRX device does not need to allow OSPF in the security policies or enable OSPF under the VPN zone for GRE-encapsulated traffic.

Security Policies:

The GRE traffic (IP protocol 47) must be permitted through the security policies.

OSPF runs inside the GRE tunnel and does not require additional configuration under the VPN zone.

Reference:

Juniper Networks Documentation:

"When using GRE over IPsec, routing protocols run over GRE and do not require separate security policies for their control traffic."

Source: Security Policies for GRE over IPsec

Option C: Overlapping addresses are allowed between remote networks.

Overlapping IP addresses can cause routing conflicts and are generally not recommended.

In a GRE over IPsec scenario, overlapping addresses can lead to issues in routing protocol adjacency and data forwarding.

Best Practice:

Ensure unique IP addressing schemes between remote networks to prevent routing issues.

Reference:

Juniper Networks Documentation:

"Overlapping IP address spaces can lead to routing ambiguities and should be avoided when configuring GRE tunnels."

Source: Avoiding Overlapping IP Addresses

Conclusion:

Correct Answers: A and D

Rationale:

Option A is correct because using lo0 as endpoints for GRE provides stability and reliability.

Option D is correct because the GRE interface must be included in the OSPF configuration to enable OSPF over the tunnel.

Question: 73

You are asked to set up advanced policy-based routing.
Which type of routing instance is designed to support this scenario?

A. forwarding
B. virtual switch
C. virtual router
D. non-forwarding

Answer: A

Explanation:

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

Understanding Advanced Policy-Based Routing (APBR):

APBR: Allows routing decisions based on application-level information and policies.

Objective: Direct specific application traffic through different paths based on policies.

Routing Instances in Junos OS:

Forwarding Instance:

Used for features like filter-based forwarding (FBF) and APBR.

Provides a separate forwarding table but shares the global routing table.

Supports APBR.

Virtual Router:

Provides a separate routing table and forwarding table.

Used for logical separation of routing domains.

Does not support APBR directly.

Virtual Switch:

Operates at Layer 2.

Used for VLAN separation and Layer 2 switching.

Not applicable to routing or APBR.

Non-Forwarding Instance:

Used for management purposes.

Does not forward transit traffic.

Not suitable for APBR.

Option A: forwarding

Correct.

A forwarding routing instance is specifically designed to support advanced policy-based routing.

It allows the SRX device to direct traffic based on policies to different forwarding instances. Reference:

Juniper Networks Documentation:

"To configure advanced policy-based routing, you must create a forwarding-type routing instance."

Source: Configuring Advanced Policy-Based Routing

Why Other Options Are Incorrect:

Option B: virtual switch

Incorrect.

Virtual switch instances are for Layer 2 switching and VLAN separation.

They do not support routing or APBR.

Option C: virtual router Incorrect.

Virtual router instances are used for isolating routing tables.

While they support routing, they are not designed for APBR.

Option D: non-forwarding Incorrect.

Non-forwarding instances do not handle transit traffic.

They are used for management routing tables and cannot be used for APBR. Conclusion:

Correct Answer: A. forwarding

Explanation:

Rationale:

A forwarding routing instance is the appropriate type to support advanced policy-based routing.

Question: 74

Click the Exhibit button.

user0srx2> show chassis high-availability services-redundancy-group 1 SRG failure event codes: BF BFD monitoring IP IP monitoring IF Interface monitoring CP Control Plane monitoring

Services Redundancy Group: 1
Deployment Type: SWITCHING Status:
BACKUP Activeness Priority: 100 Preemption:
DISABLED Process Packet In Backup State: NO
Control Plane State: READY System Integrity
Check: COMPLETE Failure Events: NONE Peer
Information: Peer Id: 1 Status : ACTIVE Health
Status: HEALTHY Failover Readiness: N/A
Virtual IP Info: Index: 2

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

A. This device is the backup node for SRG1.
B. The ge-0/0/3.0 and ge-0/0/4.0 interfaces are not active and will not respond to ARP requests to the virtual IP MAC address.
C. This device is the active node for SRG1.
D. The ge-0/0/3.0 and ge-0/0/4.0 interfaces are active and will respond to ARP requests to the virtual IP MAC address.

Answer: CD

Explanation:

Question: 75

You have a multinode HA default mode deployment and the ICL is down.

In this scenario, what are two ways that the SRX Series devices verify the activeness of their peers? (Choose two.)

A. Custom IP addresses may be configured for the activeness probe.
B. Fabric link heartbeats are used to verify the activeness of the peers.
C. Each peer sends a probe with the virtual IP address as the destination IP address.
D. Each peer sends a probe with the virtual IP address as the source IP address and the upstream router as the destination IP address.

Answer: A, D

Explanation:

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

Understanding the Scenario:

Multinode HA Default Mode Deployment:

In a chassis cluster, two SRX devices operate together to provide high availability.

ICL (Inter-Cluster Link) is Down:

The control and fabric links between the nodes are not operational.

Objective:

Determine how the SRX devices verify each other's activeness without the ICL.

Option A: Custom IP addresses may be configured for the activeness probe.

When the control link is down, SRX devices use an ICMP ping-based activeness probe to check the peer's status.
Custom IP addresses can be configured as probe targets to verify the peer's activeness.

Reference:

"You can configure the SRX Series device to send activeness probes to a configured IP address to verify the peer's state when the control link is down."

Source: Juniper Networks Documentation - Control Link Failure Detection

Option D: Each peer sends a probe with the virtual IP address as the source IP address and the upstream router as the destination IP address.

The SRX devices send ICMP probes to an upstream device using the redundancy group's virtual IP address as the source.

This helps determine if the peer node is still active by verifying network reachability.

Reference:

"When the control link fails, each node sends ICMP pings to the configured probe addresses using the redundancy group's virtual IP address as the source."

Source: Juniper Networks Documentation - Chassis Cluster Control Link Failure

Why Options B and C are Incorrect:

Option B: Fabric link heartbeats cannot be used because the ICL (which includes the fabric link) is down.

Option C: Probes are sent to upstream devices, not using the virtual IP address as the destination. Conclusion:

The correct options are A and D because they accurately describe how SRX devices verify activeness without the ICL.

Question: 76

Click the Exhibit button.

```
user@SRX>show security flow session
Session ID: ICO, Policy name: Ll-to-L9/ll, Timeout: 36, Session State: Valid
    In: 10.10.101.10/1 —> 10.10.102.10/1;icmp, Conn Tag: 0x0, If: ge-0/0/4.0, Pkts: 1, Bytes: 84,
    Out: 10.10.102.10/1 —> 10.10.101.10/1,‘icmp, Conn Tag: 0x0, If: ge-0/0/5.0, Pkts: 0, Bytes: 0,
```

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

A. The traffic is permitted.
B. The traffic was initiated by the 10.10.102.10 address.
C. The destination device is not responding.
D. The traffic is denied.

Answer: A, C

Explanation:

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

Understanding the Session Output:

Session State: Valid

Indicates that the session is active and permitted by security policies.

Policy Name: L1-to-L9/11

Shows the policy that allowed the session.

In Direction:

Source: 10.10.101.10

Destination: 10.10.102.10

Packets: 1

Bytes: 84

Out Direction:

Packets: 0

Bytes: 0

Indicates no return traffic.

Option A: The traffic is permitted.

The session state is Valid, and a policy name is specified.

This means the SRX device allowed the traffic.

Reference:

"A session with a Valid state and an associated policy name indicates permitted traffic."

Source: Juniper TechLibrary - Understanding Security Flow Sessions

Option C: The destination device is not responding.

The lack of packets in the Out direction suggests that the destination (10.10.102.10) is not responding.

Reference:

"If there are no packets in the reverse direction, it may indicate that the destination host is not responding."

Source: Juniper KB - Troubleshooting Traffic Flows

Why Options B and D are Incorrect:

Option B: The traffic was initiated by 10.10.101.10, not 10.10.102.10.

Option D: The session is valid and permitted; the traffic is not denied.

Conclusion:

The correct options are A and C because they accurately describe the state of the session.

Question: 77

You are setting up multinode HA for redundancy.

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

A. Dynamic routing is active on one device at a time.
B. Dynamic routing is active on both devices.
C. Physical connections are used for the control and fabric links.
D. ICL links require Layer 3 connectivity between peers.

Answer: A, C

Explanation:

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

Understanding Multinode HA:
Chassis Cluster in Active/Passive Mode:
One node is active, and the other is standby.
Dynamic Routing Protocols:
Run on the active node only.
Option A: Dynamic routing is active on one device at a time.

In active/passive HA, dynamic routing protocols run only on the primary (active) node.
Reference:
"In a chassis cluster, the primary node handles all control plane tasks, including dynamic routing."
Source: Juniper TechLibrary - Chassis Cluster Overview
Option C: Physical connections are used for the control and fabric links.

Control and fabric links are direct physical connections between cluster nodes.
Reference:
"The control and fabric links must be connected using physical interfaces between the nodes."
Source: Juniper TechLibrary - Chassis Cluster Components
Why Options B and D are Incorrect:
Option B: Dynamic routing is not active on both devices simultaneously in active/passive mode.
Option D: The Inter-Cluster Link (ICL) uses Layer 2 connectivity, not Layer 3.
Conclusion:
The correct options are A and C.

Question: 78

You want to configure the SRX Series device to map two peer interfaces together and ensure that there is no switching or routing lookup to forward traffic.
Which feature on the SRX Series device is used to accomplish this task?

A. Transparent mode
B. Secure wire
C. Mixed mode
D. Switching mode

Answer: B

Explanation:

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

Understanding Secure Wire:
Secure Wire Feature:
Connects two interfaces directly without any Layer 2 or Layer 3 processing.
No routing or switching lookup occurs.
Use Case:
Ideal for scenarios where traffic needs to pass through the SRX device transparently.
Option B: Secure wire

Secure wire creates a bidirectional link between two interfaces.
Traffic flows between the interfaces as if they are connected by a physical wire.

Reference:

"The secure wire feature allows traffic to pass between two interfaces without any security processing or route lookups."

Source: Juniper TechLibrary - Secure Wire Overview

Why Other Options Are Incorrect:

Option A: Transparent mode involves Layer 2 switching.

Option C: Mixed mode combines Layer 2 and Layer 3 but doesn't prevent switching/routing lookups.

Option D: Switching mode operates at Layer 2 and includes switching lookups.

Conclusion:

Secure wire is the correct feature to map two interfaces together without switching or routing lookups.

Question: 79

Click the Exhibit button.

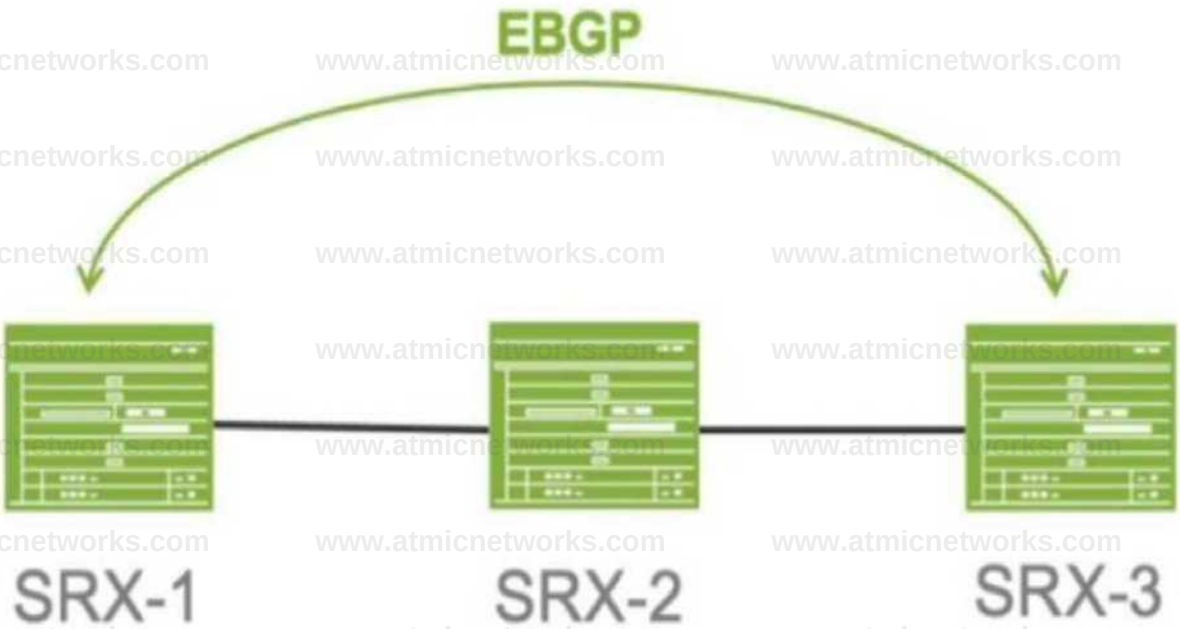

Referring to the exhibit. SRX-1 and SRX-3 have to be connected using EBGP. The BGP configuration on SRX-1 and SRX-3 is verified and correct.

Which configuration on SRX-2 would establish an EBGP connection successfully between SRX-1 and SRX-3?

A. The host-inbound-traffic statements do not allow EBGP traffic to traverse SRX-2.
B. The security policy to allow SRX-1 and SRX-3 to communicate on TCP port 79 should be configured.
C. The security policy to allow SRX-1 and SRX-3 to communicate on TCP port 169 should be configured.
D. The security policy to allow SRX-1 and SRX-3 to communicate on TCP port 179 should be configured.

Answer: D

Explanation:

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

Understanding the Scenario:

SRX-1 and SRX-3:
Need to establish an EBGP session through SRX-2.
Issue:
BGP session is not coming up despite correct configurations on SRX-1 and SRX-3.
Option D: The security policy to allow SRX-1 and SRX-3 to communicate on TCP port 179 should be configured.

BGP uses TCP port 179 for establishing sessions.
SRX-2 must have a security policy allowing traffic between SRX-1 and SRX-3 on TCP port 179.
Reference:
"Security policies must permit BGP traffic (TCP port 179) to allow BGP sessions through the SRX device."
Source: Juniper TechLibrary - Configuring Security Policies for Transit Traffic
Why Other Options Are Incorrect:
Option A: Host-inbound-traffic affects traffic destined to SRX-2, not transit traffic.
Option B and C: TCP ports 79 and 169 are unrelated to BGP.
Conclusion:
The correct option is D, configuring a security policy to allow TCP port 179.

Question: 80

You are attempting to ping an interface on your SRX Series device, but the ping is unsuccessful.
What are three reasons for this behavior? (Choose three.)

A. The interface is not assigned to a security zone.

B. The interface's host-inbound-traffic security zone configuration does not permit ping

C. The ping traffic is matching a firewall filter.

D. The device has J-Web enabled.

E. The interface has multiple logical units configured.

Answer: ABC

Explanation:

A. The interface is not assigned to a security zone.
SRX Series devices rely heavily on security zones for traffic management. If an interface isn't assigned to a zone, the device won't know how to handle traffic arriving on that interface, including ping requests (ICMP echo requests).
Reference: Security Zones, Zone Properties, and Global Properties [invalid URL removed]
B. The interface's host-inbound-traffic security zone configuration does not permit ping.
Even if an interface is in a zone, you must explicitly allow ICMP ping traffic within the zone's host-inbound-traffic settings. By default, most zones block ping for security reasons.

Reference: Configuring Host-Inbound Traffic [invalid URL removed]

C. The ping traffic is matching a firewall filter.

Firewall filters (configured using the security policies hierarchy) can block specific traffic types, including ICMP. If a filter is applied to the interface or zone, and it doesn't have a rule to permit ping, the ping will be unsuccessful.

Reference: Firewall Filters [invalid URL removed]

Why other options are incorrect:

D. The device has J-Web enabled. J-Web is a web-based management interface and has no direct impact on the device's ability to respond to pings.

E. The interface has multiple logical units configured. Logical units divide a physical interface into multiple virtual interfaces. While this can affect routing and traffic flow, it doesn't inherently prevent ping responses as long as the relevant zones and policies are correctly configured.

Troubleshooting Steps:

If you're unable to ping an SRX interface, here's a systematic approach to troubleshoot:

Verify Interface Status: Ensure the interface is up and operational using show interfaces terse.

Check Zone Assignment: Confirm the interface belongs to a security zone using show security zones. Examine host-inbound-traffic: Verify that the zone's host-inbound-traffic settings allow ping (e.g., set security zones security-zone trust host-inbound-traffic system-services ping).

Analyze Firewall Filters: Review any firewall filters applied to the interface or zone to ensure they allow ICMP ping traffic. Use show security policies and monitor traffic to diagnose filter behavior. Test from Different Zones: Try pinging the interface from devices in different zones to isolate potential policy issues.

By systematically checking these aspects, you can identify the root cause and resolve the ping issue on your SRX Series device.

Question: 81

You are deploying IPsec VPNs to securely connect several enterprise sites with ospf for dynamic routing. Some of these sites are secured by third-party devices not running Junos.

Which two statements are true for this deployment? (Choose two.)

A. OSPF over IPsec can be used for intersite dynamic routing.

B. Sites with overlapping address spaces can be supported.

C. OSPF over GRE over IPsec is required to enable intersite dynamic routing

D. Sites with overlapping address spaces cannot be supported.

Answer: BC

Explanation:

Understanding the Scenario:

Objective: Deploy IPsec VPNs connecting multiple enterprise sites using OSPF for dynamic routing.

Challenge: Some sites use third-party devices not running Junos OS.

Considerations:

Compatibility between Juniper and third-party devices.

Support for dynamic routing protocols (OSPF) over IPsec VPNs.

Handling overlapping IP address spaces.

Option Analysis:

Option A: OSPF over IPsec can be used for intersite dynamic routing.

OSPF Characteristics:

OSPF uses multicast addresses (224.0.0.5 and 224.0.0.6) for neighbor discovery and routing updates. IPsec Limitations:

Standard IPsec tunnel mode does not support multicast traffic natively.

Multicast traffic cannot traverse IPsec tunnels unless encapsulated.

Juniper Solution:

Juniper devices can use routed VPNs (route-based VPNs) with st0 interfaces, allowing OSPF over IPsec.

However, this requires support from both ends of the VPN tunnel.

Third-Party Devices:

May not support OSPF over IPsec without additional configurations.

Conclusion:

Option A is not universally true in this scenario due to third-party device limitations.

Reference:

"OSPF can be run over IPsec VPNs using route-based VPNs, but interoperability with third-party devices must be verified."

Source: Juniper TechLibrary - OSPF over IPsec VPNs

Option B: Sites with overlapping address spaces can be supported.

Overlapping IP Address Spaces:

Occurs when different sites use the same IP subnets.

Can cause routing ambiguities and conflicts.

Solution:

NAT over VPN:

Use Network Address Translation (NAT) to translate overlapping IP addresses to unique addresses.

Juniper devices support NAT over IPsec VPNs.

Third-Party Device Considerations:

Need to ensure third-party devices support NAT over IPsec.

Many enterprise-grade devices provide this functionality.

Conclusion:

Option B is true; overlapping address spaces can be supported using NAT.

Reference:

"When sites have overlapping IP addresses, NAT can be used over IPsec VPNs to resolve address conflicts."

Source: Juniper TechLibrary - NAT with IPsec VPNs

Option C: OSPF over GRE over IPsec is required to enable intersite dynamic routing.

GRE Tunnels:

Generic Routing Encapsulation (GRE) can encapsulate multicast and broadcast traffic.

Allows OSPF packets to be transmitted over IPsec VPNs.

IPsec Encryption:

GRE tunnels can be encrypted using IPsec for secure communication.

Interoperability:

GRE over IPsec is a common method to support OSPF between devices from different vendors.

Third-party devices are more likely to support GRE over IPsec than OSPF over IPsec directly.

Conclusion:

Option C is true; using OSPF over GRE over IPsec is required in this scenario.

Reference:

"To run OSPF between devices that do not support multicast over IPsec, GRE tunnels can be used over IPsec VPNs."

Source: Juniper TechLibrary - Configuring GRE over IPsec

Option D: Sites with overlapping address spaces cannot be supported.

Contradicts Option B.

As established, overlapping address spaces can be supported using NAT over IPsec VPNs. Conclusion:

Option D is false.

Conclusion:

Correct Answers: B and C

Option B: Overlapping address spaces can be supported using NAT over IPsec VPNs.

Option C: OSPF over GRE over IPsec is required to enable intersite dynamic routing, especially when third-party devices are involved.

Additional Detailed

Why OSPF over IPsec May Not Be Feasible (Option A):

Multicast Traffic:

OSPF relies on multicast for neighbor discovery and updates.

IPsec in tunnel mode does not natively support multicast traffic.

Third-Party Devices:

May not support proprietary extensions or configurations required to run OSPF directly over IPsec.

Workaround:

Encapsulate OSPF multicast packets within GRE tunnels, which can carry multicast traffic over unicast IPsec tunnels.

Why OSPF over GRE over IPsec Is Necessary (Option C):

GRE Tunnels:

Encapsulate multicast/broadcast traffic into unicast packets.

Allow routing protocols like OSPF to function over IPsec VPNs.

Compatibility:

GRE is a widely supported protocol across different vendors.

Facilitates interoperability between Juniper and third-party devices.

Supporting Overlapping Address Spaces (Option B):

NAT over IPsec:

Translates private IP addresses to unique addresses across the VPN.

Prevents routing conflicts and allows communication between sites with overlapping subnets. Considerations:

Requires proper configuration on both ends of the VPN tunnel.

Third-party devices must support NAT over IPsec.

Reference to Juniper Security Concepts:

Route-Based VPNs:

"Route-based VPNs use virtual tunnel interfaces (st0) and support dynamic routing protocols over IPsec."
Source: Juniper TechLibrary - Route-Based VPNs
GRE over IPsec:
"GRE over IPsec allows the transmission of multicast and non-IP protocols over IPsec tunnels."
Source: Juniper TechLibrary - GRE over IPsec Overview
NAT with IPsec VPNs:
"NAT can be applied to IPsec VPN traffic to resolve overlapping address issues and facilitate communication between sites."
Source: Juniper TechLibrary - NAT with IPsec
Final Notes:
Interoperability:
When working with third-party devices, always verify compatibility for protocols and features. Best Practices:
Use GRE over IPsec for dynamic routing protocols requiring multicast support across IPsec VPNs.
Implement NAT over VPN when dealing with overlapping address spaces.

Question: 82

Exhibit:

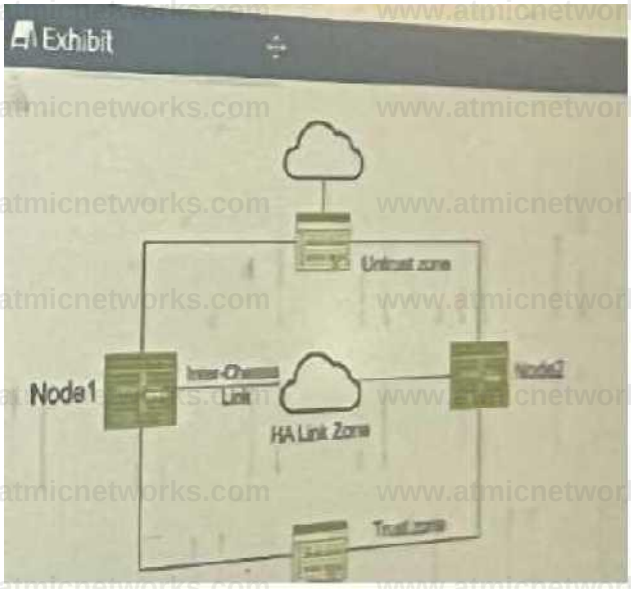

You have deployed a pair of SRX series devices in a multimode HA environment. You need to enable IPsec encryption on the interchassis link.
Referring to the exhibit, which three steps are required to enable ICL encryption? (Choose three.)

A. Install the Junos IKE package on both nodes.

B. Enable OSPF for both interchassis link interfaces and tum 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.

Question: 83

Exhibit:

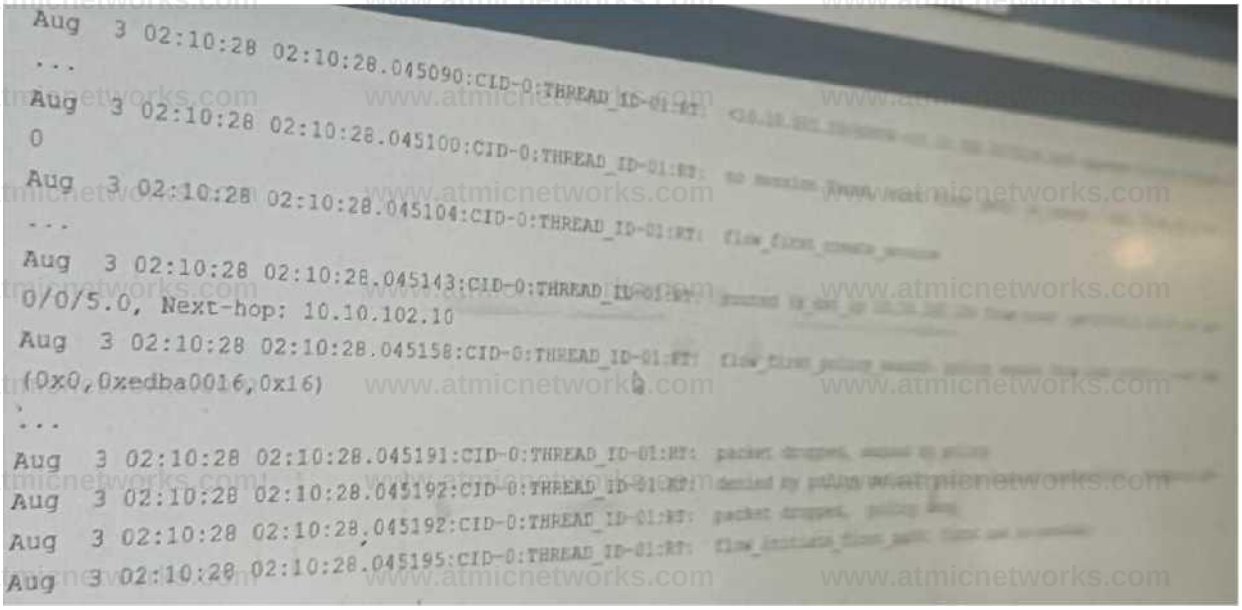

Which two statements are correct about the output shown in the exhibit. (Choose Two)

A. The data shown requires a traceoptions flag of basic-datapath.

B. The data shown requires a traceoptions flag of host-traffic.

C. The packet is dropped by the default security policy.

D. The packet is dropped by a configured security policy.

Answer: AC

Explanation:

Question: 84

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: Applications: Match based on specific applications or application groups.
URL categories: Match based on URL categories provided by a web filtering service.
Other criteria: You can also match based on source/destination IP addresses, ports, protocols, etc.
C. Then action: This specifies the action to take when traffic matches the rule. The primary action in APBR is: routing-instance: This redirects the matching traffic to a specific routing instance, allowing you to steer traffic through different paths based on the application or URL category.
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.

Question: 85

Referring to the exhibit, you are attempting to set up a remote access VPN on your SRX series devices.

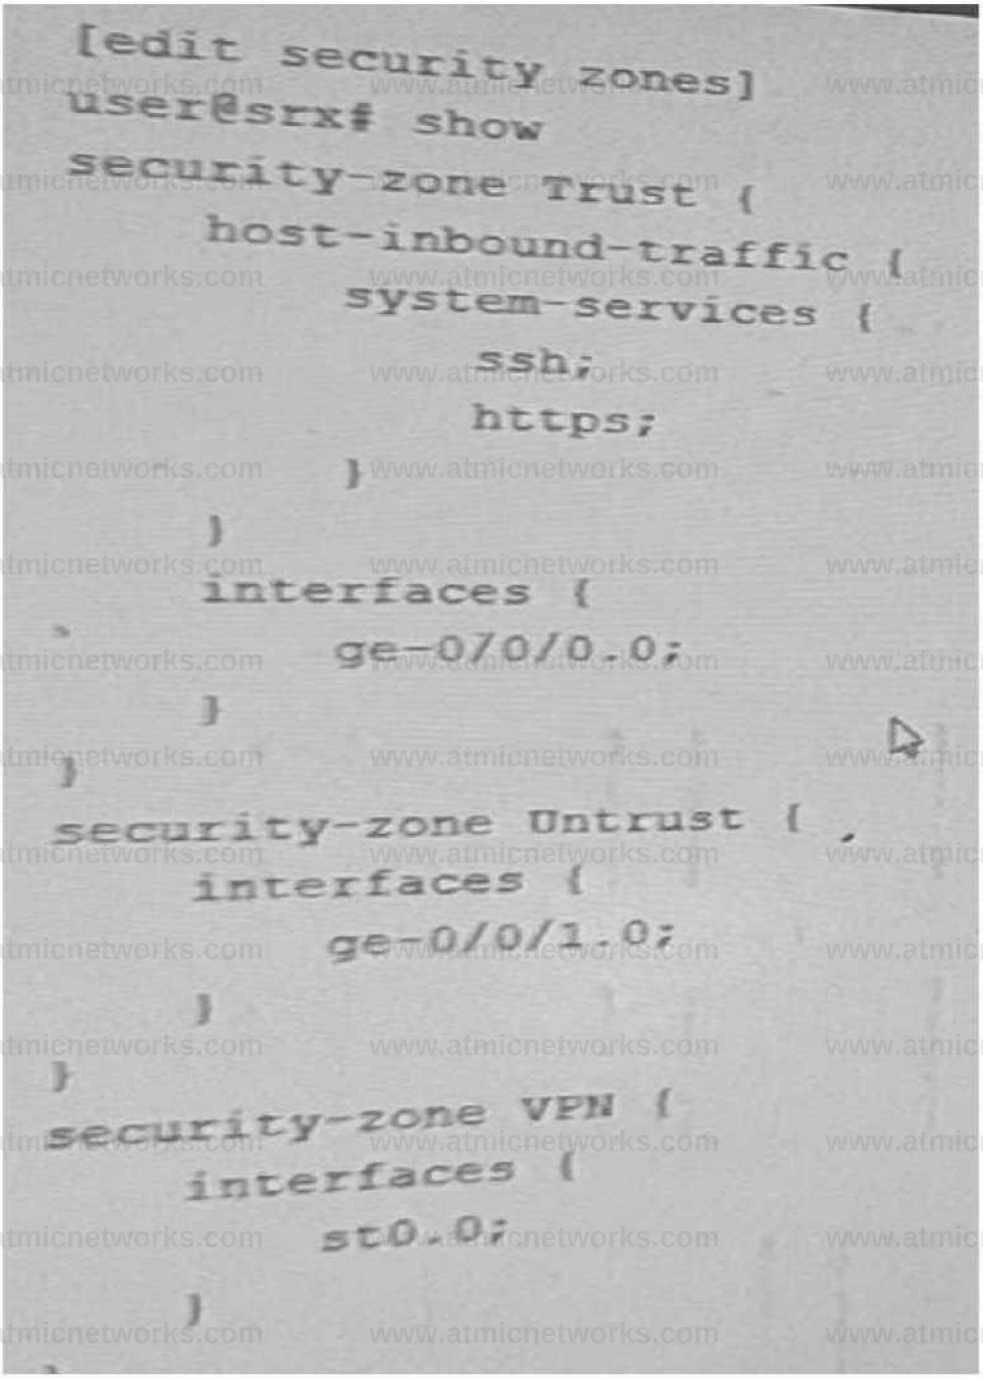

```
[edit security zones]
user@srx# show
security-zone Trust {
    host-inbound-traffic {
        system-services {
            ssh;
            https;
        }
    }
    interfaces {
        ge-0/0/0.0;
    }
}
security-zone Untrust {
    interfaces {
        ge-0/0/1.0;
    }
}
security-zone VPN {
    interfaces {
        st0.0;
    }
```

However you are unsure of which system services you should allow and in which zones they should be allowed to correctly finish the remote access VPN configuration
Which two statements are correct? (Choose two.)

A. You should add the host-inbound-traffic system-service ike statement to the Untrust zone.

B. You should add the host-inbound-traffic system-service ike statement to the VPN zone.

C. You should add the host-inbound-traffic system-service tcp-encap statement to the Untrust zone

D. You should add the host-inbound-traffic system-service tcp-encap statement to the VPN zone

Answer: AC

Explanation:

Question: 86

What are three configurable monitor components for a service redundancy group? (Choose two)

A. Interface

B. BFD

C. hardware alarm

D. IP

D. ARP

Answer: ADE

Explanation:

Question: 87

The SRX series device is performing static NAT. you want to ensure that host A can reach the internal webserver www.juniper.net using domain name.

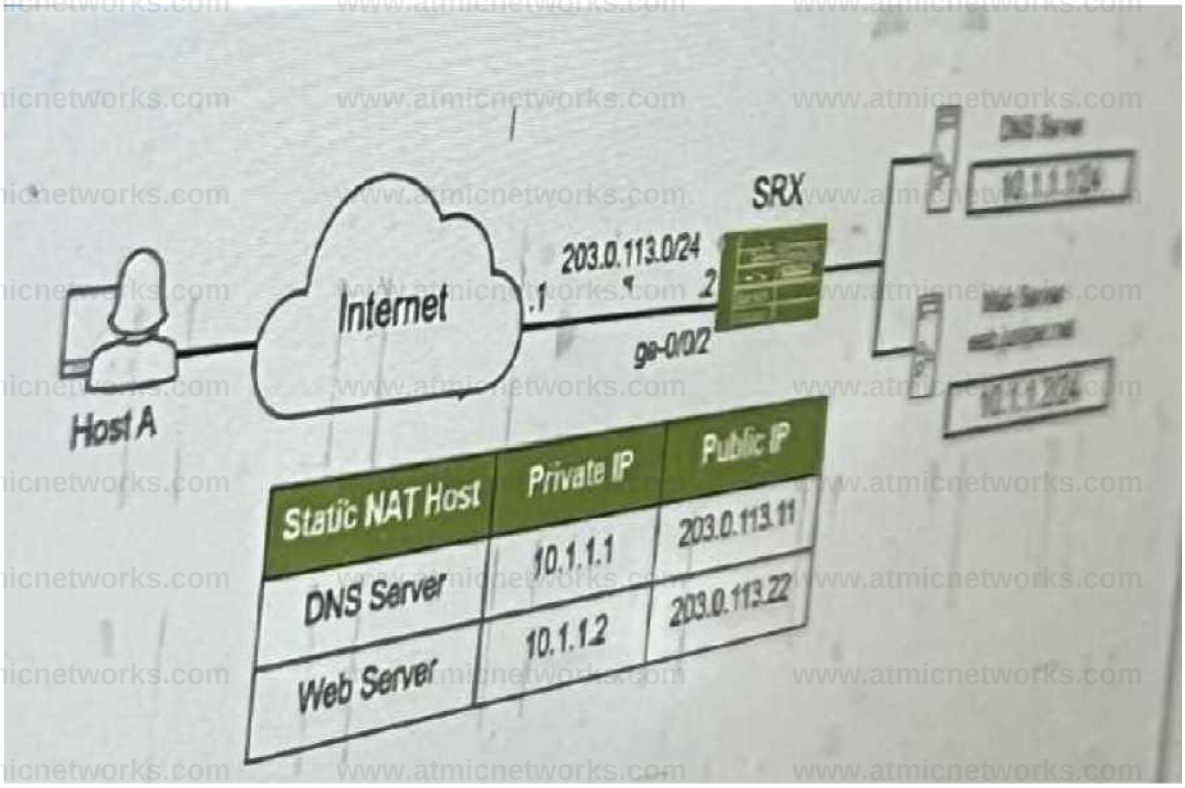

Referring to the exhibit, which two Junos features are required to accomplish this task? (Choose two.)

A. DNS doctoring

B. proxy ARP

C. persistent NAT

D. STUN

Answer: AB

Explanation:

Question: 88

You want to enable transparent mode on your SRX series device.

In this scenario, which three actions should you perform? (Choose three.)

A. Enable the ethernet-switching family on your Layer 2 interfaces

B. Install a Layer 2 feature license.

C. Reboot the SRX device.

D. Ensure that no IRB interfaces are configured on the device.

E. Add your Layer 2 interfaces to a security zone.

Answer: ACE

Explanation:

Question: 89

Referring to the exhibit, you have been assigned the user LogicalSYS1 credentials shown in the configuration.

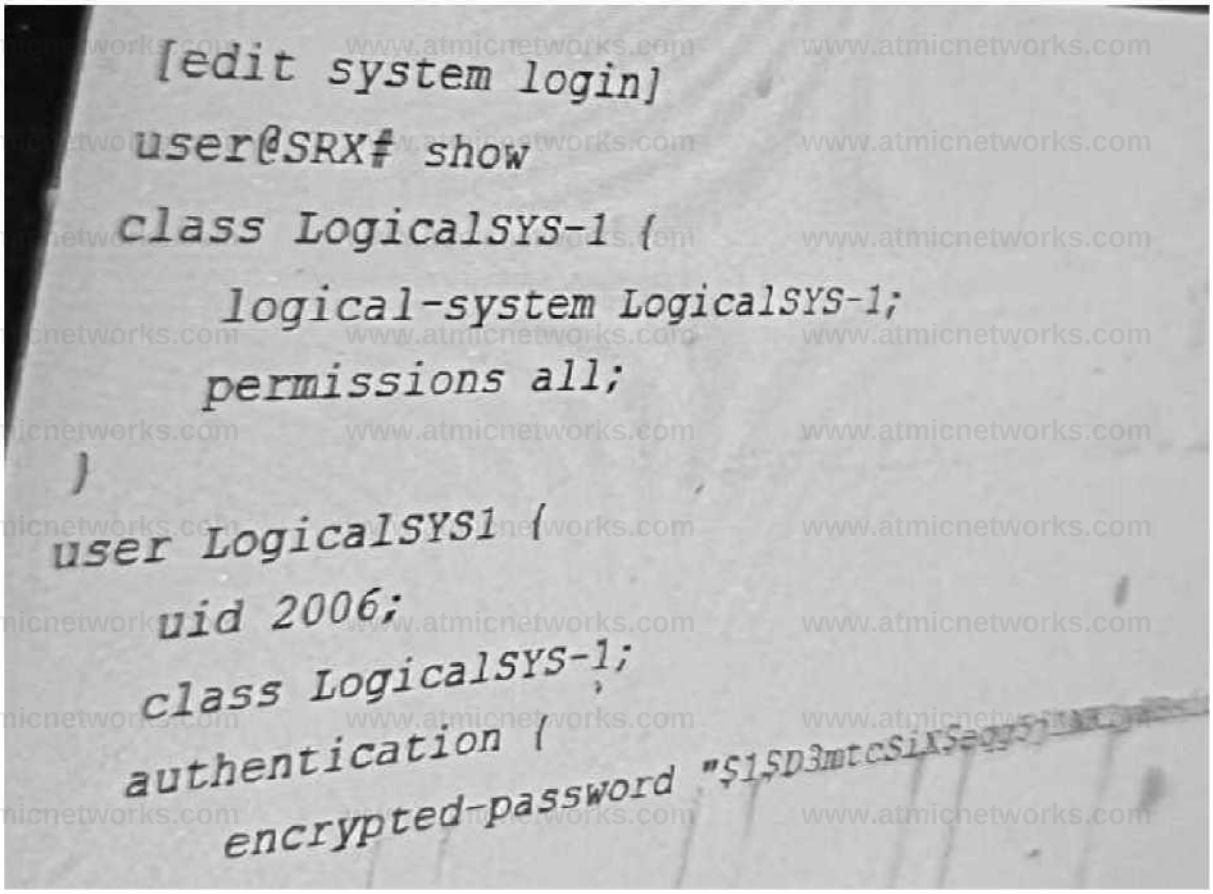

In this scenario, which two statements are correct? (Choose two.)

A. When you log in to the device, you will be permitted to view all routing tables available on the SRX device
B. When you login to the device, you will be permitted to view only the routing tables for Logic
C. When you login to the device, you will be located at the operational mode of the Logic
D. When you login to the device, you will be located at the operational mode of the main system

Answer: BC

Explanation:

Question: 90

Exhibit:

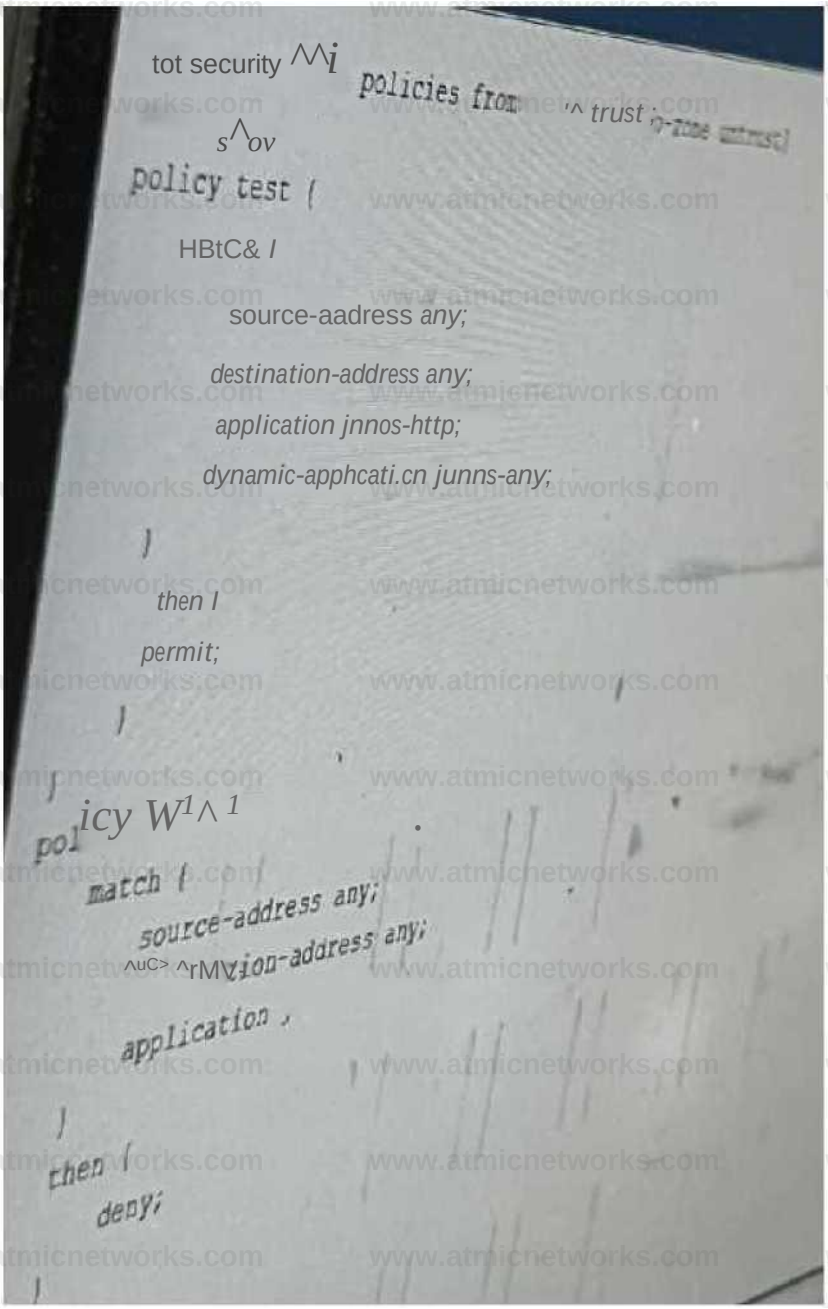

You created a Unified security policy called test on the network edge srx series firewall.
According to the firewall, this new security policy is not passing traffic.

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

A. The test policy should be the last policy.

B. A match exists on the test policy, but the dynamic application is waiting to be discovered

C. The source address cannot be any when a dynamic application is configured.

D. The drop-http policy is a terminating rule and will drop the traffic.

Answer: BD

Explanation:

Question: 91

Exhibit:

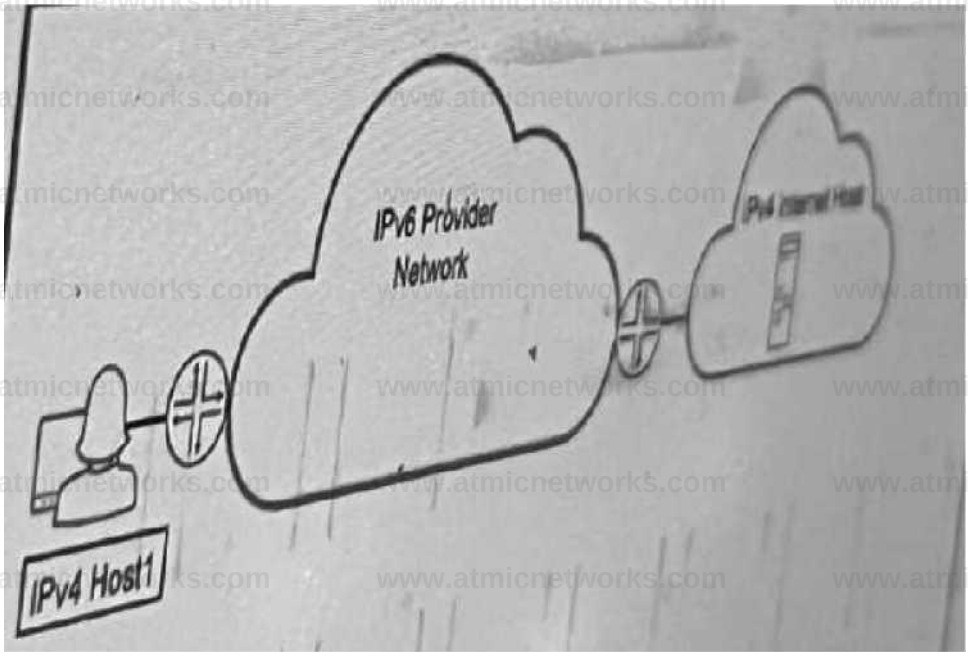

Referring to the exhibit, which technology would you use to provide communication between IPv4 host1 and ipv4 internal host

A. DS-Lite

B. NAT444

C. NAT46

D. full cone NAT

Answer: A

Explanation:

Question: 92

You are attempting to ping the IP address that is assigned to the loopback interface on the SRX series device shown in the exhibit.

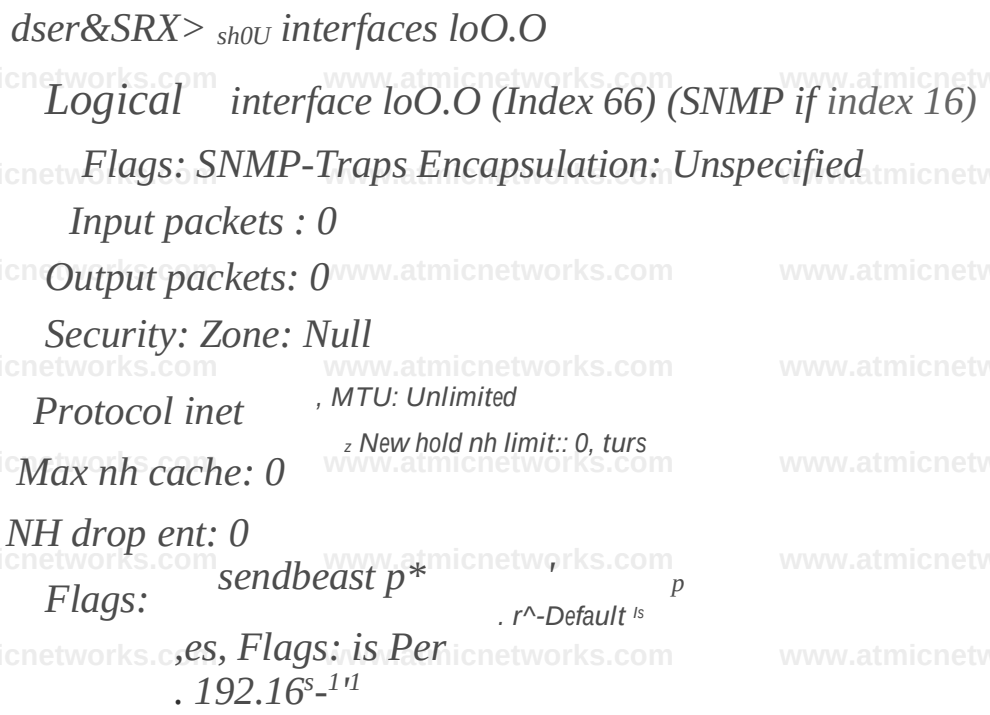

What is causing this problem?

A. The loopback interface requires encapsulation.

B. The loopback interface is not assigned to a security zone.

C. The incorrect interface index ID is assigned to the loopback interface.

D. The IP address on the loopback interface is a private address.

Answer: C

Explanation:

Question: 93

What are three requirements to run OSPF over GRE over IPsec? (Choose Three)

A. The GRE interface must be configured in OSPF Area 0.

B. The OSPF interface must be placed in a zone and must have GRE configured

C. Overlapping addresses should exist between remote networks.

D. The GRE interface must be placed in a zone and must have OSPF configured in is host

E. Overlapping addresses should not exist between remote networks.

Answer: BDE

Explanation:

Question: 94

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

Which two configuration elements are required when you generate your certificate request? (Chose 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

Explanation:

Question: 95

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

Explanation:

Question: 96

Which two statements about transparent mode and Ethernet switching mode on an SRX series device are correct.

A. In Ethernet switching mode, Layer 2 interfaces must be placed in a security zone.

B. In Ethernet switching mode, IRB interfaces must be placed in a security zone.

C. In transparent mode, Layer 2 interfaces must be placed in a security zone.

D. In transparent mode, IRB interfaces must be placed in a security zone.

Answer: BC

Explanation:

Question: 97

A customer wants to be able to initiate a return connection to an internal host from a specific Server.

Which NAT feature would you use in this scenario?

A. target-host

B. any-remote-host

C. port-overloading

D. target-server

Answer: A

Explanation:

Question: 98

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

Explanation:

Question: 99

You are configuring advanced policy-based routing. You have created a static route with next hop of an interface in your inet.0 routing table

U'CrfSfA'i show routing—instances APBRinstance (instance-type forwarding;

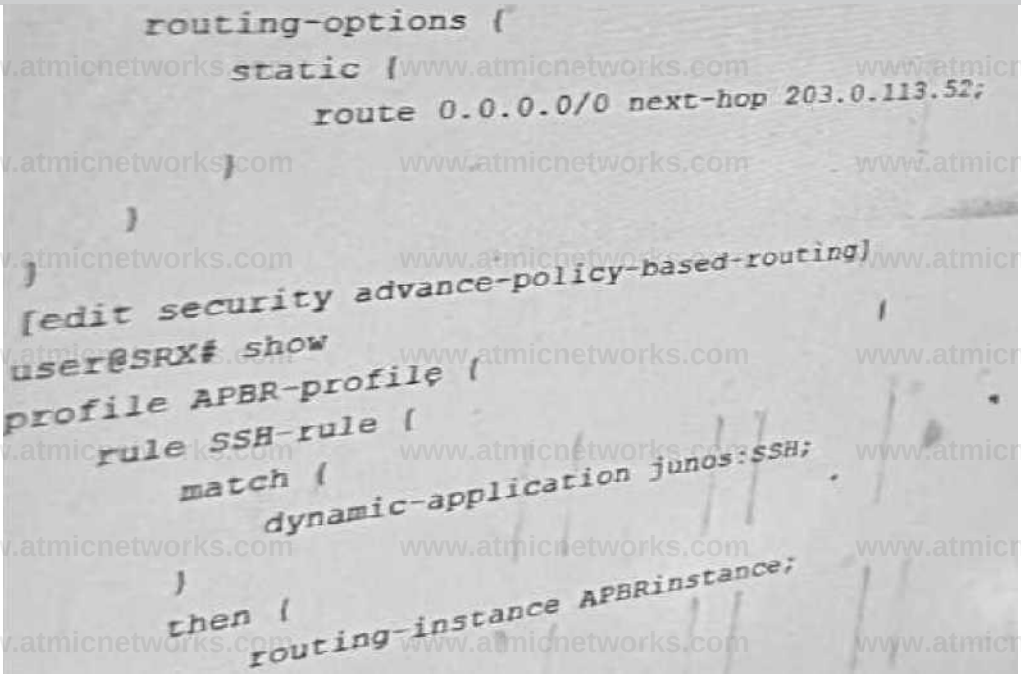

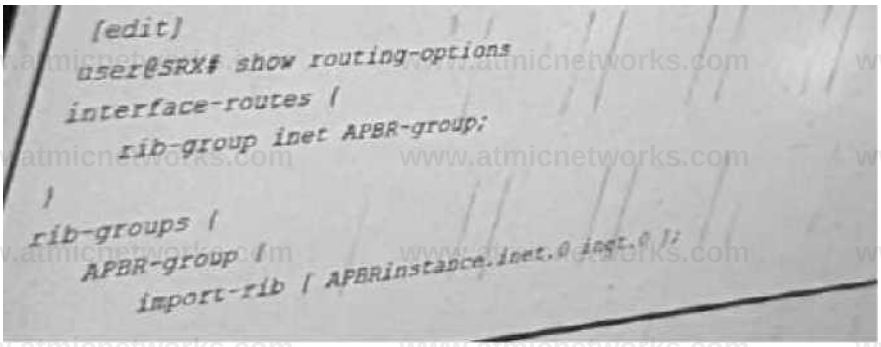

Referring to the exhibit, what should be changed to solve this issue?

A. You should change the routing instance type to virtual-router.

B. You should move the static route configuration to the main routing instance.

C. You should move the inet. o table before the routing instance table in your rib-groups configuration.

D. You should delete the interface-routes configuration under the routing-options hierarchy.

Answer: C

Explanation:

Question: 100

What are three attributes that APBR queries from the application system cache module. (Choose Three)

A. TTL

B. destination port

C. service

D. DSCP

E. protocol type

Answer: BCE

Explanation:

Question: 101

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

Explanation:

Question: 102

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

Explanation:

Question: 103

You Implement persistent NAT to allow any device on the external side of the firewall to initiate traffic.

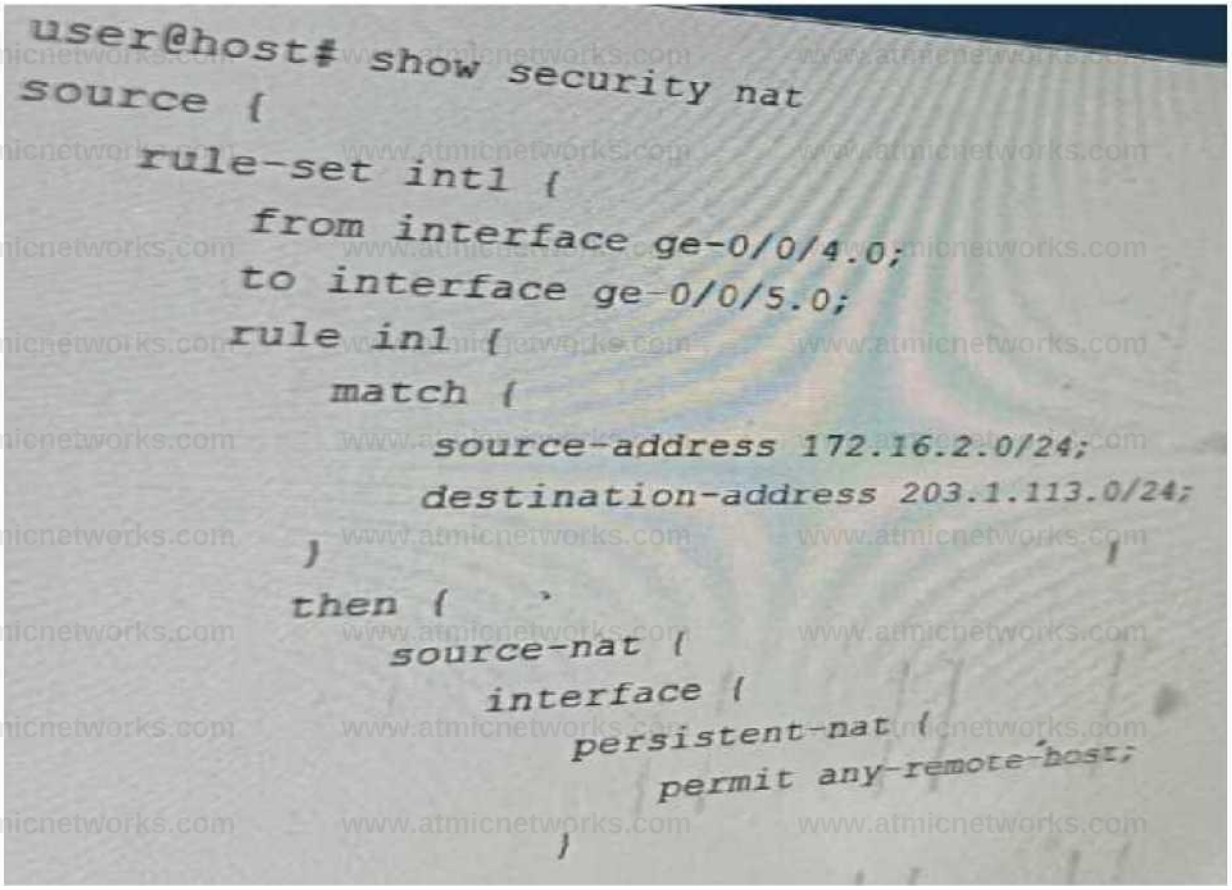

Referring to the exhibit, which statement is correct?

A. The target-host parameter should be used instead of the any-remote-host parameter.

B. The port-overloading parameter needs to be turned off in the NAT source interface configuration

C. The target-host-port parameter should be used instead of the any-remote-host parameter

D. The any-remote-host parameter does not support interface-based NAT and needs an IP pod to work.

Answer: D

Explanation:

Question: 104

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

Explanation:

Question: 105

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

Explanation:

Question: 106

A user reports that a specific application is not working properly. This application makes multiple connection to the

server and must have the same address every time from a pool and this behavior needs to be changed.

What would solve this problem?

A. Use STUN.

B. Use DNS doctoring.

C. Use the address-persistent parameter.

D. Use the persistent-nat parameter.

Answer: D

Explanation:

Question: 107

You have cloud deployments in Azure, AWS, and your private cloud. You have deployed multicloud using security director with policy enforcer to. Which three statements are true in this scenario? (Choose three.)

A. You can run Juniper ATP scans only on traffic from your private cloud.

B. You can run Juniper ATP scans for all three domains.

C. You must secure the policies individually by domain.

D. The Policy Enforcer is able to flag infected hosts in all three domains.

E. You can simultaneously manage the security policies in all three domains.

Answer: BDE

Explanation:

Question: 108

Which two statements describe the behavior of logical systems? (Choose two.)

A. Each logical system shares the routing protocol process.

B. A default routing instance must be manually created for each logical system

C. Each logical system has a copy of the routing protocol process.

D. A default routing instance is automatically created for each logical system.

Answer: CD

Explanation:

Question: 109

Which two statements are correct about advanced policy-based routing?

A. It can use the application system cache to route traffic.

B. The associated routing instance should be configured as a virtual router instance.

C. It cannot use the application system cache to route traffic.

D. The associated routing instance should be configured as a forwarding instance.

Answer: AD

Explanation:

Question: 110

You have deployed a new site as shown in the exhibit. Hosts in the 10.10.10.0/24 network must access the DB1 server. The DB1 server must also have internet access the DB1 server encrypted.

Which two configuration statements will be required as part of the configuration on SRX1 to satisfy this requirement? (Choose two)

A. set security macsec interfaces ge-0/0/1 connectivity association access-sw

B. set protocols 12-learning global mode transpårent-bridge

C. set security forwarding-options secure-wire access-sw interface ge-0/0/1.0

D. set security macsec connectivity-association access-sw security-mode static-cak

Answer: AD

Explanation:

Question: 111

Exhibit:

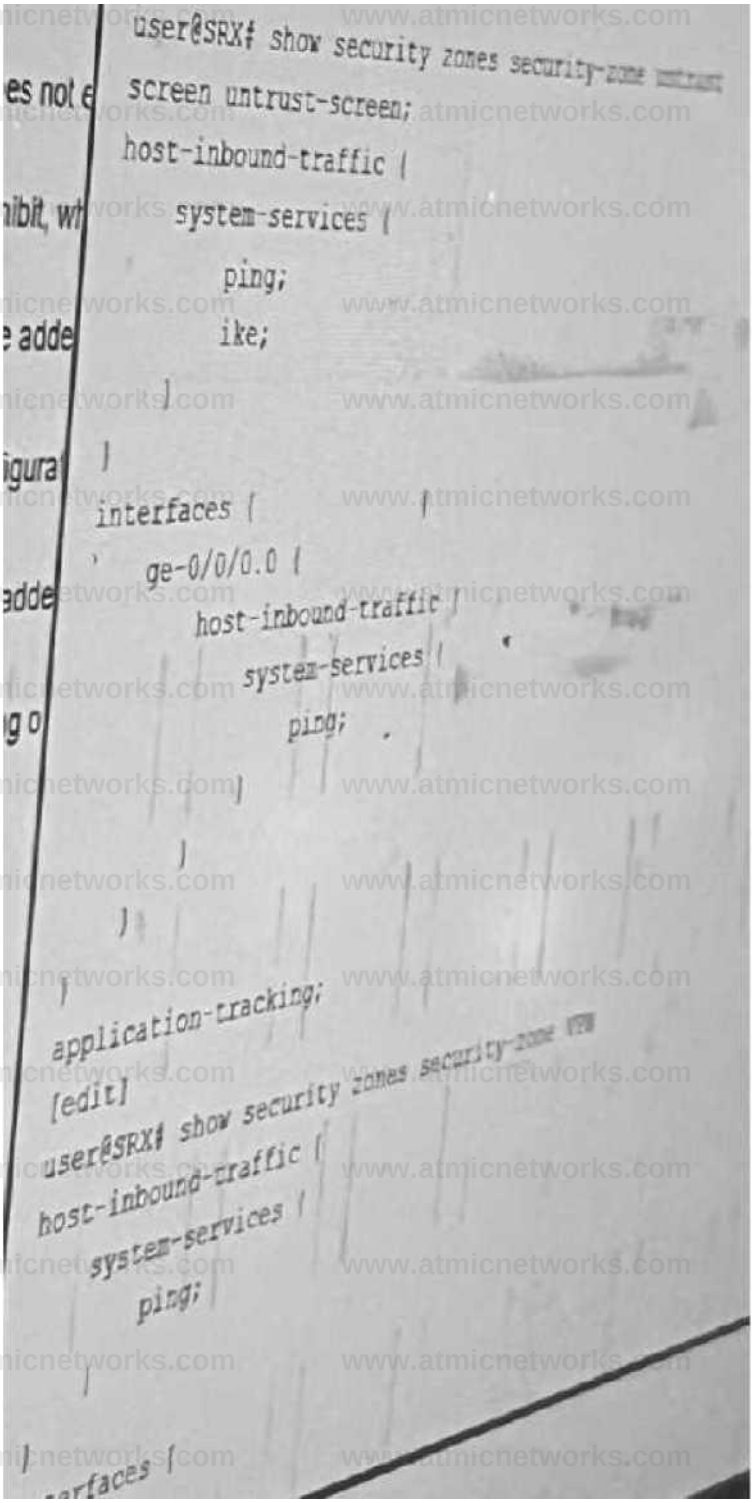

The Ipsec VPN does not establish when the peer initiates, but it does establish when the SRX series device initiates. Referring to the exhibit, what will solve this problem?

A. IKE needs to be added for the host-inbound traffic on the VPN zone.

B. The screen configuration on the untrust zone needs to be modified.

C. IKE needs to be added to the host-inbound traffic directly on the ge-0/0/0 interface.

D. Application tracking on the untrust zone needs to be removed.

Answer: C

Explanation:

Question: 112

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

Explanation:

Question: 113

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

Explanation:

Question: 114

Which encapsulation type must be configured on the lt-0/0/0 logical units for an interconnect logical systems VPLS switch?

A. encapsulation ethernet-bridge

B. encapsulation ethernet

C. encapsulation ethernet-vpls

D. encapsulation vlan-vpls

Answer: C

Explanation:

Question: 115

Referring to the exhibit, which two statements are true ?

show

P Statistics:

Encrypted bytes

Encrypted packets:

Decrypted packets:

AH Statistics:

Output bytes: . Input packets: Output packets:

nation fail-es: 0 ^

tlcation failures: o. -

0, Bad trailera- -

A. Every VPN packet that the SRX receives from the VPN peer is outside the ESP sequence window

B. The SRX is sending traffic into the tunnel and out toward the VPN peer.

C. The SRX is not sending any packets to the VPN peer.

D. The SRX is not receiving any packets from the VPN peer.

Answer: BD

Explanation: