DUMPSARENA

Security, Professional (JNCIP-SEC)

Juniper JN0-636

Version Demo

Total Demo Questions: 16

Total Premium Questions: 161

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Security Policy	40
Topic 2, NAT	19
Topic 3, IPsec VPNs	22
Topic 4, Advanced Threat Prevention	52
Topic 5, Security Management	26
Topic 6, Mix Questions	2
Total	161

QUESTION NO: 1

What is the purpose of the Switch Microservice of Policy Enforcer?

A. to isolate infected hosts

B. to enroll SRX Series devices with Juniper ATP Cloud

C. to inspect traffic for malware

D. to synchronize security policies to SRX Series devices

ANSWER: A

Explanation:

The purpose of the Switch Microservice of Policy Enforcer is **to isolate infected hosts**. Policy Enforcer integrates threat intelligence and endpoint context with enforcement infrastructure. When a host is identified as infected or otherwise compromised, the Switch Microservice communicates with supported access switches to apply the containment action at the point where that endpoint connects to the network. This enables the affected device to be placed into an isolated or quarantine network segment, limiting its access to production resources while allowing the organization to investigate and remediate it.

This switch-based enforcement is particularly useful for endpoints that are directly attached to campus or branch access ports, because it can contain a threat quickly without requiring a manual switch-port change. The isolation workflow is part of Policy Enforcer's automated response capabilities: threat detection supplies the security event, and the Switch Microservice performs the network-access enforcement needed to prevent lateral movement. Juniper describes Policy Enforcer as a solution for automated threat response and enforcement across the network. See the Juniper Policy Enforcer product information and Juniper's technical documentation portal.

QUESTION NO: 2

You are asked to configure a security policy on the SRX Series device. After committing the policy, you receive the "Policy is out of sync between RE and PFE ." error.

Which command would be used to solve the problem?

A. request security policies resync

B. request service-deployment

C. request security polices check

D. restart security-intelligence

ANSWER: A

Explanation:

`request security policies resync` is the operational-mode command used to synchronize the security-policy configuration from the Routing Engine (RE) to the Packet Forwarding Engine (PFE)/Services Processing Unit (SPU). Security policies are configured on the RE, but policy enforcement requires the corresponding compiled policy information to be installed on the data-plane processing components. The reported message means that this synchronization did not complete successfully for one or more SPUs, leaving the control-plane policy state and forwarding-plane policy state inconsistent. Running the resynchronization command requests that Junos resend and reinstall the policy information to the affected security processing hardware, restoring the expected RE-to-PFE policy consistency without requiring the policy to be manually recreated. This command is specifically intended for correcting an out-of-sync security-policy condition after a commit. Juniper documents the command in the request security policies resync CLI reference; the associated SRX troubleshooting guidance is available in Juniper KB30443.

QUESTION NO: 3

You are asked to deploy filter-based forwarding on your SRX Series device for incoming traffic sourced from the 10.10 100 0/24 network in this scenario, which three statements are correct? (Choose three.)

A. You must create a forwarding-type routing instance.

B. You must create and apply a firewall filter that matches on the source address 10.10.100.0/24 and then sends this traffic to your routing instance.

C. You must create and apply a firewall filter that matches on the destination address 10 10.100.0/24 and then sends this traffic to your routing instance.

D. You must create a RIB group that adds interface routes to your routing instance.

E. You must create a VRF-type routing instance.

ANSWER: A B D

Explanation:

Filter-based forwarding on an SRX Series device uses an ingress firewall filter to classify packets and select an alternate routing table for the matched traffic. Therefore, “You must create a forwarding-type routing instance” is correct: a forwarding instance provides the separate routing table used by the firewall filter’s `then routing-instance` action. “You must create and apply a firewall filter that matches on the source address 10.10.100.0/24 and then sends this traffic to your routing instance” is also correct because the filter is applied inbound on the receiving interface and its source-prefix match identifies precisely the traffic that requires alternate forwarding treatment. Packets that match are route-looked-up in the selected forwarding instance’s routing table.

“You must create a RIB group that adds interface routes to your routing instance” is correct when the forwarding instance needs awareness of directly connected routes. The RIB group imports the required interface routes from the main routing table into the forwarding instance’s `inet.0` table, allowing the selected routing table to resolve next hops and forward traffic correctly. This is a standard component of an SRX FBF deployment when connected-route reachability is required. Juniper describes the forwarding-instance and firewall-filter model in its Filter-Based Forwarding Overview; RIB import behavior is documented in RIB Groups.

QUESTION NO: 4

You have a webserver and a DNS server residing in the same internal DMZ subnet. The public Static NAT addresses for the servers are in the same subnet as the SRX Series devices internet-facing interface. You implement DNS doctoring to ensure remote users can access the webserver.Which two statements are true in this scenario? (Choose two.)

A. The DNS doctoring ALG is not enabled by default.

B. The Proxy ARP feature must be configured.

C. The DNS doctoring ALG is enabled by default.

D. The DNS CNAME record is translated.

ANSWER: B C

Explanation:

The Proxy ARP feature must be configured. Static NAT maps a public address to the server’s private address, but the SRX must also be able to receive Ethernet frames sent by upstream devices to that public address. Because the public static-NAT addresses are on the same Layer 2 subnet as the Internet-facing SRX interface, proxy ARP causes the SRX to answer ARP requests for those translated addresses. This directs traffic for the published web server and DNS server to the SRX, where the static-NAT translation and security-policy processing can occur. Juniper documents proxy ARP as the required mechanism when a translated address resides in the same subnet as the ingress interface: Static NAT Overview.

The DNS doctoring ALG is enabled by default. The SRX DNS application-layer gateway is enabled by default and supports DNS doctoring for static NAT. When an internal DNS response contains an address record for a host with a static-NAT mapping, the ALG can replace the private address in the response with that host's public mapped address. Consequently, remote clients receive a routable public address for the web server while the server remains addressed privately in the DMZ. Juniper's DNS ALG documentation describes its default enablement and DNS response address translation behavior: DNS ALG Overview.

QUESTION NO: 5

You configure both a static NAT rule and a destination NAT rule that could match the same incoming packet on an SRX Series device.

Which NAT rule type is evaluated first?

A. Static NAT

B. Destination NAT

C. Source NAT

D. Proxy ARP

ANSWER: A

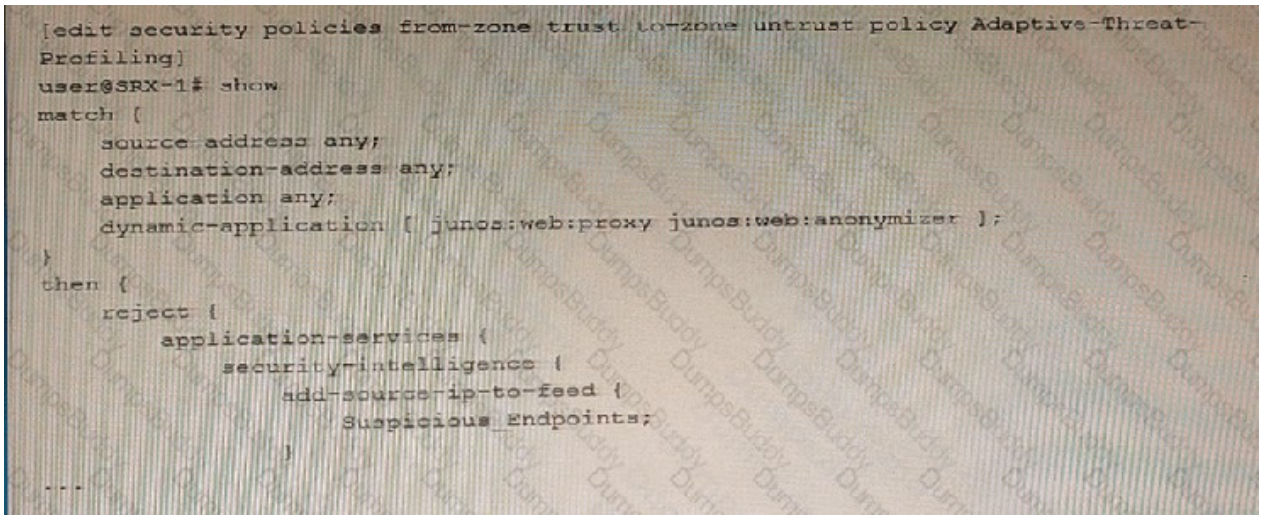
Explanation:

Static NAT is evaluated before destination NAT. Junos performs NAT rule processing in a defined order. For inbound traffic, static NAT rule matching occurs first; destination NAT is considered only if no static NAT rule matches the packet. This precedence prevents a destination NAT rule from overriding an applicable one-to-one static mapping.

Static NAT provides bidirectional address translation between an internal address and an external address, whereas destination NAT is normally used to translate inbound traffic to an internal server address or service. Administrators should account for static NAT precedence when designing overlapping NAT rule matches. See Juniper's static NAT documentation and destination NAT documentation.

QUESTION NO: 6

Exhibit



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

A. The suspicious_Endpoints feed is only usable by the SRX-1 device.

B. You must manually create the suspicious_Endpoints feed in the Juniper ATP Cloud interface.

C. The suspicious_Endpoints feed is usable by any SRX Series device that is a part of the same realm as SRX-1.

D. Juniper ATP Cloud automatically creates the suspicious_Endpoints feed after you commit the security policy.

ANSWER: C D

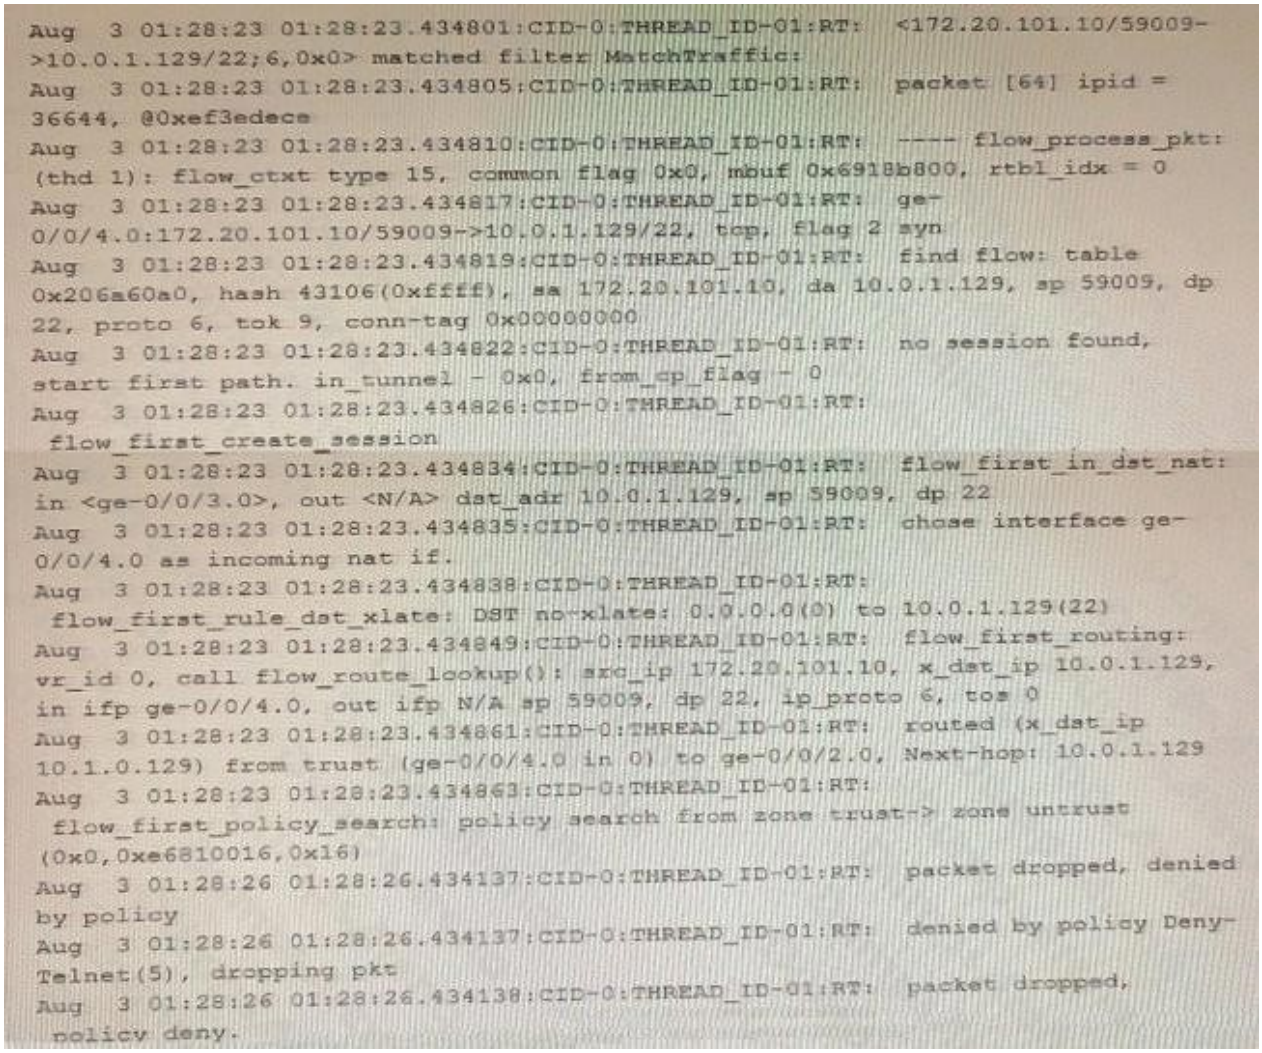
Explanation:

The *suspicious_Endpoints* feed is usable by any SRX Series device in the same Juniper ATP Cloud realm as SRX-1. ATP Cloud realm membership provides the shared context for cloud-delivered threat intelligence and IoT policy-enforcement data, allowing other enrolled SRX devices in that realm to consume the dynamically maintained endpoint information. This lets administrators apply consistent policy controls to endpoints identified through IoT discovery without limiting the feed to the device that first observed the traffic.

Juniper ATP Cloud automatically creates the *suspicious_Endpoints* feed after the security policy referencing it is committed. The feed is part of the IoT Device Discovery and Policy Enforcement workflow: SRX devices send relevant IoT telemetry for cloud analysis, ATP Cloud classifies devices and identifies suspicious endpoints, and the resulting dynamic endpoint data becomes available for policy use. Referencing the feed in a committed policy establishes the required configuration, while ATP Cloud maintains the feed contents as classifications and risk assessments change. Juniper documents this workflow in its IoT Device Discovery overview and IoT policy-enforcement configuration guide.

QUESTION NO: 7

Exhibit



Referring to the exhibit, which statement is true?

A. This custom block list feed will be used before the Juniper SecIntel

B. This custom block list feed cannot be saved if the Juniper SecIntel block list feed is configured.

C. This custom block list feed will be used instead of the Juniper SecIntel block list feed

D. This custom block list feed will be used after the Juniper SecIntel block list feed.

ANSWER: D

Explanation:

This custom block list feed will be used after the Juniper SecIntel block list feed. The feed ordering shown in the configuration determines the sequence in which Security Intelligence processes the configured block-list sources. In this case, the Juniper-provided SecIntel block-list feed has precedence in the displayed order, and the custom feed is evaluated subsequently. This ordering lets administrators supplement Juniper's continuously maintained intelligence with organization-specific indicators, such as internally identified malicious addresses or domains, without replacing the Juniper SecIntel source. Both sources remain part of the security-intelligence configuration and can contribute intelligence used to identify traffic associated with known threats. Security Intelligence is designed to use threat feeds to make enforcement decisions early, helping prevent traffic associated with listed indicators from reaching protected resources. Juniper's Security Intelligence documentation describes the use of Juniper threat intelligence and custom feeds within a security-intelligence policy: Juniper Security Intelligence documentation. For the broader product documentation and current release-specific configuration guidance, consult the Juniper Documentation portal.

QUESTION NO: 8

You want an SRX Series device to generate a security-policy log entry when a session is created after matching a permit policy.

Which logging option should be configured in the policy?

A. session-init

B. session-close

C. log-invalid-sessions

D. packet-log

ANSWER: A

Explanation:

session-init is the correct logging option. Configuring `session-init` under the policy logging statement creates a log entry when the SRX creates a session that matches the policy. This provides immediate visibility that the policy was selected for the flow.

By comparison, `session-close` generates a log record after the session terminates and includes session-end information such as duration and byte counts. Both options can be configured when operational visibility is needed at session establishment and termination. See Juniper's security policy log statement reference.

QUESTION NO: 9

What are two valid modes for the Juniper ATP Appliance? (Choose two.)

A. flow collector

B. event collector

C. all-in-one

D. core

ANSWER: A C

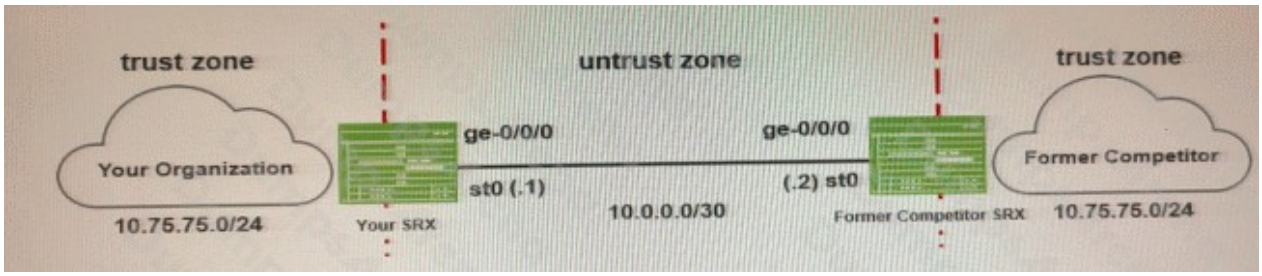
Explanation:

The valid Juniper ATP Appliance operating modes are **flow collector** and **all-in-one**. In all-in-one mode, one ATP Appliance performs the full on-premises threat-prevention workflow. It receives applicable traffic or files, extracts and submits samples for analysis, applies the available detection capabilities, and provides the management and reporting functions needed for the deployment. This mode is appropriate when the required capacity and resiliency can be provided by a single appliance deployment.

Flow collector mode is used when the appliance is dedicated to collecting flows and files from protected network devices, such as SRX Series devices. The collector obtains candidate files from observed traffic and forwards them for ATP analysis according to the configured architecture and policies. This permits the file-collection function to be placed close to monitored traffic sources and supports deployments that separate collection from other ATP processing functions. Juniper's ATP Appliance documentation describes the appliance deployment and operating model, including all-in-one and flow-collector use cases. See the Juniper ATP Appliance overview and the Juniper ATP Appliance documentation portal for configuration guidance.

QUESTION NO: 10

Exhibit



Your company recently acquired a competitor. You want to use using the same IPv4 address space as your company.

Referring to the exhibit, which two actions solve this problem? (Choose two)

A. Configure static NAT on the SRX Series devices.

B. Connect the competitor network using IPsec policy-based VPNs.

C. Identify two neutral IPv4 address spaces for address translation.

D. Configure IPsec Transport mode.

ANSWER: A C

Explanation:

Configure static NAT on the SRX Series devices and **Identify two neutral IPv4 address spaces for address translation** are the required actions when the two organizations have overlapping IPv4 prefixes. Static NAT provides consistent, one-to-one address mappings, so a host can be represented by a translated address when it communicates with the acquired company while retaining its existing address within its own local network. This preserves the existing addressing plan and avoids requiring an immediate enterprise-wide renumbering effort.

Two neutral, nonoverlapping address ranges must be selected for the translations. Each organization can then use translated addresses that are unique from the perspective of the other organization. The SRX devices apply the corresponding static NAT mappings at the interconnection boundary, allowing sessions to be forwarded and return traffic to be translated back to the original internal addresses. In other words, the neutral ranges supply unique identities for communication across the boundary, and static NAT enforces the deterministic mapping between those identities and the duplicated internal addresses. Juniper documents static NAT as a fixed mapping between original and translated addresses, which is appropriate where predictable bidirectional reachability is needed. See Juniper static NAT documentation and Juniper NAT overview.

QUESTION NO: 11

You want to enable inter-tenant communica ???? on with tenant system.

In this Scenario, Which two solutions will accomplish this task?

A. interconnect EVPN switch

B. interconnect VPLS switch

C. external router

D. logical tunnel interface

ANSWER: C D

Explanation:

external router and **logical tunnel interface** are the supported methods for providing communication between tenant systems. Tenant systems are designed to maintain separate routing and security administration contexts, so traffic that must pass from one tenant to another requires an explicitly configured Layer 3 path. An external router can provide that path by connecting to interfaces assigned to the relevant tenant systems. Routes and any required security policies are then configured so that the router forwards traffic between the tenant networks while preserving a clear, physical separation point.

A logical tunnel interface provides the equivalent connectivity internally on the Junos device. Logical tunnel interface units are paired and assigned to the appropriate tenant contexts, creating a point-to-point virtual link. Each side can be addressed and routed independently, allowing traffic to be exchanged through deliberate routing and policy configuration without requiring a separate physical router. This approach is particularly useful where inter-tenant connectivity is needed but external cabling or an additional routing device is undesirable. Juniper's tenant-system documentation identifies external routing and logical tunnel interfaces as the mechanisms used for inter-tenant communication. See the Juniper Tenant Systems documentation and the Juniper logical tunnel interfaces documentation.

QUESTION NO: 12

Exhibit

A. What are two appropriate mitigation actions for the selected incident? (Choose two.)

B. Immediate response required: Block malware IP addresses (download server or CnC server)

C. Immediate response required: Wipe infected endpoint hosts.

D. Immediate response required: Deploy IVP integration (if configured) to confirm if the endpoint has executed the malware and is infected.

E. Not an urgent action: Use IVP to confirm if machine is infected.

ANSWER: B D

Explanation:

For an incident at the *Download* stage of the kill chain, the appropriate immediate actions are to block the identified malware IP addresses and to deploy IVP integration, when it is configured. Blocking the download-server or command-and-control IP addresses disrupts the malicious infrastructure associated with the observed file and prevents additional downloads or subsequent communications from reaching that infrastructure. This containment action can be applied promptly through the security controls integrated with the ATP workflow.

Deploying IVP integration provides endpoint-focused validation of whether the downloaded sample was actually executed and whether the host is infected. That confirmation is important because a download event establishes exposure, while endpoint telemetry and integration can establish execution and infection status. The resulting endpoint context supports accurate containment and incident handling for affected devices. Juniper Advanced Threat Prevention correlates threat-analysis results with mitigation workflows to help security teams contain threats and investigate their impact. See Juniper's

Advanced Threat Prevention overview and the Juniper technical documentation portal for product and operational documentation.

QUESTION NO: 13

you are connecting two remote sites to your corporate headquarters site. You must ensure that traffic

passes corporate headquarter.

In this scenario, which VPN should be used?

A. full mesh IPsec VPNs with tunnels between all sites

B. a full mesh Layer 3 VPN with the BGP route reflector behind the corporate firewall device

C. a Layer 3 VPN with the corporate firewall acting as the hub device

D. hub-and-spoke IPsec VPN with the corporate firewall acting as the hub device

ANSWER: D

Explanation:

hub-and-spoke IPsec VPN with the corporate firewall acting as the hub device is the appropriate design because it explicitly centralizes connectivity and intersite forwarding at the headquarters. Each remote site establishes an IPsec tunnel to the corporate firewall at the headquarters, which is the hub. When one remote site needs to communicate with the other, its encrypted traffic is sent first to that hub; the headquarters firewall then applies the organization's security policy, performs any required routing or inspection, and forwards the traffic through the applicable spoke tunnel. This provides the required deterministic traffic path through the corporate headquarters rather than permitting a direct remote-site-to-remote-site path.

Using the firewall as the hub also gives the organization one central enforcement point for access control, logging, threat-prevention services, and route policy. IPsec supplies peer authentication, integrity protection, confidentiality, and anti-replay protection across the untrusted network, while the hub-and-spoke topology keeps tunnel and policy administration practical as additional branches are added. Juniper documents IPsec VPN configuration and operational concepts for SRX devices in its IPsec VPN documentation. The SRX platform's role in securely connecting distributed locations is also described in the Juniper SRX Series overview.

QUESTION NO: 14

You are connecting two remote sites to your corporate headquarters site; you must ensure that all traffic is secured and

only uses a single Phase 2 SA for both sites.

In this scenario, which VPN should be used?

A. An IPsec group VPN with the corporate firewall acting as the hub device.

B. Full mesh IPsec VPNs with tunnels between all sites.

C. A hub-and-spoke IPsec VPN with the corporate firewall acting as the hub device.

D. A full mesh Layer 3 VPN with the corporate firewall acting as the hub device.

ANSWER: A

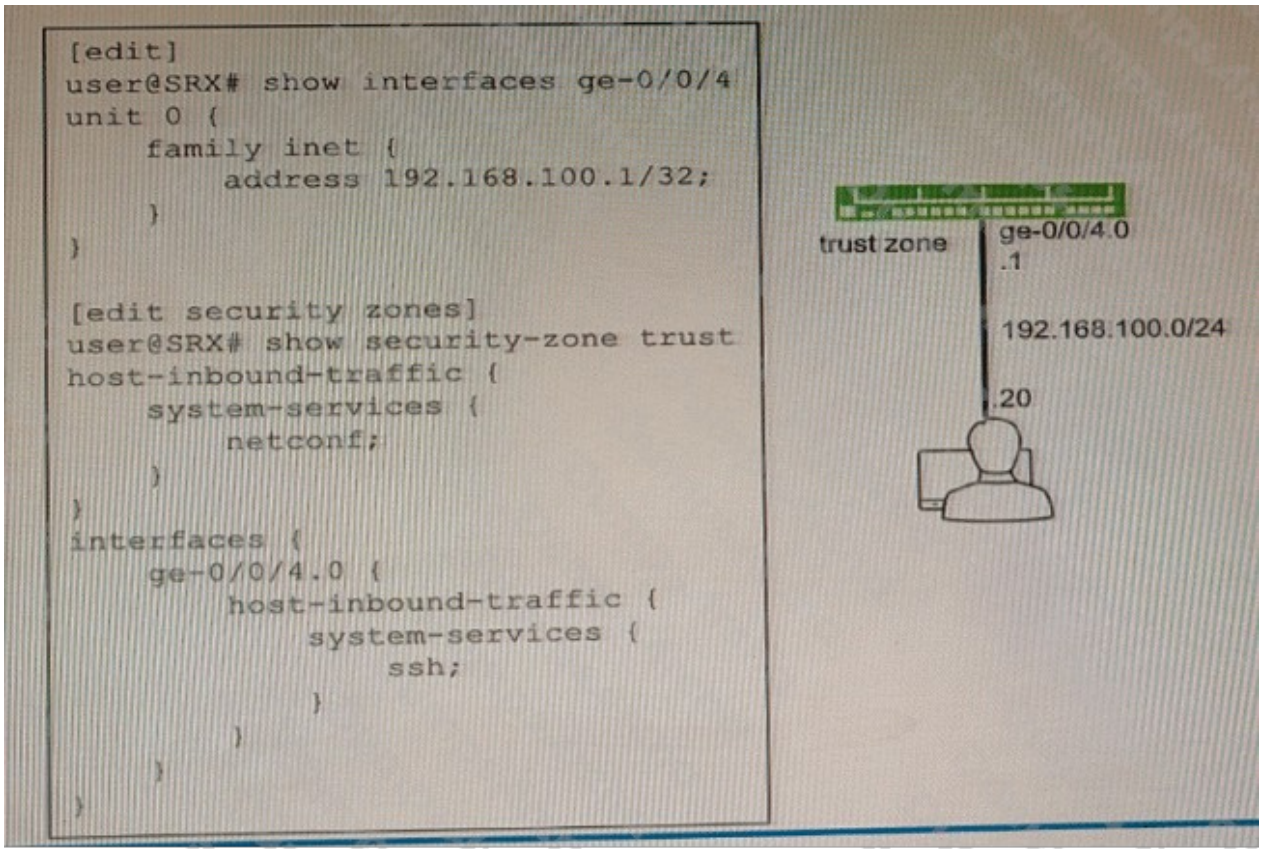
Explanation:

An IPsec group VPN with the corporate firewall acting as the hub device is correct because Group VPN is designed for secure, scalable communication among multiple VPN members while minimizing IPsec SA usage at the central device. In a Group VPN deployment, the headquarters firewall operates as the group controller or hub. It authenticates members and distributes the group security information required for protected traffic. The remote sites can then participate in the same

protected group rather than requiring the hub to maintain a distinct Phase 2 IPsec SA for every individual remote-site relationship. This shared-SA model is specifically useful where several remote devices need secure connectivity to a central site and the design requirement is to conserve VPN and security-association resources. The Phase 2 SA provides the IPsec security parameters used to protect the actual user traffic, including the negotiated ESP or AH protection settings. Juniper documents Group VPN as a multipoint VPN capability in which group members share IPsec security associations, with the group controller managing group membership and keys. See Juniper's Group VPN documentation and the IPsec security association overview.

QUESTION NO: 15

Exhibit



You are not able to ping the default gateway of 192.168 100 1 (or your network that is located on your SRX Series firewall.

Referring to the exhibit, which two commands would correct the configuration of your SRX Series device? (Choose two.)

A)

```
[edit security zones security-zone trust]
user@SRX# set interfaces ge-0/0/4.0 host-inbound-traffic system-services ping
```

B)

```
[edit interfaces ge-0/0/4]
user@SRX# replace pattern 32 with 24
```

C)

```
[edit security zones security-zone trust]
user@SRX# set host-inbound-traffic system-services ping
```

D)

```
[edit security zones security-zone trust]
user@SRX# set host-inbound-traffic system-services ping except
```

A. Option A

B. Option B

C. Option C

D. Option D

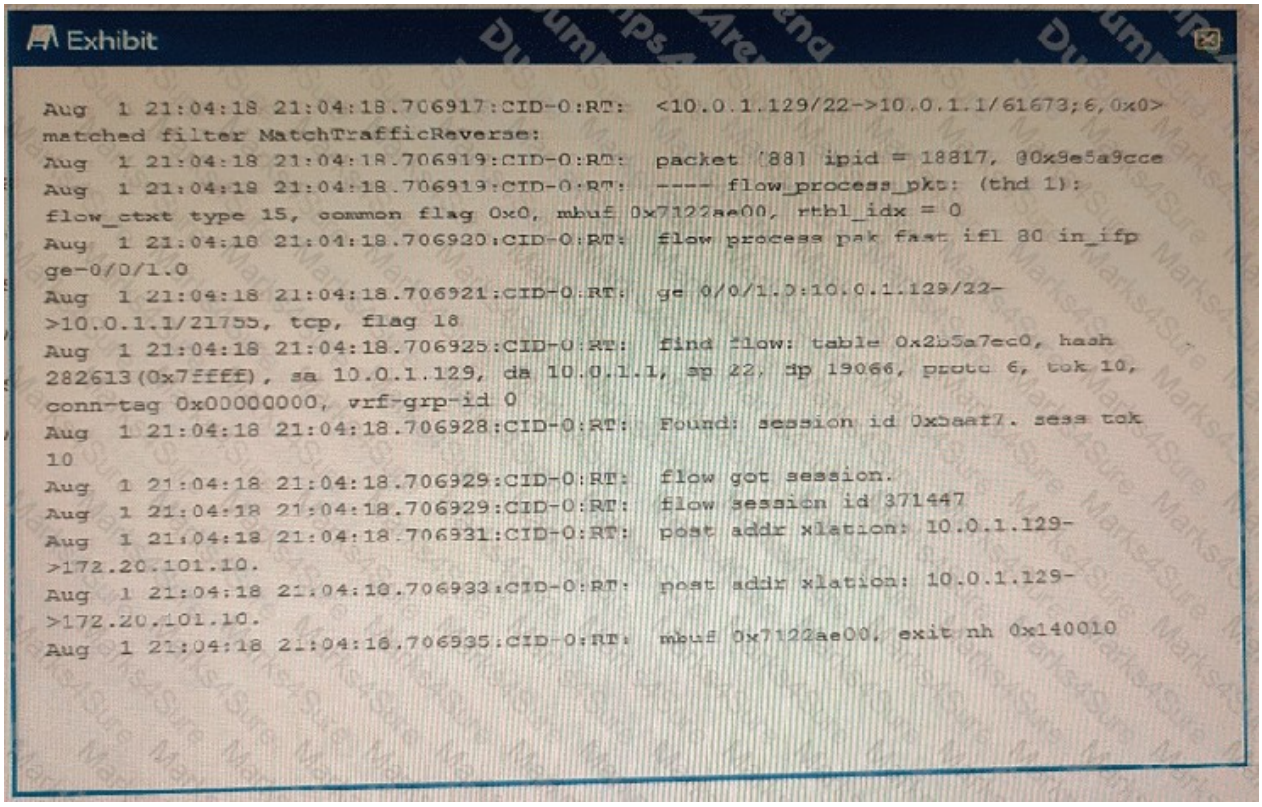
ANSWER: C

Explanation:

Option C is the correct selection because it represents the required configuration correction for restoring reachability between the SRX device and its gateway on the 192.168.100.0 network. For an SRX interface to communicate on a directly connected IPv4 subnet, the interface must have an address and prefix length that place it in that same subnet, and the applicable interface and routing configuration must be committed. Once the interface has a valid address on the gateway's subnet, Junos installs the directly connected route automatically. The SRX can then send ARP requests for 192.168.100.1 and transmit ICMP echo packets through the correct Layer 3 interface. This is the fundamental prerequisite for a successful ping to a directly connected default gateway; no static route is required merely to reach a neighbor on the same IP subnet. Juniper's interface-address configuration documentation describes assigning IPv4 addresses to logical units, and its routing documentation explains that direct routes are created from configured interface addresses. See Juniper interface configuration documentation and Juniper routing documentation.

QUESTION NO: 16

Exhibit



You are using trace options to verity NAT session information on your SRX Series device

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

A. This packet is part of an existing session.

B. The SRX device is changing the source address on this packet.

C. This is the first packet in the session

D. The SRX device is changing the destination address on this packet from 10.0.1.1 to 172.20.101.10.

ANSWER: A D

Explanation:

This packet is part of an existing session. The flow trace indicates that the packet is matched against session state already maintained by the SRX device, rather than causing creation of a new flow-session entry. SRX processing is stateful: after the initial packet establishes a session, later packets are associated with that session and use its stored forwarding and NAT information. This allows the device to apply the same translation consistently for the duration of the flow.

The SRX device is changing the destination address on this packet from 10.0.1.1 to 172.20.101.10. The NAT trace explicitly records the translation in the direction `10.0.1.1 -> 172.20.101.10`. In this packet's processing direction, the original destination is rewritten to the internal translated destination, which is destination NAT behavior. The existing session retains this NAT mapping, so subsequent matching traffic can be translated according to the session's established state. Juniper's flow-trace documentation describes using flow traceoptions to inspect session lookup and packet-processing details, while its NAT documentation describes destination translation and the session-based handling of NAT traffic. See Juniper flow tracing documentation and Juniper NAT documentation.