

DUMPS ARENA

Security, Professional (JNCIP-SEC)

Juniper JN0-637

Version Demo

Total Demo Questions: 13

Total Premium Questions: 132

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Security Policies	17
Topic 2, NAT	27
Topic 3, IPsec VPNs	27
Topic 4, Advanced Security Features	36
Topic 5, High Availability	17
Topic 6, Troubleshooting	6
Topic 7, Mix Questions	2
Total	132

QUESTION NO: 1

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: C

Explanation:

Use the address-persistent parameter. In a Junos source NAT pool, address persistence causes the device to consistently select the same translated public address from that pool for a particular internal source address. This is appropriate when an application establishes several separate sessions to a server but requires its externally visible source address to remain stable across those sessions. The device therefore retains the source-address-to-translated-address association rather than selecting a different available address from the configured pool for later connections. This behavior is particularly important for applications or services that associate a client's state, authorization, or transactions with the observed translated IP address. Address persistence applies to the translated address selected from the pool; normal session processing still creates the individual session and port translations required for each connection. Configure the setting under the relevant source NAT pool so that all source NAT rules using that pool receive the consistent-address behavior. Juniper documents this pool-level setting in its [address-persistent CLI reference](#) and describes the related source NAT pool behavior in the [Junos source NAT documentation](#).

QUESTION NO: 2

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: D

Explanation:

One security profile can be applied to multiple logical systems. A security profile is a reusable configuration object for defining resource limits for user logical systems. After the profile is configured, it can be associated with more than one logical system, allowing a common set of limits to be consistently enforced across logical systems without duplicating the resource-limit configuration in each logical-system hierarchy. This model is particularly useful in multitenant deployments, where several tenants require the same allocation policy and administrators need a centralized way to maintain that policy. Changing the limits in the shared profile updates the policy definition used by the logical systems to which that profile is assigned, making profile-based resource control more manageable than independently defining equivalent limits per tenant. Junos logical systems are configured under the `logical-systems` hierarchy, and security-profile configuration is part of the logical-system resource-control model. See Juniper's [logical-systems CLI hierarchy reference](#) and the [Juniper Junos documentation portal](#) for the applicable platform and release documentation.

QUESTION NO: 3

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 is the SRX feature designed to correct DNS address information when DNS traffic crosses a NAT boundary. In this situation, the DNS reply contains Web1's RFC 1918 private address, but Host A must receive the server's mapped public address to establish a connection through the SRX. DNS doctoring inspects the DNS response and rewrites the applicable address record, substituting the public, statically mapped address for the internal address before the reply reaches the client. The client can therefore continue to use the same hostname, **web.acme.com**, while receiving an address that is valid and routable from its location. This avoids requiring clients to know the server's internal addressing and preserves the expected name-to-public-address mapping for access through the firewall. On Junos OS, DNS doctoring is used with static NAT and is applied to DNS responses traversing the relevant security zones. The SRX can then maintain the address translation for the subsequent web session consistently with the DNS response it provided. Juniper documents this behavior under DNS doctoring and static NAT configuration guidance: [Juniper DNS Doctoring documentation](#) and [Juniper Static NAT documentation](#).

QUESTION NO: 4

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:

Softwire concentrator is correct because an SRX Series device supporting Dual-Stack Lite (DS-Lite) operates at the provider edge as the Address Family Transition Router (AFTR). In this role, it terminates the IPv4-in-IPv6 softwire tunnels established from customer-side Basic Bridging Broadband (B4) elements. The SRX decapsulates IPv4 subscriber traffic carried across the IPv6 access network, then forwards that traffic toward IPv4 destinations. It also provides the IPv4 address-and-port translation function required for DS-Lite, allowing multiple subscribers to share provider IPv4 address resources while their access network can remain IPv6-based. This centralized softwire termination and translation function is precisely what is meant by a softwire concentrator. Junos documentation describes DS-Lite as a transition mechanism in which the AFTR encapsulation/decapsulation endpoint resides in the service-provider network and supports transport of IPv4 packets over an IPv6 infrastructure. On SRX platforms, this deployment combines the tunneling role with security-policy and NAT capabilities, enabling controlled subscriber access to IPv4 services during IPv6 migration. See Juniper's [Dual-Stack Lite documentation](#) and the IETF's defining [RFC 6333](#).

QUESTION NO: 5

Referring to the exhibit,

```
SRX(ttyp0)
login: User1
Password:
--- JUNOS 22.4R1.9 built 2023-03-24 12:52:33 UTC
User1@SRX:LSYS-1>
```

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: A B

Explanation:

The command prompt in the exhibit identifies the active login context as `User1@SRX:LSYS-1>`. In Junos OS, the logical-system name appended to the prompt confirms that the session is operating in that logical system, so User1 is logged in to logical system LSYS-1. Logical systems virtualize a single physical device into separate administrative and routing environments. A user working in a logical-system context works with that logical system's configuration hierarchy and operational context rather than the configuration of a different logical system. Therefore, User1 has access to the configuration specific to their assigned logical system. This separation is central to logical-system administration: configuration, interfaces assigned to the logical system, protocols, and associated operational data are scoped to the relevant logical system. The precise commands a user can issue within that context still depend on the Junos login class and permissions granted to the account, but the prompt establishes the logical-system scope of this session. Juniper documents that logical systems provide independent administrative domains on one device and that users can log in directly to a

configured logical system. See the [Juniper Logical Systems Overview](#) and the [Juniper User Access and Authentication documentation](#).

QUESTION NO: 6

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: A C

Explanation:

In a multinode HA deployment, signal routes provide an additional mechanism for nodes to communicate HA state and to avoid an unnecessary role transition when the interchassis link (ICL) is unavailable. When the ICL fails while a backup signal route IP is configured, the node that is already active retains its active role. This preserves forwarding continuity and prevents the ICL failure alone from causing an avoidable active/backup switchover.

The backup node responds to the ICL-loss condition by changing the route preference toward the peer node to its medium-priority value. That route-preference adjustment is part of the signal-route behavior used to maintain deterministic node-role handling during the failure. Together, retaining the existing active role and applying the backup node's medium-priority routing behavior ensure that the cluster can distinguish an ICL failure from a condition that genuinely requires a role change. Juniper's multinode HA documentation describes the role of signal routes and ICL connectivity in maintaining cluster availability; see the [Juniper multinode high availability documentation](#) and the [Juniper chassis cluster documentation](#).

QUESTION NO: 7

You require a hub-and-spoke VPN design in which spokes initially communicate through a hub. When ongoing traffic between two spokes justifies a more efficient path, the spokes must dynamically establish a direct IPsec shortcut tunnel.

Which VPN technology satisfies this requirement?

- A. AutoVPN
- B. ADVPN
- C. Policy-based VPN
- D. DS-Lite

ANSWER: B


Explanation:

ADVPN is designed to establish dynamic spoke-to-spoke shortcut tunnels while retaining hub-and-spoke connectivity. Initial spoke-to-spoke traffic can use the hub path while the ADVPN control plane enables the direct shortcut VPN to be created. Subsequent eligible traffic can then use the direct tunnel.

AutoVPN supports scalable hub-and-spoke VPN deployment, but spoke-to-spoke traffic continues through the hub. A conventional route-based VPN does not dynamically create spoke-to-spoke shortcuts by itself.

QUESTION NO: 8

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: A B D

Explanation:


The required configuration uses transparent bridging so that the SRX forwards traffic at Layer 2 among hosts in the Local network without becoming their Layer 3 gateway. Configuring `set protocols l2-learning global-mode transparent-bridge` places Layer 2 learning in transparent-bridge mode, allowing the SRX to learn MAC addresses and bridge frames for the local segment. Because the devices are members of VLAN 10, `set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan-members 10` makes the connected interface a switched Layer 2 port in that VLAN. In the complete deployment, each relevant local-facing switched port would be configured as a member of the same VLAN.

Assigning the connected interface to the Local security zone with `set security zones security-zone Local interfaces ge-0/0/1.0` is also essential. A security zone associates the interface with the SRX security-policy framework, enabling policies for traffic entering and leaving the Local zone, including intra-zone flows when such policies are configured. Together, the Layer 2 VLAN membership, transparent bridging mode, and zone membership provide local switched connectivity while retaining the ability to enforce SRX security policy on the traffic. Juniper documents Layer 2

transparent bridging and its L2-learning configuration under [Layer 2 Transparent Mode](#) and describes security-zone interface assignment in the [Security Zones documentation](#).

QUESTION NO: 9

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: A C

Explanation:

“You should add the host-inbound-traffic system-service ike statement to the Untrust zone.” is correct because the remote VPN client initiates IKE negotiation toward the SRX device’s Internet-facing address. On an SRX, traffic destined to the device itself is controlled by the host-inbound-traffic configuration of the zone containing the receiving interface. Enabling the `ike` system service in the Untrust zone permits the SRX to receive and process the IKE exchanges required to establish the IPsec security association.

“You should add the host-inbound-traffic system-service tcp-encap statement to the Untrust zone” is also correct. TCP encapsulation carries IPsec-related traffic over TCP when that transport is configured for the remote-access VPN. Since the remote client reaches the SRX through its external, untrusted interface, the SRX must permit the `tcp-encap` host-inbound system service in that same Untrust zone. These are control-plane services directed at the SRX itself, rather than transit traffic governed by interzone security policies. Junos documents host-inbound system services under security zones and identifies IKE as an available service: [Juniper host-inbound-traffic CLI reference](#). See also [Juniper IPsec VPN documentation](#).

QUESTION NO: 10

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: A B C

Explanation:

To automate switch-port isolation for a host identified as infected, the SRX Series device must be enrolled with Juniper ATP Cloud. This enrollment allows the firewall to submit suspicious files for analysis and receive malware verdicts and infected-host intelligence that can initiate an enforcement workflow. Juniper ATP Cloud is designed to provide this threat intelligence to participating enforcement points; see the [Juniper Advanced Threat Prevention](#) overview.

Security Director with Policy Enforcer is also required because Policy Enforcer centrally coordinates the infected-host response. It translates the threat event into enforcement actions across the relevant security and network infrastructure, including quarantining the endpoint by disabling or restricting its connected switch port. This policy-driven orchestration is the mechanism that connects detection with automated containment; see the [Juniper Security Director](#) product information.

Where the access switch shown in the exhibit is a supported third-party switch rather than a Juniper switch managed through native integration, a third-party connector is required. The connector supplies the device-specific integration that enables Policy Enforcer to communicate the containment instruction to that switch and apply the port-level isolation action for the infected host.

Exhibit:

```
user@peer1> show chassis high-availability information
Node failure codes:
HW Hardware monitoring LB Loopback monitoring
MB Mbuf monitoring SP SPU monitoring
CS 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: 0
Current State: ONLINE
Peer Information:
Peer Id: 2
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: ACTIVE
Activeness Priority: 200
Preemption: ENABLED
Process Packet In Backup State: NO
```

```
Control Plane State: READY
System Integrity Check: N/A
Failure Events: NONE
Peer Information:
Peer Id: 2
Status : BACKUP
Health Status: HEALTHY
Failover Readiness: READY
```

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:**

If SRG1 moves to peer 2, peer 1 will forward packets sent to the SRG1 interfaces. In a chassis cluster, a redundancy group has one active peer that owns and processes traffic for the interfaces associated with that group. After SRG1 fails over to peer 2, peer 2 becomes the active processing peer for those interfaces. Traffic that arrives at an SRG1 interface on peer 1 is therefore sent across the cluster fabric path to peer 2, where it can be processed by the active instance. This forwarding behavior maintains reachability when traffic enters through the non-active peer and is a key part of chassis-cluster data-plane operation. The cluster's control and fabric connectivity allow the peers to synchronize the redundancy-group role and carry the required inter-node traffic during and after failover. Juniper describes redundancy groups as the mechanism that determines active and secondary node roles for clustered interfaces and services; see the [Juniper chassis cluster redundancy group documentation](#). Additional chassis-cluster architecture and traffic-handling information is available in the [Juniper chassis cluster overview](#).

QUESTION NO: 12

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.

```
Mar 7 02:38:15 02:38:15.353921:CID=0:THREAD_ID=01:RT: <192.168.2.1/54882->192.168.1.1/179;6,0x0 > matched filter ibgp-traffic:
...
Mar 7 02:38:15 02:38:15.353933:CID=0:THREAD_ID=01:RT: ge-0/0/3.0:192.168.2.1/54882->192.168.1.1/179, tcp, flag 2 syn
Mar 7 02:38:15 02:38:15.353935:CID=0:THREAD_ID=01:RT: find flow: table 0x206a60a0, hash 6145(0xffff), sa 192.168.2.1, da 192.168.1.1, sp 54882, dp 179, proto 6, tok 9, conn-tag 0x00000000
Mar 7 02:38:15 02:38:15.353938:CID=0:THREAD_ID=01:RT: no session found, start first path. in_tunnel = 0x0, from_op_flag = 0
Mar 7 02:38:15 02:38:15.353941:CID=0:THREAD_ID=01:RT: flow_first_create_session
...
Mar 7 02:38:15 02:38:15.353964:CID=0:THREAD_ID=01:RT: Doing DESTINATION addr route-lookup
Mar 7 02:38:15 02:38:15.353971:CID=0:THREAD_ID=01:RT: flow_ipv4_rt_lookup success 192.168.1.1, ifl 0x47, oifl 0x0
Mar 7 02:38:15 02:38:15.353975:CID=0:THREAD_ID=01:RT: Changing out-ifp from .local..0 to lo0.0 for dst: 192.168.1.1 in vr_id:0
Mar 7 02:38:15 02:38:15.353976:CID=0:THREAD_ID=01:RT: routed (x_dst_ip 192.168.1.1) from untrust (ge-0/0/3.0 in 0) to lo0.0, Next-hop: 192.168.1.1
Mar 7 02:38:15 02:38:15.353978:CID=0:THREAD_ID=01:RT: flow_first_policy_search: policy search from zone untrust-> zone trust (0x0,0xd66200b3,0xb3)
Mar 7 02:38:15 02:38:15.353986:CID=0:THREAD_ID=01:RT: Policy lookup: ways 0 zone(5:global) -> zone(5:global) scope:0
...
Mar 7 02:38:15 02:38:15.354000:CID=0:THREAD_ID=01:RT: permitted by policy allow-bgp(6)
Mar 7 02:38:15 02:38:15.354048:CID=0:THREAD_ID=01:RT: flow_first_final_check: in 0/3.0>, out
Mar 7 02:38:15 02:38:15.354050:CID=0:THREAD_ID=01:RT: In flow_first_complete_session
Mar 7 02:38:15 02:38:15.354051:CID=0:THREAD_ID=01:RT: flow_first_complete_session, pak_ptr: 0x2c5fd40, nsp: 0x2a140340, in_tunnel: 0x0
...
Mar 7 02:38:15 02:38:15.353978:CID=0:THREAD_ID=01:RT: flow_first_policy_search: policy search from zone untrust-> zone trust (0x0,0xd66200b3,0xb3)
Mar 7 02:38:15 02:38:15.353986:CID=0:THREAD_ID=01:RT: Policy lookup: ways 0 zone(5:global) -> zone(5:global) scope:0
...
Mar 7 02:38:15 02:38:15.354000:CID=0:THREAD_ID=01:RT: permitted by policy allow-bgp(6)
Mar 7 02:38:15 02:38:15.354048:CID=0:THREAD_ID=01:RT: flow_first_final_check: in 0/3.0>, out
Mar 7 02:38:15 02:38:15.354050:CID=0:THREAD_ID=01:RT: In flow_first_complete_session
Mar 7 02:38:15 02:38:15.354051:CID=0:THREAD_ID=01:RT: flow_first_complete_session, pak_ptr: 0x2c5fd40, nsp: 0x2a140340, in_tunnel: 0x0
...
Mar 7 02:38:15 02:38:15.353978:CID=0:THREAD_ID=01:RT: flow_first_policy_search: policy search from zone untrust-> zone trust (0x0,0xd66200b3,0xb3)
Mar 7 02:38:15 02:38:15.353986:CID=0:THREAD_ID=01:RT: Policy lookup: ways 0 zone(5:global) -> zone(5:global) scope:0
...
Mar 7 02:38:15 02:38:15.354000:CID=0:THREAD_ID=01:RT: permitted by policy allow-bgp(6)
Mar 7 02:38:15 02:38:15.354048:CID=0:THREAD_ID=01:RT: flow_first_final_check: in 0/3.0>, out
Mar 7 02:38:15 02:38:15.354050:CID=0:THREAD_ID=01:RT: In flow_first_complete_session
Mar 7 02:38:15 02:38:15.354051:CID=0:THREAD_ID=01:RT: flow_first_complete_session, pak_ptr: 0x2c5fd40, nsp: 0x2a140340, in_tunnel: 0x0
...
Mar 7 02:38:15 02:38:15.354055:CID=0:THREAD_ID=01:RT: Session (id:20395) created for first pak 2
Mar 7 02:38:15 02:38:15.354073:CID=0:THREAD_ID=01:RT: flow_first_in_dst_nat: in . out A> dst_addr 192.168.1.1, sp 54882, dp 179
Mar 7 02:38:15 02:38:15.354075:CID=0:THREAD_ID=01:RT: chose interface lo0.0 as incoming nat if.
Mar 7 02:38:15 02:38:15.354075:CID=0:THREAD_ID=01:RT: packet dropped: for self but not interested
Mar 7 02:38:15 02:38:15.354076:CID=0:THREAD_ID=01:RT: packet dropped: for self but not interested.
Mar 7 02:38:15 02:38:15.354079:CID=0:THREAD_ID=01:RT: flow_first_install_session: Loopback session processing aborted
Mar 7 02:38:15 02:38:15.354080:CID=0:THREAD_ID=01:RT: first path session installation failed
Mar 7 02:38:15 02:38:15.354081:CID=0:THREAD_ID=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:

Add BGP to the lo0 host-inbound-traffic configuration. An IBGP peer establishes its transport session using TCP port 179, and when the BGP neighbor is addressed through the device's loopback interface, that traffic is destined to the SRX Routing Engine rather than merely transiting the firewall. Traffic destined to the device itself is controlled by the security zone's `host-inbound-traffic` configuration. The trace output indicates that BGP connection traffic reaches the device but the control-plane service is not permitted to accept it, preventing the TCP/BGP session from completing. Enabling the `bgp` system service under host-inbound traffic for the security zone associated with the loopback interface authorizes the SRX to receive BGP packets for that local interface. For example, the required configuration is conceptually `set security zones security-zone <zone-name> host-inbound-traffic system-services bgp`, with the zone selected according to the interface's zone assignment. Juniper documents host-inbound traffic as the mechanism for permitting protocol traffic directed to an SRX device, and lists BGP as a supported system service. See Juniper's [security zone and host-inbound traffic documentation](#) and the [system-services host-inbound-traffic CLI reference](#).

QUESTION NO: 13

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: A C

Explanation:

"You are unable to apply stateful security features to traffic that is switched between the two interfaces." is correct because traffic switched locally at Layer 2 within the VLAN is handled through Layer 2 learning and forwarding. Stateful SRX security processing depends on an ingress and egress security-zone context so that flow processing can create sessions and evaluate policies. With these Layer 2 interfaces not assigned to a security zone, there is no security-zone context through which to apply stateful policy features to the switched traffic.

"The interfaces will not forward traffic by default." is also correct. On an SRX device operating with the default Layer 2 learning behavior, Layer 2 interfaces must be associated with a security zone for traffic to be forwarded through the security forwarding architecture. Merely placing the interfaces in the same VLAN establishes the VLAN membership; it does not create the required zone association for forwarding. Assigning the interfaces appropriately to a security zone provides the forwarding context and enables the SRX to apply its security services where applicable. Juniper's security documentation describes the relationship between interfaces, security zones, and flow-based security processing: [Juniper Security Services Documentation](#) and [Juniper Security Zone Configuration Documentation](#).