PassQueen
a.k.a dumpsqueen.com

Security, Specialist (JNCIS-SEC)

Juniper JN0-335

Version Demo

Total Demo Questions: 12

Total Premium Questions: 122

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Security Policies	26
Topic 2, Network Address Translation	6
Topic 3, IPsec VPNs	7
Topic 4, High Availability	21
Topic 5, Intrusion Prevention System	7
Topic 6, AppSecure	21
Topic 7, Unified Threat Management	3
Topic 8, Security Monitoring and Reporting	20
Topic 9, Mix Questions	11
Total	122

QUESTION NO: 1

You configure a destination NAT rule to translate traffic destined to a public address to an internal web server.

Which destination address must the security policy match to permit the translated traffic?

A. The translated destination address.

B. The original public destination address.

C. The source address after source NAT processing.

D. The address assigned to the egress interface.

ANSWER: A

Explanation:

The security policy must match the translated destination address. Destination NAT is performed before the security-policy lookup for inbound traffic. After the SRX Series device translates the public destination address to the private address of the protected server, security-policy evaluation uses the translated destination address when determining whether to permit the session.

This behavior allows the security policy to be written using the actual address of the internal resource rather than its external, pretranslation address. The destination NAT rule remains responsible for matching the public address and applying the translation. See Juniper's destination NAT documentation.

QUESTION NO: 2

You are implementing an SRX Series device at a branch office that has low bandwidth and also uses a cloud-based VoIP solution with an outbound policy that permits all traffic.

Which service would you implement at your edge device to prioritize VoIP traffic in this scenario?

A. AppFW

B. SIP ALG

C. AppQoE

D. AppQoS

ANSWER: D

Explanation:

AppQoS is the appropriate service because it provides application-aware quality-of-service handling on supported SRX Series devices. In a low-bandwidth branch environment, simply allowing all outbound traffic does not ensure that latency-sensitive voice packets receive timely treatment when the WAN link becomes congested. AppQoS can identify traffic by application and apply an application traffic-control profile, allowing the SRX device to prioritize the cloud-based VoIP application relative to less time-sensitive traffic. This helps protect voice quality by giving the identified voice flow preferential handling during periods of contention, reducing the likelihood of delay, jitter, and packet loss that can make calls unusable. The policy can remain broadly permissive for outbound access while AppQoS supplies the required differentiation and prioritization at the edge. Juniper documents AppQoS as an application-based QoS feature that classifies traffic using application identification and applies traffic-control policies to it. See Juniper's Application Quality of Service overview and AppQoS configuration documentation.

QUESTION NO: 3

Which two statements are correct about screen options on an SRX Series device? (Choose two.)

A. Screen options are configured under a security zone.

B. Screen options inspect traffic entering interfaces in the associated zone.

C. Screen options are configured under an individual security policy.

D. Screen options perform URL categorization for HTTP traffic.

ANSWER: A B

Explanation:

Screen options are configured under a security zone and inspect traffic entering interfaces assigned to that zone. This makes the zone the ingress security context for screen processing. For example, screen options configured under an untrust zone can detect and protect against malformed or suspicious packets arriving from an external network.

Screen options are intended to detect network-layer and transport-layer attacks and anomalies, such as IP spoofing, ICMP floods, TCP SYN floods, and invalid TCP flags. When traffic matches an enabled screen condition, the SRX Series device can drop the offending traffic before it establishes a normal flow session. See Juniper's security screens documentation.

QUESTION NO: 4

You want to require a new Diffie-Hellman key exchange when the IPsec security association is rekeyed.

Which IPsec feature should you configure?

A. perfect forward secrecy

B. dead peer detection

C. NAT traversal

D. proxy identity

ANSWER: A

Explanation:

The correct answer is **perfect forward secrecy**. Perfect forward secrecy (PFS) requires the VPN peers to perform an additional Diffie-Hellman exchange when new Phase 2 IPsec SAs are created. The new key material is therefore not derived solely from the existing IKE SA.

PFS improves key separation between IPsec security associations. Both VPN peers must use compatible PFS settings for the Phase 2 negotiation to succeed. See Juniper's IPsec VPN documentation.

QUESTION NO: 5

You are troubleshooting a flow-processing problem and need to capture detailed packet-processing information for selected traffic on an SRX Series device.

Which feature should you configure?

A. security flow traceoptions

B. security log stream

C. services accounting

D. routing-options forwarding-table export

ANSWER: A

Explanation:

You should configure **security flow traceoptions**. Flow traceoptions provide detailed tracing for SRX packet processing, including session creation, policy lookup, NAT processing, routing decisions, and session teardown information. A packet filter can be used to restrict tracing to the relevant addresses, ports, or protocols.

Using a packet filter is important because unrestricted flow tracing can generate a large volume of log data and consume system resources. Filtering the trace to the affected traffic provides focused diagnostic information while minimizing the operational impact on the device. See Juniper's security flow traceoptions documentation.

QUESTION NO: 6

Which two statements are true about the fab interface in a chassis cluster? (Choose two.)

A. The fab link does not support fragmentation.

B. The physical interface for the fab link must be specified in the configuration.

C. The fab link supports traditional interface features.

D. The Junos OS supports only one fab link.

ANSWER: A B

Explanation:

The fab link does not support fragmentation. A chassis-cluster fabric link transports session synchronization and data-plane traffic between cluster nodes, and traffic sent across it must fit within the applicable fabric-link MTU. It is therefore important to use compatible link characteristics and an appropriate MTU on the physical interfaces allocated to the fabric connection. The physical interface for the fab link must be specified in the configuration. During chassis-cluster configuration, Junos OS requires the administrator to identify the physical interface assigned as each node's fabric link. That assignment establishes the underlying connection used by the internal fabric interface and enables the nodes to exchange the required clustered traffic. Juniper's chassis-cluster documentation describes configuring fabric links by naming the physical interface for each node, rather than treating the fabric connection as an automatically selected network port. This explicit mapping also makes the fabric-link design, cabling, and failover behavior deterministic. For configuration syntax and fabric-link requirements, see Juniper's fabric interface configuration documentation and its fabric interface overview.

QUESTION NO: 7

While working on an SRX firewall, you execute the show security policies policy-name detail command.

Which function does this command accomplish?

A. It displays details about the default security policy.

B. It identifies the different custom policies enabled.

C. It shows the system log files for the local SRX Series device.

D. It shows policy counters for a configured policy.

ANSWER: D

Explanation:

The `show security policies policy-name <name> detail` command displays detailed operational information for the specified configured security policy, including its policy counters. These counters provide visibility into traffic that has matched the policy, such as packet and byte counts, and session statistics including sessions created, denied, and closed.

This makes the command useful when validating that traffic is reaching the intended rule and when troubleshooting policy enforcement or session behavior on an SRX Series device. The policy name limits the output to the particular policy of interest, while the `detail` keyword requests the extended information that includes these statistics. Depending on the Junos release and policy configuration, the detailed output can also show the policy's match criteria, action, logging settings, and other operational attributes alongside the counters. Juniper documents the command syntax and its detailed output in the show security policies CLI reference. For context on how SRX security policies process and control traffic between security zones, see Juniper's security policy configuration documentation.

QUESTION NO: 8

You want to deploy a virtualized SRX in your environment.

In this scenario, why would you use a vSRX instead of a cSRX? (Choose two.)

A. The vSRX supports Layer 2 and Layer 3 configurations.

B. Only the vSRX provides clustering.

C. The vSRX has faster boot times.

D. Only the vSRX provides NAT, IPS, and UTM services

ANSWER: A B

Explanation:

The vSRX supports Layer 2 and Layer 3 configurations because it can operate in deployments requiring Ethernet switching behavior as well as routed security services. This makes it appropriate where the virtual firewall must participate in Layer 2 designs, such as transparent or bridged network segments, in addition to conventional Layer 3 routing and security-policy enforcement. The vSRX also supports chassis clustering for high availability. A chassis cluster combines two vSRX instances into a resilient firewall pair, using control-link and fabric-link connectivity to maintain state and provide failover capabilities. This is a significant selection factor when the virtual firewall itself must provide Junos OS chassis-cluster-based redundancy rather than relying solely on external platform or orchestration resiliency. Juniper documents vSRX Layer 2 capabilities and its supported deployment features in the vSRX Features documentation. Juniper's cSRX Features documentation is also useful for comparing the containerized firewall's supported feature set and deployment model.

QUESTION NO: 9

Which solution enables you to create security policies that include user and group information?

A. JIMS

B. ATP Appliance

C. Network Director

D. NETCONF

ANSWER: A

Explanation:

JIMS is correct because Juniper Identity Management Service supplies identity context that can be used in access-control decisions. It gathers, correlates, and maintains user, group, and device identity information from enterprise identity sources, including Microsoft Active Directory. JIMS then makes that information available to supported SRX Series security devices so that policy evaluation can associate network activity with an authenticated user and that user's directory-group membership, rather than relying only on IP addresses, ports, and protocols.

This identity-aware capability is particularly useful when users move among devices or receive dynamically assigned addresses. Administrators can define and enforce security policy based on organizational identity attributes, such as allowing

a particular directory group to reach a protected application or network resource. The firewall can therefore apply consistent access controls aligned with the organization's existing user and group structure. Juniper describes JIMS as an identity-management solution for providing user, device, and group information to SRX platforms, enabling identity-based policy enforcement. See the Juniper Identity Management Service overview and Juniper's user firewall policy documentation.

QUESTION NO: 10

Which two statements are true about the vSRX? (Choose two.)

A. It does not have VMXNET3 vNIC support.

B. It has VMXNET3 vNIC support.

C. UNIX is the base OS.

D. Linux is the base OS.

ANSWER: B D

Explanation:

The vSRX Virtual Firewall supports VMXNET3 virtual network interface cards when deployed in supported VMware environments. VMXNET3 is VMware's paravirtualized, high-performance virtual NIC type and is used to provide efficient packet I/O for virtual appliances. When configuring a vSRX on ESXi, Juniper documentation identifies VMXNET3 as a supported adapter type, subject to the applicable vSRX release and deployment requirements.

Linux is the base operating-system environment for vSRX. Unlike physical SRX Series devices, which use the traditional Junos OS architecture, vSRX is implemented as a virtual security appliance optimized for operation on virtualization platforms. Its architecture uses a Linux-based host environment together with Junos OS components to provide firewall, VPN, NAT, and other security services in virtualized deployments. This design enables vSRX to run on supported hypervisors and cloud infrastructure while retaining Junos security-policy and operational capabilities. See Juniper's vSRX overview documentation and the vSRX VMware installation documentation for platform and virtual-interface details.

QUESTION NO: 11

Which two statements are true about application identification? (Choose two.)

A. Application identification can identity nested applications that are within Layer 7.

B. Application identification cannot identify nested applications that are within Layer 7.

C. Application signatures are the same as IDP signatures.

D. Application signatures are not the same as IDP signatures.

ANSWER: A D

Explanation:

Application identification can identity nested applications that are within Layer 7. Juniper Application Identification (AppID) examines traffic beyond basic Layer 3/Layer 4 attributes and uses predefined application signatures to recognize application behavior and protocol characteristics. This enables the device to identify applications embedded within, tunneled through, or otherwise carried by another application protocol at the application layer. The resulting application information can then be used by security policies to control traffic with greater precision than port- or protocol-only matching.

Application signatures are not the same as IDP signatures. AppID signatures are intended to classify and identify applications so that application-aware security policies can match the traffic. IDP signatures serve a different security function: they detect known attack patterns, protocol anomalies, and exploit attempts as part of intrusion detection and prevention processing. Although both capabilities inspect traffic and rely on signature-based mechanisms, their signature sets, purposes, and policy actions are distinct. Juniper documents the predefined AppID signature capability and its role in

recognizing applications at Application Identification Predefined Signatures. Juniper's IDP documentation further describes IDP as a threat-detection and prevention service at Intrusion Detection and Prevention Overview.

QUESTION NO: 12

Which two statements are correct about chassis clustering? (Choose two.)

A. The node ID value ranges from 1 to 255.

B. The node ID is used to identify each device in the chassis cluster.

C. A system reboot is required to activate changes to the cluster.

D. The cluster ID is used to identify each device in the chassis cluster.

ANSWER: B C

Explanation:

"The node ID is used to identify each device in the chassis cluster" is correct because a chassis cluster is composed of two member devices, designated node 0 and node 1. The node ID distinguishes the individual cluster members and is used in operational output, configuration context, and cluster-management functions. Junos assigns these node identifiers as part of chassis-cluster configuration, allowing the two devices to operate together while retaining an unambiguous per-device identity.

"A system reboot is required to activate changes to the cluster" is also correct in the context of configuring or changing fundamental chassis-cluster settings. A device must reboot to enter chassis-cluster mode after the cluster configuration is committed. Likewise, changing cluster identity parameters requires the cluster members to reboot for the changed cluster configuration to take effect. This reboot allows Junos to initialize the chassis-cluster control and fabric infrastructure using the active cluster settings. Juniper's chassis-cluster documentation describes configuring the cluster ID and node ID and then rebooting the devices to complete cluster formation. See Juniper: Configure a Chassis Cluster and Juniper: cluster-id CLI reference.