DUMPSARENA

Security, Associate (JNCIA-SEC)

Juniper JN0-232

Version Demo

Total Demo Questions: 12

Total Premium Questions: 126

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Security Concepts	6
Topic 2, Junos OS Security	32
Topic 3, Security Policies	37
Topic 4, Security Zones	21
Topic 5, Network Address Translation (NAT)	26
Topic 6, IPsec VPNs	4
Total	126

QUESTION NO: 1

You are deploying an internal web server using static NAT. The server address 10.10.20.10 is mapped to public address 198.51.100.10.

What is the primary benefit of static NAT in this scenario?

A. It provides a persistent one-to-one, bidirectional address mapping.

B. It permits multiple internal hosts to share one public address by translating ports.

C. It translates only the destination address of inbound sessions.

D. It permits inbound traffic without requiring a security policy.

ANSWER: A

Explanation:

Static NAT provides a persistent one-to-one, bidirectional address mapping. The internal server has a consistent public address for inbound access, and traffic initiated from the internal server can use the corresponding translated address. This makes static NAT suitable when a host requires a fixed translated identity.

Source NAT is generally used to translate the source addresses of outbound sessions and is not inherently a persistent bidirectional mapping. Destination NAT is used to translate destinations for inbound services, but static NAT specifically provides the fixed one-to-one relationship described in the scenario. Security policies are still required to permit the traffic.

QUESTION NO: 2

A URL is not found in the local allow list, block list, or local cache during the NextGen Web Filtering process. Which action does the SRX Series Firewall take in this scenario?

A. It allows the URL by default.

B. It sends a TCP reset message to the client.

C. It forwards the URL to the NextGen Web Filtering application in the Juniper cloud.

D. It blocks the URL by default.

ANSWER: C

Explanation:

It forwards the URL to the NextGen Web Filtering application in the Juniper cloud. NextGen Web Filtering uses a lookup sequence designed to make policy decisions efficiently while retaining the ability to classify previously unknown URLs. The SRX Series device first evaluates local exceptions, including configured allow-list and block-list entries. It also checks its locally stored URL categorization cache, which avoids a cloud query for URLs that have already been classified recently. When none of those local sources produces a match, the firewall submits the URL information to Juniper's cloud-based NextGen Web Filtering service. The cloud service supplies the URL's categorization and reputation data, allowing the SRX device to apply the web-filtering profile and associated security policy to the session. This cloud lookup is an essential part of maintaining current classifications without requiring every categorization record to be permanently held on the firewall. Juniper describes NextGen Firewall capabilities and cloud-delivered security intelligence at Juniper SRX Series Next-Generation Firewalls. Additional Juniper documentation on web-filtering and UTM security services is available at Junos Web Filtering Documentation.

QUESTION NO: 3

You have created a series of security policies permitting access to a variety of services. You now want to create a policy that blocks access to all other services for all user groups.

What should you create in this scenario?

A. global security policy

B. Juniper ATP policy

C. IDP policy

D. integrated user firewall policy

ANSWER: A

Explanation:

A **global security policy** is appropriate because it lets an administrator define a policy that is not limited to one particular source-zone-to-destination-zone policy context. It can therefore provide the broad, catch-all control needed after the more specific permit policies have been defined. The policy can match all relevant sources, destinations, and applications, with a deny action, so traffic that does not match an earlier intended service-access rule is explicitly blocked. This approach makes the security intent clear and provides a centrally managed default-deny rule for traffic across the firewall policy domain. Policy ordering must ensure that the specific rules permitting the required services are evaluated before the broad deny rule; otherwise, the catch-all match would prevent the intended access. Junos security policies use ordered matching, so designing a final deny rule is a standard way to enforce least-privilege access while allowing only the explicitly approved services. Juniper documents the scope and use of global policies in its Global Security Policies documentation and describes ordered policy matching in the Security Policy Overview.

QUESTION NO: 4

Referring to the exhibit,

which two statements are correct? (Choose two.)

A. The SRX Series Firewall is performing destination NAT.

B. The SRX Series Firewall is performing source NAT.

C. The SRX Series Firewall is not performing PAT.

D. The SRX Series Firewall is performing PAT.

ANSWER: A C

Explanation:

The SRX Series Firewall is performing destination NAT because the session's original packet direction uses the public destination address 203.0.113.199 on port 22, while the corresponding translated session direction identifies the internal destination as 10.10.101.10 on the same service port. This is the expected flow-session representation for destination translation: traffic is addressed to a public or externally reachable address, and the SRX rewrites that destination to the protected server's private address before forwarding it internally. The session pair therefore provides direct evidence that the firewall is publishing an internal host through a destination NAT rule.

The SRX Series Firewall is not performing PAT because the relevant port remains 22 throughout the displayed translation. Port Address Translation requires a transport-layer port value to be translated, allowing multiple connections to share an address by differentiating flows with ports. Here, only the destination IP address changes between the public and private sides of the session, while SSH continues to use port 22. Juniper documents destination NAT as translating the destination address and, when configured, optionally the destination port; the unchanged port shown here confirms address-only destination translation. See Juniper's destination NAT documentation and flow session documentation.

QUESTION NO: 5

You are troubleshooting traffic traversing the SRX Series Firewall and require detailed information showing how the flow module is handling the traffic.

How would you accomplish this task?

A. Review the flow session table.

B. Review the forwarding table.

C. Enable flow trace options.

D. Enable firewall filters.

ANSWER: C

Explanation:

Enable flow trace options. On an SRX Series device, flow tracing is the purpose-built diagnostic capability for observing how the security flow module processes traffic. Configuring `security flow traceoptions` enables trace output for selected traffic and processing events, allowing an administrator to follow a packet or session through the flow engine. The trace can show important decisions such as session creation and lookup, route lookup, security-policy evaluation, source or destination NAT processing, and whether the packet is permitted, denied, or dropped. This makes it especially useful when the issue is not simply whether a session exists, but why the device handled traffic in a particular way. Traceoptions can be constrained with match criteria, flags, and a trace file so collection is focused on the affected traffic and does not create unnecessary diagnostic output. Juniper documents flow traceoptions under the security flow configuration hierarchy and describes the available tracing flags for flow-processing events. See the Juniper flow traceoptions documentation and the security flow traceoptions CLI reference.

QUESTION NO: 6

Click the Exhibit button.

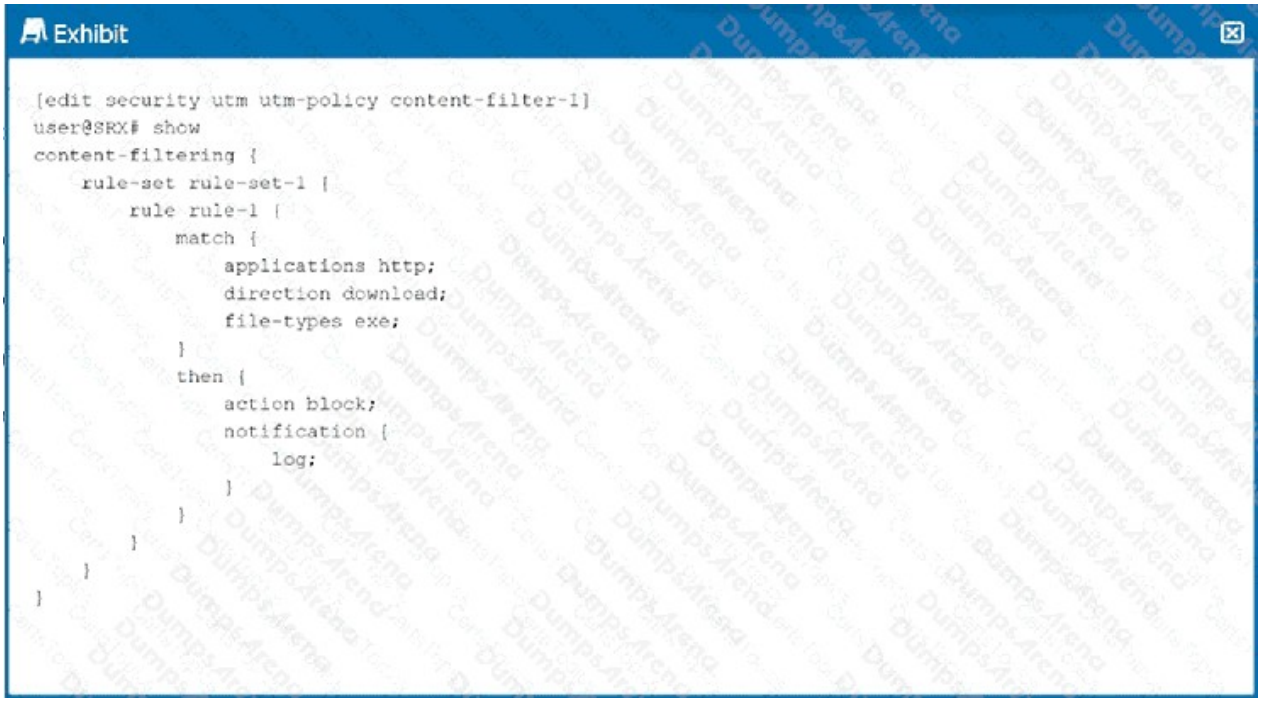

Which two statements are correct about the content filter shown in the exhibit? (Choose two.)

A. .exe files will not be allowed to be uploaded over HTTP.

B. .exe files will not be allowed to be downloaded over HTTP.

C. There will be a notice added to the SRX log file about the file being blocked.

D. There will be an e-mail sent to the user about why the SRX is blocking the file.

ANSWER: B C

Explanation:

The statement ".exe files will not be allowed to be downloaded over HTTP." is correct because the configured match criteria identify HTTP traffic in the download direction and select the executable file type. When a transfer matches those conditions, the configured `block` action prevents delivery of the executable to the client. The policy is therefore evaluated for HTTP downloads, with the file-type match providing the specific control over executable content.

The statement "There will be a notice added to the SRX log file about the file being blocked." is also correct. The content-filter action includes a log notification, so the SRX generates a log event when the content-filter rule is triggered and blocks matching content. This provides an operational record that administrators can use to monitor blocked file transfers and investigate policy events. Junos UTM content filtering supports matching traffic characteristics such as application, direction, and file type, then applying actions and notifications to matching traffic. See Juniper's Content Filtering documentation and the UTM Content Filtering overview.

QUESTION NO: 7

What is the purpose of assigning logical interfaces to separate security zones in Junos OS?

A. to simplify the configuration of network interfaces

B. to manage routing protocols and updates

C. to control traffic between security zones using security policies

D. to enable network monitoring through SNMP

ANSWER: C

Explanation:

Assigning logical interfaces to separate security zones is used to classify network traffic by its security context and to apply stateful security policies when traffic moves between zones. On an SRX device, a zone represents a collection of interfaces with a common trust level or role, such as internal, external, or DMZ connectivity. After interfaces are associated with zones, Junos OS evaluates the source zone and destination zone for a session and uses the applicable security policy to determine whether the traffic is permitted, denied, or subject to configured services such as logging, NAT, or application inspection.

This design lets an administrator create clear security boundaries even where the connected networks use different VLANs, IP subnets, or functional segments. The logical interface is the attachment point through which Junos identifies the traffic's zone membership; the security policy is the mechanism that controls communication across that boundary. Juniper documents security zones as the basis for defining interface membership and applying policies to traffic between zones. See the Juniper security zone configuration documentation and the Juniper security policy configuration documentation.

QUESTION NO: 8

What is a purpose for creating multiple routing instances on an SRX Series Firewall device?

A. to enable network monitoring through SNMP

B. to maintain separation of routing information for security purposes

C. to manage routing protocols and updates

D. to simplify the configuration of network interfaces

ANSWER: B

Explanation:

The correct answer is “to maintain separation of routing information for security purposes.” On an SRX Series Firewall, a routing instance provides a distinct routing table and routing domain. Creating multiple instances lets an administrator isolate route information for separate networks, such as different customers, departments, security zones, or tenant environments. Routes learned or configured in one routing instance are kept independent of routes in another instance unless an administrator deliberately configures controlled route exchange or interconnection. This separation supports segmentation by preventing one routing domain from automatically using, advertising, or being affected by another domain’s route information. In practical deployments, virtual-router and VRF routing instances can associate interfaces with separate routing tables, enabling the SRX to handle overlapping or otherwise independent IP networks while preserving clear administrative and security boundaries. Juniper documents routing instances as constructs that create separate routing tables and identifies virtual-router instances as independent routing domains. More information is available in the Juniper routing instances overview and the instance-type configuration reference.

QUESTION NO: 9

Which two statements are correct about a Juniper Routing Engine? (Choose two.)

A. The Routing Engine is managed by the Packet Forwarding Engine.

B. The Routing Engine manages the Packet Forwarding Engine.

C. The Routing Engine creates the routing and forwarding tables.

D. The Routing Engine is responsible for forwarding transit traffic.

ANSWER: B C

Explanation:

The Routing Engine manages and controls the Packet Forwarding Engine by installing the forwarding information that the Packet Forwarding Engine uses to handle traffic. It is the control-plane component of a Junos device: it runs the Junos OS processes, handles device-management functions, runs routing protocols, and communicates the resulting forwarding state to the forwarding plane. This separation lets the device continue high-speed packet handling in the Packet Forwarding Engine while the Routing Engine performs route computation and control functions.

The Routing Engine creates and maintains the routing table from routing-protocol information, static routes, and other route sources. It then derives the active forwarding information and installs a copy of the forwarding table in the Packet Forwarding Engine. In Ethernet switching contexts, the Routing Engine also provides the control-plane processing that programs the hardware forwarding state. Juniper describes the Routing Engine as maintaining routing and forwarding tables and the Packet Forwarding Engine as using the forwarding table to move packets. See Juniper’s Junos OS architecture overview and its router architecture documentation.

QUESTION NO: 10

Which two statements are correct about Junos OS? (Choose two.)

A. It is a network operating system.

B. It is supported on physical and virtual devices.

C. It is a network management system for Juniper devices.

D. It is a network operating system for IoT devices.

ANSWER: A B

Explanation:

Junos OS is Juniper Networks’ network operating system, providing the software foundation for routing, switching, and security functions across the company’s portfolio. It supplies the control-plane, management-plane, and data-plane capabilities needed to configure devices, establish routing and security policies, monitor operations, and maintain network

services. A consistent Junos OS architecture and command-line interface also help administrators apply common operational practices across different Juniper platforms.

Junos OS is supported on both physical and virtual devices. Physical deployments include many Juniper routing, switching, and SRX security platforms, while virtualized forms of Junos, such as vSRX and vMX, allow Junos-based network and security functions to run in virtualized environments. This portability lets organizations use similar Junos operational concepts whether a function is delivered by dedicated hardware or a virtual network function. Juniper describes Junos OS as its network operating system and documents its use across a broad range of platforms and deployment models. See the Juniper Junos OS overview and the Junos OS documentation portal.

QUESTION NO: 11

Click the Exhibit button.

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

A. The URL matches a predefined Web filtering category.

B. The NextGen Web Filtering type is being used.

C. The SRX firewall does not have an SSL proxy configuration.

D. This is a custom Web filtering block message.

ANSWER: A B

Explanation:

The statement *The URL matches a predefined Web filtering category.* is supported directly by the block-page details. The displayed reason, `BY_PRE_DEFINED`, identifies a block resulting from a configured predefined category action rather than from an individually specified URL. The category shown on the page, `NG_Reference`, is the classification that caused the request to be denied.

The statement *The NextGen Web Filtering type is being used.* is also supported by that category name. Juniper Enhanced Web Filtering, also called Next-Generation Web Filtering, uses category identifiers with the `NG_` prefix. Thus, `NG_Reference` identifies the site as belonging to the Reference category in the Next-Generation categorization service. When an SRX policy applies a predefined action to that category, the web-filtering service produces the displayed block response.

Juniper documents the Enhanced Web Filtering service and its predefined URL categories in its Web Filtering documentation. Configuration and operational details for the UTM web-filtering feature are also available in the Web Filtering Overview.

QUESTION NO: 12

You just made a configuration change to a security policy on your SRX Series Firewall. Your users alert you that an application that uses FTP is no longer working.

```
[edit security policies from-zone Trust to-zone Untrust]
user@SRX# show
inactive: policy ftp {
    match {
        source-address any;
        destination-address Server-1;
        application junos-ftp;
    }
    then {
        permit;
    }
}
policy web-smtp {
    match {
        source-address any;
        destination-address [ Web-1 Mail-Server-1 ];
        application [ junos-http junos-https junos-smtp ];
    }
    then {
        permit;
    }
}
```

Referring to the exhibit, what are two ways to solve this problem? (Choose two.)

A. Enter the rollback 1 command followed by a commit command.

B. Activate the FTP security policy and commit the configuration.

C. Insert the FTP security policy before the web-smtp security policy.

D. Change the destination address in the FTP security policy to any and commit the configuration.

ANSWER: A B

Explanation:

The displayed FTP security policy is configured as inactive. In Junos, an inactive configuration statement remains visible in the candidate and committed configuration, but Junos does not apply it operationally. As a result, traffic that would otherwise match the FTP policy does not receive the policy action and service handling required for the FTP application. Activating the FTP security policy and committing the configuration directly restores the policy to active use while retaining the intended policy definition. Junos provides the `activate` configuration-mode command specifically to re-enable statements that were previously marked inactive.

Using `rollback 1` followed by `commit` is also a valid recovery method when the immediately preceding committed configuration was the known working configuration. The rollback command replaces the current candidate configuration with the selected prior committed version; committing it restores that prior policy state, including an active FTP policy. This approach is especially useful when the recent change might have included more than the single policy-state modification. Juniper documents inactive configuration behavior and activation in the activate command reference and describes restoring previous configurations in the rollback command reference.