DUMPSARENA

Palo Alto Networks Next-Generation Firewall Engineer

Palo Alto Networks NGFW-Engineer

Version Demo

Total Demo Questions: 16

Total Premium Questions: 165

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Plan and Design a Next-Generation Firewall Solution	18
Topic 2, Configure and Deploy a Next-Generation Firewall Solution	130
Topic 3, Operate and Maintain a Next-Generation Firewall Solution	10
Topic 4, Troubleshoot a Next-Generation Firewall Solution	7
Total	165

QUESTION NO: 1

An engineer needs to prevent internal users from uploading executable files to public web-based file-sharing applications while allowing other permitted web traffic.

Which Security profile should be attached to the applicable Security policy rule?

A. Anti-Spyware profile

B. File Blocking profile

C. URL Filtering profile

D. Vulnerability Protection profile

ANSWER: B

Explanation:

A **File Blocking profile** is used to control files based on file type, application, direction, and action. The engineer can configure a rule in the profile to match executable file types for uploads to the applicable web-based file-sharing applications and set the action to block.

File Blocking provides prevention at the point where a file transfer is detected. WildFire Analysis profiles are used to forward eligible files for analysis, whereas Data Filtering profiles are used to detect or control sensitive data patterns. Palo Alto Networks documents file-transfer controls in the File Blocking documentation.

QUESTION NO: 2

After an engineer configures an IPSec tunnel with a Cisco ASA, the Palo Alto Networks firewall generates system messages reporting the tunnel is failing to establish.

Which of the following actions will resolve this issue?

A. Ensure that an active static or dynamic route exists for the VPN peer with next hop as the tunnel interface.

B. Configure the Proxy IDs to match the Cisco ASA configuration.

C. Check that IPSec is enabled in the management profile on the external interface.

D. Validate the tunnel interface VLAN against the peer's configuration.

ANSWER: B

Explanation:

Configure the Proxy IDs to match the Cisco ASA configuration. A Cisco ASA commonly uses policy-based VPN definitions, in which the crypto access list defines the protected local and remote networks, also called the encryption domains or traffic selectors. For the IPSec security association to be negotiated successfully, the Palo Alto Networks firewall must propose corresponding Proxy IDs whose local and remote subnets align with the selectors configured on the ASA. A mismatch in either subnet, address direction, or protocol can prevent the peers from agreeing on the Phase 2/IPSec traffic selectors, causing tunnel-establishment failures or an IPSec SA that does not handle the intended traffic.

On PAN-OS, Proxy IDs are configured under the IPSec tunnel and are specifically intended for interoperability with policy-based VPN peers. The firewall can use route-based tunnel interfaces for forwarding, but defining matching Proxy IDs remains necessary when the peer requires explicit Phase 2 selectors. Configure each required local/remote network pair exactly as represented by the ASA encryption-domain configuration, then commit the configuration and renegotiate the tunnel. Palo Alto Networks documents Proxy IDs and their role in matching peer traffic selectors at Proxy ID for IPSec Tunnels. Cisco also describes the ASA requirement for matching protected-network definitions at Cisco ASA-to-ASA IPSec VPN Configuration Example.

QUESTION NO: 3

Which forwarding methods can be used on the Objects tab when configuring the Log Forwarding profile?

A. Panorama, syslog, email

B. Syslog, HTTP, NetFlow

C. Panorama, ADEM, syslog

D. SNMP, HTTP, RADIUS

ANSWER: A

Explanation:

Panorama, syslog, email is correct because these are supported forwarding actions that can be selected in a PAN-OS Log Forwarding profile. A Log Forwarding profile defines filters for specific log types and severities, then associates each matching log entry with one or more configured forwarding destinations. Panorama forwarding sends logs to the centralized Palo Alto Networks management and logging platform, allowing aggregated visibility, reporting, and log storage. Syslog forwarding sends matching log entries to an external syslog server using the configured server profile, which supports integration with SIEM platforms and other operational monitoring systems. Email forwarding uses an Email Server profile to deliver alert messages for logs that meet the configured match criteria. Before these actions can be used in the Log Forwarding profile, the corresponding server profiles or Panorama logging configuration must be available and properly configured. Palo Alto Networks documents log forwarding as a policy-driven mechanism that can forward matching logs to Panorama, syslog servers, and email recipients. See Configure Log Forwarding and Configure Email Alerts.

QUESTION NO: 4

Which two statements apply to configuring required security rules when setting up an IPSec tunnel between a Palo Alto Networks firewall and a third- party gateway? (Choose two.)

A. For incoming and outgoing traffic through the tunnel, creating separate rules for each direction is optional.

B. The IKE negotiation and IPSec/ESP packets are allowed by default via the intrazone default allow policy.

C. For incoming and outgoing traffic through the tunnel, separate rules must be created for each direction.

D. The IKE negotiation and IPSec/ESP packets are denied by default via the interzone default deny policy.

ANSWER: C D

Explanation:

“For incoming and outgoing traffic through the tunnel, separate rules must be created for each direction.” is correct because Palo Alto Networks Security policy is evaluated based on the ingress and egress security zones. A rule permitting protected traffic from an internal zone to the VPN tunnel zone does not automatically permit traffic in the reverse direction, where the source and destination zones are reversed. To support bidirectional communication across a site-to-site VPN, create an appropriate Security policy rule for each required direction, using the tunnel interface’s zone and the relevant local zone.

“The IKE negotiation and IPSec/ESP packets are denied by default via the interzone default deny policy.” is also correct in the context of traffic crossing different zones without an explicit matching allow rule. The default interzone rule denies unmatched traffic, so the firewall policy design must account for the IKE and IPsec communications needed to establish and maintain the VPN, as well as the protected traffic that will use the tunnel. Ensuring that the tunnel interface is assigned to the intended zone and that directional policy rules are present is central to a functional, least-privilege deployment. See Palo Alto Networks’ site-to-site VPN configuration guidance and its Security policy overview.

QUESTION NO: 5

A Managed Security Service Provider (MSSP) is creating a new VSYS for a customer.

To prevent this customer's traffic from overwhelming the firewall's state table, which resource limit should the MSSP configure for the new VSYS?

A. Max security profiles

B. Max bandwidth

C. Max sessions

D. Max Log Forwarding profiles

ANSWER: C

Explanation:

Max sessions is the correct resource limit because the firewall state table tracks active sessions. Every permitted connection consumes a session entry while it is active, including the associated state and inspection information needed for Stateful firewall processing. In a multi-tenant deployment using virtual systems, all tenants ultimately use the same physical firewall session capacity. Configuring a maximum-session quota for the new virtual system reserves a defined ceiling for that customer and prevents an unusually high volume of connections from consuming a disproportionate share of the platform's state-table resources.

This limit is particularly important for an MSSP because tenant traffic patterns can vary significantly over time. A per-virtual-system session limit provides predictable resource isolation: when the tenant reaches its configured quota, additional new sessions for that virtual system are constrained, while established sessions and capacity allocated to other tenants are protected. The appropriate value should be sized according to the customer's expected concurrent-connection demand and the firewall model's overall session capacity. Palo Alto Networks documents virtual-system resource allocation and limits in its Virtual System Capabilities documentation and platform capacity planning in the deployment planning guide.

QUESTION NO: 6

A network administrator is establishing a site-to-site VPN between a Palo Alto Networks firewall and a partner's Check Point Security Gateway. The partner has provided a specific list of local and remote IP address subnets that are permitted through the tunnel. The initial tunnel configuration on the PAN-OS firewall fails during the IKE Phase 2 exchange.

Which configuration step is essential to ensure compatibility with the policy-based Check Point gateway?

A. Define the local and remote subnets provided by the partner in the Proxy ID settings.

B. Create individual Security policies for each pair of local and remote subnets.

C. Assign a specific IP address to the tunnel interface to match the Check Point gateway.

D. Enable Dead Peer Detection (DPD) in the IKE Gateway configuration.

ANSWER: A

Explanation:

Define the local and remote subnets provided by the partner in the Proxy ID settings. A policy-based IPsec peer such as a Check Point gateway commonly uses traffic selectors to identify the precise protected networks for an IPsec security association. On a Palo Alto Networks firewall, Proxy IDs define those Phase 2 traffic selectors, including the local subnet, remote subnet, and protocol. The values must correspond to the encryption-domain or traffic-selector values configured by the peer. If the selectors proposed by PAN-OS do not match what the Check Point gateway expects, the peer can reject the Phase 2 negotiation and no IPsec Security Association is established for the intended traffic. Configure one Proxy ID for each required local/remote subnet pair when the partner requires separate selector pairs, using the appropriate protocol setting—typically `any` unless the partner specifies otherwise. This configuration is particularly important when PAN-OS interoperates with a policy-based VPN implementation, because Proxy IDs provide the explicit subnet definitions that policy-based peers require rather than relying on route-based tunnel behavior alone. After the IPsec tunnel is established, normal security policy and routing configuration controls whether matching traffic is permitted and forwarded into the tunnel. See Palo Alto Networks' Proxy ID for IPsec VPNs documentation and its IPsec tunnel configuration guide.

QUESTION NO: 7

An organization requires GlobalProtect users to authenticate with an external SAML identity provider. The administrator is configuring the firewall for SAML authentication.

Which two configuration tasks are required on the firewall? (Choose two.)

A. Create a SAML Identity Provider Server Profile.

B. Create an Authentication Profile that references the SAML server profile.

C. Configure a RADIUS Server Profile for the identity provider.

D. Configure a captive portal authentication sequence.

E. Create a dynamic address group for every SAML user.

ANSWER: A B

Explanation:

The firewall requires a **SAML Identity Provider Server Profile** that defines the identity provider information used for the SAML exchange. The administrator must also create an **Authentication Profile** that references the SAML identity provider server profile and then apply the authentication profile to the GlobalProtect portal or gateway configuration.

The SAML identity provider must separately be configured to trust the GlobalProtect portal or gateway as a service provider, using the appropriate service provider metadata or values. On the firewall, the identity provider server profile and the authentication profile provide the required authentication integration. Palo Alto Networks documents the configuration process in Configure SAML Authentication.

QUESTION NO: 8

A network administrator is configuring an Aggregate Ethernet (AE) interface on an active/passive high availability (HA) pair. To reduce network downtime during a failover, the administrator wants the passive firewall's AE interface to be fully negotiated with the switch before it becomes active.

Which Link Aggregation Control Protocol (LACP) setting achieves this administrator's goal?

A. LACP Mode active

B. Enable in HA passive state

C. System Priority: 1

D. Transmission Rate: fast

ANSWER: B

Explanation:

Enable in HA passive state is the setting designed for this active/passive HA requirement. Normally, the passive firewall does not actively participate in LACP negotiation for an aggregate Ethernet interface because its data interfaces are in a passive state. Enabling LACP in the HA passive state allows that firewall to exchange LACP protocol data units with the connected switch while it remains the standby peer. Consequently, the switch can maintain the LAG association and the passive firewall has an already-negotiated aggregate when an HA failover occurs.

Because negotiation is completed before the role transition, the newly active firewall can begin forwarding through its AE interface with less delay. This specifically addresses downtime attributable to waiting for LACP to establish and synchronize after failover. Palo Alto Networks documents this behavior as using LACP pre-negotiation on the passive firewall and notes that it helps minimize traffic disruption during HA failover. Configure the setting in the aggregate interface's LACP configuration and ensure the connected switch supports the intended LACP/port-channel behavior.

See the Palo Alto Networks documentation for configuring aggregate interfaces and high availability interface configuration.

QUESTION NO: 9

An administrator is using Panorama to centrally configure common network settings for a group of firewalls.

Which two configuration elements are normally managed through a template or template stack? (Choose two.)

A. Layer 3 interfaces

B. Virtual routers

C. Security policy pre-rules

D. Device group hierarchy

E. Shared address objects

ANSWER: A B

Explanation:

Layer 3 interfaces are managed through templates because interfaces are part of the firewall network configuration. A template can define interface type, addressing, zones, management profiles, and other device-specific network settings.

Virtual routers are also template-managed objects. They contain routing configuration, including static routes and dynamic routing settings, and are pushed through the relevant template or template stack. In contrast, centrally managed Security policy rules and shared policy objects are generally configured through device groups.

See the Palo Alto Networks Manage Templates and Template Stacks documentation.

QUESTION NO: 10

A Security policy rule allows the application ssl and uses application-default as its service setting.

How does the firewall enforce the service portion of this rule after it identifies the application?

A. It allows the application only on its default ports.

B. It allows the application on all TCP and UDP ports.

C. It dynamically opens every port observed for the application.

D. It allows the application only after a custom service object is selected.

ANSWER: A

Explanation:

The firewall allows the application only on its default ports. The application-default service setting restricts a permitted application to the ports that PAN-OS defines as the application's standard ports. This reduces exposure compared to allowing the application on any service or port.

App-ID identification and service enforcement work together in this policy design. If the firewall identifies ssl on a nonstandard port, the rule does not match when its service is set to application-default; the traffic must match another applicable rule. This approach is a common best practice for rules that allow known applications while limiting their use to expected ports.

See the Palo Alto Networks Security Policy Overview.

QUESTION NO: 11

Palo Alto Networks NGFWs use SSL/TLS profiles to secure which two types of connections? (Choose two.)

A. NAT tables

B. User Authentication

C. GlobalProtect Gateways

D. GlobalProtect Portal

ANSWER: C D

Explanation:

GlobalProtect Gateways and the GlobalProtect Portal use SSL/TLS service profiles because each operates as a TLS service endpoint on the firewall. The portal is the initial client-facing GlobalProtect endpoint: it delivers client configuration, supports client download and updates where applicable, and handles the initial connection and authentication workflow. Assigning an SSL/TLS service profile to the portal determines the server certificate presented to connecting clients and the TLS protocol and cipher settings used to protect that session.

GlobalProtect Gateways also require an SSL/TLS service profile to secure their client-facing TLS connections. The gateway authenticates GlobalProtect users and establishes the secure remote-access connection through which the client receives access to protected network resources. The selected profile provides the gateway identity certificate and defines the TLS settings used during client negotiation. Administrators should use a certificate whose subject or subject alternative name matches the public hostname clients use for the portal or gateway, and should select protocols and ciphers consistent with the organization's security requirements.

Palo Alto Networks documents SSL/TLS service profiles as the mechanism for specifying a certificate and TLS settings for services including GlobalProtect. See Configure an SSL/TLS Service Profile and GlobalProtect Portals.

QUESTION NO: 12

A network architect is planning the deployment of a new IPSec VPN tunnel to connect a local data center to a cloud environment. The plan must include all necessary Security policy configurations for both tunnel negotiation and data transit.

Which two Security policy requirements must be included in the implementation plan? (Choose two answers)

A. The default interzone-default security policy is sufficient to allow the tunnel negotiation traffic between the firewall and the remote peer.

B. A pair of policies is required to control the flow of data traffic into and out of the security zone assigned to the tunnel interface.

C. A policy must explicitly permit only the IKE application between the external-facing zone and local zone.

D. A policy must explicitly permit the IPSec container application between the external-facing zone and local zone.

ANSWER: B D

Explanation:

An IPSec VPN deployment requires Security policy planning for two distinct traffic stages: establishing the VPN and carrying the protected traffic after decryption. "A pair of policies is required to control the flow of data traffic into and out of the security zone assigned to the tunnel interface" is required because a tunnel interface is assigned to a security zone. Once packets are decapsulated, the firewall evaluates them as normal routed traffic between the tunnel zone and the internal, cloud, or other destination zone. Policy must therefore permit the intended protected applications and networks in each required direction, following the organization's least-privilege design.

"A policy must explicitly permit the IPSec container application between the external-facing zone and local zone" addresses VPN control-plane traffic received at the firewall's external VPN endpoint. The policy must allow the applicable IPSec application family so that the firewall can establish and maintain the site-to-site VPN with its peer. This is separate from policies governing the decrypted payload; permitting tunnel establishment alone does not authorize user or application traffic through the tunnel. Palo Alto Networks documents IPSec VPN configuration and traffic protection concepts at IPSec VPN documentation and explains zone-based rule evaluation in the Security Policy Overview.

QUESTION NO: 13

Which two actions in the IKE Gateways will allow implementation of post-quantum cryptography when building VPNs between multiple Palo Alto Networks NGFWs? (Choose two.)

A. Select IKE v2, enable the Advanced Options PQ PPK, then set a 64+ character string for the post-quantum pre shared key.

B. Ensure Authentication is set to œcertificate, then import a post-quantum derived certificate.

C. Select IKE v2 Preferred, enable the Advanced Options PQ KEM, then add one or more œRounds.

D. Select IKE v2, enable the Advanced Options PQ KEM, then create an IKE Crypto Profile with Advanced Options adding one or more œRounds.

ANSWER: A D

Explanation:

Selecting IKE v2, enabling the Advanced Options PQ PPK setting, and supplying a 64+ character post-quantum pre-shared key enables the post-quantum pre-shared key method for an IKEv2 VPN. The same PPK must be configured on both VPN peers so it can provide additional quantum-resistant key material during IKE negotiation. This is a supported post-quantum VPN approach when the peers are Palo Alto Networks NGFWs.

Selecting IKE v2, enabling the Advanced Options PQ KEM setting, and configuring one or more PQ KEM rounds in the associated IKE Crypto Profile enables the key encapsulation mechanism method. PQ KEM is negotiated through IKEv2 and uses the configured rounds to establish post-quantum keying material. Both peers must use compatible PQ KEM settings and the applicable IKE Crypto Profile for the tunnel to negotiate successfully. Palo Alto Networks documents PQ PPK and PQ KEM as the supported post-quantum mechanisms for site-to-site VPN configuration, with PQ KEM rounds configured through the IKE crypto profile's advanced options. See the Post-Quantum VPN documentation and the IKE Gateway documentation.

QUESTION NO: 14

Which two actions in the IKE Gateways will allow implementation of post-quantum cryptography when building VPNs between multiple Palo Alto Networks NGFWs? (Choose two.)

A. Select IKE v2, enable the Advanced Options PQ PPK, then set a 64+ character string for the post-quantum pre shared key.

B. Ensure Authentication is set to “certificate,” then import a post-quantum derived certificate.

C. Select IKE v2 Preferred, enable the Advanced Options PQ KEM, then add one or more “Rounds.”

D. Select IKE v2, enable the Advanced Options PQ KEM, then create an IKE Crypto Profile with Advanced Options adding oneor more “Rounds.”

ANSWER: A C

Explanation:

Selecting IKE v2 and enabling the Advanced Options PQ PPK setting with a 64-or-more-character post-quantum pre-shared key is a supported way to add post-quantum protection to IKEv2 authentication. The post-quantum pre-shared key (PPK) is mixed into IKEv2 key derivation, providing additional protection for the VPN’s security association against future cryptographically relevant quantum computers. Both peers must use the same PPK configuration for negotiation to succeed.

Selecting IKE v2 Preferred, enabling PQ KEM under Advanced Options, and adding one or more rounds is also correct. PQ KEM enables post-quantum key establishment during IKEv2 negotiation. The selected rounds define the KEM algorithms and their order of use; the corresponding IKE Crypto profile must contain compatible PQ KEM rounds so the peers can negotiate a mutually supported post-quantum exchange. Palo Alto Networks documents PQ PPK and PQ KEM as distinct IKEv2 quantum-safe capabilities, allowing an administrator to deploy either approach as appropriate for interoperability and security requirements.

See the Palo Alto Networks IPsec VPN documentation and the IKE Gateway configuration documentation for configuration details.

QUESTION NO: 15

What must be configured before a firewall administrator can define policy rules based on users and groups?

A. User Mapping profile

B. Authentication profile

C. Group mapping settings

D. LDAP Server profile

ANSWER: C

Explanation:

Group mapping settings must be configured so the firewall can retrieve group and user-membership information from the organization's directory service, such as Microsoft Active Directory or LDAP. This makes directory groups available as selectable entries when an administrator creates a Security policy rule with User-ID-based source or destination users. The firewall uses the configured directory connection and group-mapping query parameters to learn which users belong to which groups, then maintains that information locally for policy evaluation.

Group mapping is therefore the necessary configuration for policies that use groups, because a group name in a rule must be resolved to the directory users who are members of that group. Administrators can narrow the mapping scope to relevant domains or groups and can configure update intervals so membership changes are reflected in policy enforcement. Once group information is mapped, User-ID can associate observed IP addresses with authenticated usernames, allowing the firewall to apply the rule to the appropriate users as they generate traffic.

Palo Alto Networks documents the group-mapping configuration workflow and its role in using directory groups in policy at Group Mapping. See also the User-ID administration overview for how identity data is used by firewall policy.

QUESTION NO: 16

An engineer is configuring a site-to-site IPSec VPN to a partner network. The IKE Gateway and IPSec tunnel configurations are complete, and the tunnel interface has been assigned to a security zone. However, the tunnel fails to establish, and no application traffic passes through it once it is up.

Which two Security policy configurations are required to allow tunnel establishment and data traffic flow in this scenario? (Choose two.)

A. A security rule is needed to allow IKE and IPSec traffic between the zone where the physical interface resides and the zone of the partner gateway.

B. A single bidirectional security rule must be configured to manage traffic flowing through the tunnel interface.

C. Security rules must be configured to permit application traffic from the local zone to the tunnel zone, and from the tunnel zone to the local zone.

D. An Application Override policy is needed to allow both the IKE negotiation and the encapsulated data traffic.

ANSWER: A C

Explanation:

A security rule is needed to allow IKE and IPSec traffic between the zone where the physical interface resides and the zone of the partner gateway because the firewall must permit the VPN negotiation and IPsec encapsulation packets arriving at, and leaving from, the external interface used by the IKE gateway. The policy should match the relevant external zones, peer

addresses, and IKE/IPsec applications or services as appropriate for the deployment. This allows the firewall to establish the security association with the partner gateway.

Security rules must be configured to permit application traffic from the local zone to the tunnel zone, and from the tunnel zone to the local zone because a tunnel interface is assigned to a security zone and traffic carried inside the VPN is evaluated against Security policy after it is decapsulated. Separate directional rules support the normal stateful traffic flow and also allow partner-initiated sessions when required. The policies should use the actual protected local and remote subnets, the tunnel zone, and only the applications/services authorized by the organization. Palo Alto Networks documents the need to create Security policy rules for traffic sent through an IPsec tunnel and to associate the tunnel interface with a zone. See Security Policy for VPNs and Create an IPsec Tunnel.