

DUMPS ARENA

Palo Alto Networks Next-Generation Firewall Engineer

Palo Alto Networks NGFW-Engineer

Version Demo

Total Demo Questions: 10

Total Premium Questions: 50

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

Which configuration step is required when implementing a new self-signed root certificate authority (CA) certificate for SSL decryption on a Palo Alto Networks firewall?

- A. Import the new subordinate CA certificate into the trust stores of all client devices.
- B. Set the subordinate CA certificate as the default routing certificate for all network traffic.
- C. Configure the subordinate CA to issue certificates with indefinite validity periods.
- D. Disable all existing SSL decryption rules until the new certificate is fully propagated.

ANSWER: A**Explanation:**

When you use a self-signed root CA for SSL Forward Proxy decryption, the firewall will generate (or use) a CA certificate to dynamically sign certificates for the sites users visit. If client devices don't trust that CA, they'll see browser certificate warnings and many apps will simply fail the TLS handshake.

So the key "must-do" step is getting the firewall's decryption CA certificate into the trusted root store on every client (or your enterprise trust store via AD/GPO/MDM). Once that CA is trusted, clients will accept the certificates the firewall re-signs during decryption, and traffic can be decrypted and re-encrypted cleanly.

None of the other options are required for SSL decryption: routing certificates aren't a thing here, indefinite validity is a bad practice and not needed, and you don't have to disable decryption rules while you roll out trust (though you might stage it operationally).

References: <https://docs.paloaltonetworks.com/pan-os/10-2/pan-os-admin/decryption/configure-ssl-forward-proxy-decryption> and <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIEZCA0>

QUESTION NO: 2

Which set of options is available for detailed logs when building a custom report on a Palo Alto Networks NGFW?

- A. Traffic, User-ID, URL
- B. Traffic, threat, data filtering, User-ID
- C. GlobalProtect, traffic, application statistics
- D. Threat, GlobalProtect, application statistics, WildFire submissions

ANSWER: B**Explanation:**

When you build a custom report on a Palo Alto Networks firewall, the "detailed logs" choices map to the main log databases the NGFW writes to. The core ones you can report on are Traffic logs (sessions and flows), Threat logs (IPS/AV/anti-

spyware and other threat events), Data Filtering logs (data patterns and file/property matches), and User-ID logs (user mapping and related events). That set is the standard group you'll see when you're choosing what kind of detailed log source your custom report should use.

The other options listed (like GlobalProtect, application statistics, or WildFire submissions) aren't typically presented as "detailed logs" in the same way for custom reports. Some of those exist as separate reporting areas or different log types/features, but they don't match the common "detailed logs" selection set used for custom report building.

Reference: <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring/view-and-manage-logs/log-types-and-severity-levels>

QUESTION NO: 3

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:**

On Palo Alto Networks firewalls, Security policy is evaluated by traffic direction (source zone to destination zone). So if you want traffic to work both ways across the VPN (for example, HQ to branch and branch back to HQ), you normally need a rule that permits each direction. One rule won't automatically "cover" the return direction if the zones flip.

Also, don't forget the tunnel has to come up before any user traffic can pass. By default, traffic that crosses zones hits the interzone default deny rule unless you explicitly allow it. That means IKE (UDP/500 and sometimes UDP/4500 for NAT-T) and IPSec ESP (IP protocol 50) can be blocked unless you create rules permitting them to the firewall's outside interface (often to the IKE gateway's local IP).

Palo Alto's docs call out both the need to allow IKE/IPsec and the way policy is zone-based, which is why these two statements are the right picks. References: <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/vpns/set-up-site-to-site-vpn> and <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/policy/security-policy/security-policy-overview>

QUESTION NO: 4

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:**

On Palo Alto Networks firewalls, an SSL/TLS Service Profile is basically the “certificate + TLS settings” the firewall presents when it’s acting like a server for a feature. So whenever a client connects to the firewall over HTTPS/TLS, you typically pick an SSL/TLS profile to control what cert is used and what TLS versions/ciphers are allowed.

That’s exactly why GlobalProtect uses these profiles in two key places: the GlobalProtect **Portal** (where users initially connect to download the app/config and authenticate) and the GlobalProtect **Gateway** (where the VPN tunnel and ongoing secure communications are established). In both cases, the firewall must securely terminate TLS for the client connection, and the SSL/TLS profile is what defines how that handshake is done.

NAT tables aren’t “secured” with SSL/TLS profiles—they’re just part of packet processing. And while user authentication can happen over TLS, the SSL/TLS profile itself is applied to the service endpoints (like the portal/gateway), not to “user authentication” as a standalone connection type.

References: <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/certificate-management/configure-an-ssl-tls-service-profile> and <https://docs.paloaltonetworks.com/globalprotect/11-0/globalprotect-admin/globalprotect-portals>

QUESTION NO: 5

Which two statements describe an external zone in the context of virtual systems (VSYS) on a Palo Alto Networks firewall? (Choose two.)

- A. It is associated with an interface within a VSYS of a firewall.
- B. It is a security object associated with a specific virtual router of a VSYS.
- C. It is not associated with an interface; it is associated with a VSYS itself.
- D. It is a security object associated with a specific VSYS.

ANSWER: A D**Explanation:**

In PAN-OS, a zone is something you assign to an interface (or subinterface) so the firewall can classify traffic as it comes in and out. So if you’re calling something an “external zone,” it’s really just a zone that’s attached to the interface facing the outside network (like an ISP or untrusted network). That’s why the statement about being associated with an interface inside a VSYS is spot on.

Zones are also scoped to a VSYS. Each VSYS is its own virtual firewall with its own policies and objects, and zones live inside that VSYS. So saying the external zone is a security object associated with a specific VSYS is also correct—VSYS1’s zones don’t automatically exist in VSYS2.

The virtual router is separate from zones: routing decides where packets go, while zones are used for policy and security segmentation. And a zone isn’t just “associated with the VSYS itself” without interfaces—interfaces are how zones actually get used in traffic classification.

References: <https://docs.paloaltonetworks.com/pan-os/10-2/pan-os-admin/networking/configure-interfaces/zone>
<https://docs.paloaltonetworks.com/pan-os/10-2/pan-os-admin/virtual-systems/virtual-systems-overview>

QUESTION NO: 6

Which two zone types are valid when configuring a new security zone? (Choose two.)

- A. Tunnel
- B. Intrazone
- C. Internal
- D. Virtual Wire

ANSWER: A D**Explanation:**

On a Palo Alto Networks firewall, a security zone has a specific *zone type* that matches how the interface works. A **Tunnel** zone is used when the interface is a tunnel interface (like an IPsec tunnel or GRE). Traffic that enters/leaves through that tunnel interface is mapped to the Tunnel zone, which is how policy and logging are applied.

A **Virtual Wire** zone is for vwire deployments, where the firewall is essentially “in-line” and passing traffic transparently at Layer 2. You still assign zones so you can write security policies, but the firewall isn’t routing between subnets like it would with a Layer 3 interface.

The other choices don’t fit as zone types you can pick when creating a zone. “Intrazone” is something you’ll see in policy concepts (traffic within the same zone), not a zone type. “Internal” also isn’t a selectable zone type in PAN-OS.

References: <https://docs.paloaltonetworks.com/pan-os/11-1/pan-os-admin/networking/configure-zones> and <https://docs.paloaltonetworks.com/pan-os/11-1/pan-os-admin/networking/virtual-wire-interfaces>

QUESTION NO: 7

Which interface types should be used to configure link monitoring for a high availability (HA) deployment on a Palo Alto Networks NGFW?

- A. HA, Virtual Wire, and Layer 2
- B. Tap, Virtual Wire, and Layer 3
- C. Virtual Wire, Layer 2, and Layer 3
- D. HA, Layer 2. and Layer 3

ANSWER: C**Explanation:**

For HA link monitoring, you’re basically telling the firewall, “If this data interface goes down, treat it as a real failure and react (like triggering a failover).” Palo Alto only lets you do that on normal traffic-carrying interfaces—so the supported types are Virtual Wire, Layer 2, and Layer 3.

Options that include “HA” as an interface type are just wrong, because HA is a feature/set of ports, not a data interface type you pick for link monitoring. And TAP interfaces don’t fit either—TAP is meant for passive monitoring and doesn’t participate in forwarding the way L2/L3/vwire do, so it’s not used for HA link monitoring decisions.

If you want to double-check, Palo Alto's HA docs describe link monitoring as being configured on data interfaces (vwire, L2, or L3) so the device can detect a path failure and take action. See <https://docs.paloaltonetworks.com/pan-os> and the High Availability section for link monitoring details.

QUESTION NO: 8

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 one or more `rounds`.

ANSWER: C D**Explanation:**

On Palo Alto Networks NGFWs, the post-quantum piece for IPsec VPN is handled in the IKE Gateway by turning on PQ KEM (post-quantum key encapsulation). That's the setting that adds a quantum-resistant key exchange step to IKEv2, so the VPN can negotiate keys in a way that's designed to hold up even if someone has a quantum computer later.

Option C is right because it explicitly uses IKEv2 (IKE v2 Preferred is fine) and enables PQ KEM, then lets you choose one or more "Rounds." Those rounds are basically the PQ KEM parameters used during the exchange.

Option D is also right because it still enables PQ KEM on IKEv2, and it calls out creating/using an IKE Crypto Profile where you can set the PQ KEM advanced options (including rounds). In real deployments, you commonly pair the gateway setting with a crypto profile to make sure both sides match.

The other options don't line up with Palo Alto's PQ VPN approach: PQ PPK and "post-quantum certificates" aren't what you enable here to get PQ key exchange working between NGFWs.

References: <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/vpns/site-to-site-vpns/post-quantum-vpn> and <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/vpns/ike-gateways>

QUESTION NO: 9

An organization runs multiple Kubernetes clusters both on-premises and in public clouds (AWS, Azure, GCP). They want to deploy the Palo Alto Networks CN-Series NGFW to secure east-west traffic within each cluster, maintain consistent Security policies across all environments, and dynamically scale as containerized workloads spin up or down. They also plan to use a centralized Panorama instance for policy management and visibility. Which approach meets these requirements?

- A.** Install standalone CN-Series instances in each cluster with local configuration only. Export daily policy configuration snapshots to Panorama for recordkeeping, but do not unify policy enforcement.
- B.** Configure the CN-Series only in public cloud clusters, and rely on Kubernetes Network Policies for on-premises cluster security. Synchronize partial policy information into Panorama manually as needed.

C. Use Kubernetes-native deployment tools (e.g., Helm) to deploy CN-Series in each cluster, ensuring local insertion into the service mesh or CN

D. Deploy a single CN-Series firewall in the on-premises data center to process traffic for all clusters, connecting remote clusters via VPN or peering. Manage this single instance through Panorama.

E. Manage all CN-Series firewalls centrally from Panorama, applying uniform Security policies across on-premises and cloud clusters.

ANSWER: C

Explanation:

The best fit is deploying CN-Series into every Kubernetes cluster using Kubernetes-native tooling (like Helm) and inserting it into the cluster's traffic path (for example via CNI/service insertion). That's what lets the firewall actually see and control east-west pod-to-pod traffic inside each cluster, instead of only protecting north-south traffic at the edge.

Because it's deployed as a Kubernetes app, CN-Series can scale along with the cluster. When nodes and workloads scale up or down, you can scale the CN-Series components with Kubernetes mechanisms, keeping inspection capacity aligned with demand.

Finally, tying all those CN-Series firewalls back to a centralized Panorama gives you the "single place" to push consistent Security policies and get visibility across on-prem and all clouds. Options that rely on local-only config, Kubernetes NetworkPolicy alone, or a single remote firewall won't meet the east-west and multi-cluster scaling goals.

References: <https://docs.paloaltonetworks.com/cloud-ngfw/cn-series> and <https://docs.paloaltonetworks.com/panorama>

QUESTION NO: 10

For which two purposes is an IP address configured on a tunnel interface? (Choose two.)

- A.** Use of dynamic routing protocols
- B.** Tunnel monitoring
- C.** Use of peer IP
- D.** Redistribution of User-ID

ANSWER: A B

Explanation:

An IP on a tunnel interface is mainly there so the firewall can treat the tunnel like a real Layer 3 interface. The most common reason is dynamic routing: if you want OSPF, BGP, or another routing protocol to run across the VPN, the tunnel interface needs an IP address so it can form adjacencies/neighbors and exchange routes over that logical link.

The other big reason is tunnel monitoring/health checks. When you enable monitoring, the firewall needs a source IP to send probe traffic (like ICMP) through the tunnel toward a monitor destination. That source is typically the tunnel interface IP, and it lets the firewall decide whether the tunnel is actually usable and fail over if it isn't.

The peer IP itself is configured in the IKE/IPsec gateway settings, not as the tunnel interface IP, and User-ID redistribution isn't something that depends on assigning an IP to the tunnel interface.

References: <https://docs.paloaltonetworks.com/pan-os/10-2/pan-os-admin/vpns/site-to-site-vpn-concepts/tunnel-interfaces>
and <https://docs.paloaltonetworks.com/pan-os/10-2/pan-os-admin/vpns/site-to-site-vpn-concepts/tunnel-monitoring>