



# **Paloalto-Networks**

## **Exam Questions NGFW-Engineer**

Palo Alto Networks Next-Generation Firewall Engineer

#### NEW QUESTION 1

An enterprise uses GlobalProtect with both user- and machine-based certificate authentication and requires pre-logon, OCSP checks, and minimal user disruption. They manage multiple firewalls via Panorama and deploy domain-issued machine certificates via Group Policy. Which approach ensures continuous, secure connectivity and consistent policy enforcement?

- A. Use a wildcard certificate from a public CA, disable all revocation checks to reduce latency, and manage certificate renewals manually on each firewall.
- B. Distribute root and intermediate CAs via Panorama template, use distinct certificate profiles for user versus machine certs, reference an internal OCSP responder, and automate certificate deployment with Group Policy.
- C. Configure a single certificate profile for both user and machine certificate
- D. Rely solely on CRLs for revocation to minimize complexity.
- E. Deploy self-signed certificates on each firewall, allow IP-based authentication to override certificate checks, and use default GlobalProtect settings for user / machine identification.

**Answer: B**

#### Explanation:

To ensure continuous, secure connectivity and consistent policy enforcement with GlobalProtect in an enterprise environment that uses user- and machine-based certificate authentication, the approach should:

Distribute root and intermediate CAs via Panorama templates: This ensures that all firewalls managed by Panorama share the same trusted certificate authorities for consistency and security.

Use distinct certificate profiles for user vs. machine certificates: This enables separate handling of user and machine authentication, ensuring that both types of certificates are managed and validated appropriately.

Reference an internal OCSP responder: By integrating OCSP checks, the firewall can validate certificate revocation in real-time, meeting the security requirement while minimizing the overhead and latency associated with traditional CRLs (Certificate Revocation Lists).

Automate certificate deployment with Group Policy: This ensures that machine certificates are deployed in a consistent and scalable manner across the enterprise, reducing manual intervention and minimizing user disruption.

This approach supports the requirements for pre-logon, OCSP checks, and minimal user disruption, while maintaining a secure, automated, and consistent authentication process across all firewalls managed via Panorama.

#### NEW QUESTION 2

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 onl
- B. Export daily policy configuration snapshots to Panorama for recordkeeping, but do not unify policy enforcement.
- C. Configure the CN-Series only in public cloud clusters, and rely on Kubernetes Network Policies for on-premises cluster securit
- D. Synchronize partial policy information into Panorama manually as needed.
- E. Use Kubernetes-native deployment tools (e.g., Helm) to deploy CN-Series in eachcluster, ensuring local insertion into the service mesh or CN
- F. Manage all CN-Series firewalls centrally from Panorama, applying uniform Security policies across on-premises and cloud clusters.
- G. Deploy a single CN-Series firewall in the on-premises data center to process traffic for all clusters, connecting remote clusters via VPN or peerin
- H. Manage this single instance through Panorama.

**Answer: C**

#### Explanation:

This approach meets all the requirements for securing east-west traffic within each Kubernetes cluster, maintaining consistent security policies across on-premises and cloud environments, and allowing for dynamic scaling of the CN-Series NGFWs as containerized workloads spin up or down. By using Kubernetes-native deployment tools (such as Helm), the CN-Series NGFWs can be deployed and scaled dynamically within each cluster. Local insertion into the service mesh or CNI ensures that the NGFW can inspect traffic at the appropriate points within the cluster.

Centralized management via Panorama ensures that security policies are uniform across both on-premises and cloud environments, providing visibility and control across all clusters.

#### NEW QUESTION 3

What is the purpose of assigning an Admin Role Profile to a user in a Palo Alto Networks NGFW?

- A. Allow access to all resources without restrictions.
- B. Enable multi-factor authentication (MFA) for administrator access.
- C. Define granular permissions for management tasks.
- D. Restrict access to sensitive report data.

**Answer: C**

#### Explanation:

Assigning an Admin Role Profile to a user in a Palo Alto Networks NGFW is used to define granular permissions for management tasks. This allows administrators to control what actions a user can perform on the firewall, such as configuration changes, monitoring, and logging. By assigning different admin roles, you can ensure that users have access only to the areas and tasks they need, enforcing the principle of least privilege.

#### NEW QUESTION 4

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 implementing a new self-signed root certificate authority (CA) for SSL decryption on a Palo Alto Networks firewall, the subordinate CA certificate (which is generated by the firewall) must be imported into the trust stores of all client devices. This ensures that client devices trust the firewall as a valid certificate authority, enabling the firewall to decrypt and re-encrypt SSL traffic.

Importing the subordinate CA certificate into the client devices' trust stores is necessary for those devices to trust the new self-signed root CA and properly handle SSL decryption traffic.

**NEW QUESTION 5**

Which statement applies to Log Collector Groups?

- A. Log redundancy is available only if each Log Collector has the same amount of total disk storage.
- B. Enabling redundancy increases the log processing traffic in a Collector Group by 50%.
- C. In any single Collector Group, all the Log Collectors must run on the same Panorama model.
- D. The maximum number of Log Collectors in a Log Collector Group is 18 plus two hot spares.

**Answer:** D

**Explanation:**

The maximum number of Log Collectors that can be added to a Log Collector Group is 18 plus 2 hot spares, ensuring redundancy and availability in case of failure. This allows for a total of up to 20 Log Collectors in a group, providing sufficient scalability and reliability for log collection.

**NEW QUESTION 6**

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:**

When configuring link monitoring for high availability (HA) on a Palo Alto Networks NGFW, the following interface types are supported:

Virtual Wire: Used when you have a transparent mode firewall deployment, where the firewall operates at Layer 2 to monitor traffic between two network segments.

Layer 2: Also used in transparent mode, where the firewall operates as a Layer 2 device and can be configured for link monitoring.

Layer 3: Used in routed mode, where the firewall is involved in routing traffic and can also be configured to monitor links.

**NEW QUESTION 7**

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:**

The Proxy IDs (or Traffic Selectors) define the local and remote subnets that are allowed to communicate over the IPSec tunnel. If the Proxy IDs on the Palo Alto Networks firewall do not match the configuration on the Cisco ASA, the tunnel will fail to establish because the firewalls won't agree on which traffic to encrypt.

Ensuring that the Proxy IDs match between the Palo Alto Networks firewall and the Cisco ASA will resolve the issue.

**NEW QUESTION 8**

When integrating Kubernetes with Palo Alto Networks NGFWs, what is used to secure traffic between microservices?

- A. Service graph
- B. Ansible automation modules
- C. Panorama role-based access control
- D. CN-Series firewalls

**Answer:** D

**Explanation:**

When integrating Kubernetes with Palo Alto Networks NGFWs, the CN-Series firewalls are specifically designed to secure traffic between microservices in containerized environments. These firewalls provide advanced security features like Application Identification (App-ID), URL filtering, and Threat Prevention to secure communication between containers and microservices within a Kubernetes environment.

**NEW QUESTION 9**

Which type of firewall resource can be assigned when configuring a new firewall virtual system (VSYS)?

- A. ICPU
- B. Sessions limit

- C. Memory
- D. Security profile limit

**Answer:** B

**Explanation:**

When configuring a new firewall virtual system (VSYS) on a Palo Alto Networks firewall, one of the resources that can be assigned is the sessions limit. This setting allows the administrator to control the number of active sessions that can be handled by the VSYS, ensuring that each virtual system has an appropriate allocation of resources based on its needs.

**NEW QUESTION 10**

In regard to the Advanced Routing Engine (ARE), what must be enabled first when configuring a logical router on a PAN-OS firewall?

- A. License
- B. Plugin
- C. Content update
- D. General setting

**Answer:** A

**Explanation:**

To enable the Advanced Routing Engine (ARE) on a Palo Alto Networks firewall, the license for the ARE must be applied first. Without the proper license, the firewall cannot activate and use the advanced routing features provided by ARE, such as support for more complex routing protocols (e.g., BGP, OSPF, etc.). Once the license is applied and validated, the routing engine can be configured, allowing the creation of logical routers and routing policies.

**NEW QUESTION 10**

When deploying Palo Alto Networks NGFWs in a cloud service provider (CSP) environment, which method ensures high availability (HA) across multiple availability zones?

- A. Deploying Ansible scripts for zone-specific scaling
- B. Implementing Terraform templates for redundancy within one availability zone
- C. Using load balancer and health probes
- D. Configuring active/active HA

**Answer:** C

**Explanation:**

To ensure high availability (HA) across multiple availability zones (AZs) in a cloud service provider (CSP) environment, using a load balancer with health probes is a recommended method. This setup ensures that traffic can be directed to the healthy NGFW instances across multiple availability zones. If one NGFW instance or availability zone goes down, the load balancer can redirect traffic to the available instance(s) in other zones, providing redundancy and maintaining service availability.

**NEW QUESTION 14**

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:** AB

**Explanation:**

Use of dynamic routing protocols: An IP address is needed on the tunnel interface to participate in dynamic routing protocols (like OSPF, BGP, etc.) over the tunnel. This allows the firewall to advertise routes and receive updates over the tunnel.

Tunnel monitoring: The IP address on the tunnel interface can also be used for monitoring the tunnel's status. Tunnel monitoring (such as IPSec tunnel monitoring) requires an IP address on the tunnel interface to check the health and availability of the tunnel.

**NEW QUESTION 18**

An NGFW engineer is configuring multiple Layer 2 interfaces on a Palo Alto Networks firewall, and all interfaces must be assigned to the same VLAN. During initial testing, it is reported that clients located behind the various interfaces cannot communicate with each other.

Which action taken by the engineer will resolve this issue?

- A. Configure each interface to belong to the same Layer 2 zone and enable IP routing between them.
- B. Assign each interface to the appropriate Layer 2 zone and configure a policy that allows traffic within the VLAN.
- C. Assign each interface to the appropriate Layer 2 zone and configure Security policies for interfaces not assigned to the same zone.
- D. Enable IP routing between the interfaces and configure a Security policy to allow traffic between interfaces within the VLAN.

**Answer:** B

**Explanation:**

In a Layer 2 configuration, interfaces are typically grouped into the same Layer 2 zone. When the interfaces are assigned to the same VLAN, the firewall will treat them as part of the same broadcast domain.

In a Layer 2 setup, interfaces must be in the same Layer 2 zone to allow the traffic within the same VLAN to pass. Additionally, a security policy must be configured to allow traffic within this VLAN or zone. This will resolve the issue by ensuring that traffic is permitted between clients behind different interfaces assigned to the same VLAN.

#### NEW QUESTION 20

An engineer is implementing a new rollout of SAML for administrator authentication across a company's Palo Alto Networks NGFWs. User authentication on company firewalls is currently performed with RADIUS, which will remain available for six months, until it is decommissioned. The company wants both authentication types to be running in parallel during the transition to SAML.

Which two actions meet the criteria? (Choose two.)

- A. Create a testing and rollback plan for the transition from Radius to SAML, as the two authentication profiles cannot be run in tandem.
- B. Create an authentication sequence that includes both the RADIUS Server Profile and SAML Identity Provider Server Profile to run the two services in tandem.
- C. Create and apply an authentication profile with the SAML Identity Provider Server Profile.
- D. Create and add the SAML Identity Provider Server Profile to the authentication profile for the RADIUS Server Profile.

**Answer:** BD

#### Explanation:

To enable both RADIUS and SAML authentication to run in parallel during the transition period, you need to configure an authentication sequence and an authentication profile that includes both authentication methods.

By creating an authentication sequence that includes both RADIUS and SAML server profiles, the firewall will attempt authentication with RADIUS first and, if that fails, will fall back to SAML. This enables both authentication types to function simultaneously during the transition period.

You can also configure an authentication profile that includes both the RADIUS Server Profile and the SAML Identity Provider server profile. This setup allows the firewall to use both RADIUS and SAML for authentication requests, and it will check both authentication methods in parallel.

#### NEW QUESTION 22

A multinational organization wants to use the Cloud Identity Engine (CIE) to aggregate identity data from multiple sources (on premises AD, Azure AD, Okta) while enforcing strict data isolation for different regional business units. Each region's firewalls, managed via Panorama, must only receive the user and group information relevant to that region. The organization aims to minimize administrative overhead while meeting data sovereignty requirements.

Which approach achieves this segmentation of identity data?

- A. Create one CIE tenant, aggregate all identity data into a single view, and redistribute the full dataset to all firewall
- B. Rely on per-firewall Security policies to restrict access to out-of-scope user and group information.
- C. Establish separate CIE tenants for each business unit, integrating each tenant with the relevant identity source
- D. Redistribute user and group data from each tenant only to the region's firewalls, maintaining a strict one-to-one mapping of tenant to business unit.
- E. Disable redistribution of identity data entirely
- F. Instead, configure each regional firewall to pull user and group details directly from its local identity providers (IdPs).
- G. Deploy a single CIE tenant that collects all identity data, then configure segments within the tenant to filter and redistribute only the relevant user/group sets to each regional firewall group.

**Answer:** B

#### Explanation:

To meet the requirement of data isolation for different regional business units while minimizing administrative overhead, the best approach is to establish separate Cloud Identity Engine (CIE) tenants for each business unit. Each tenant would be integrated with the relevant identity sources (such as on-premises AD, Azure AD, and Okta) for that specific region. This ensures that the identity data for each region is kept isolated and only relevant user and group data is distributed to the respective regional firewalls.

By maintaining a strict one-to-one mapping between CIE tenants and business units, the organization ensures that each region's firewall only receives the user and group data relevant to that region, thus meeting data sovereignty requirements and minimizing administrative complexity.

#### NEW QUESTION 26

During an upgrade to the routing infrastructure in a customer environment, the network administrator wants to implement Advanced Routing Engine (ARE) on a Palo Alto Networks firewall.

Which firewall models support this configuration?

- A. PA-5280, PA-7080, PA-3250, VM-Series
- B. PA-455, VM-Series, PA-1410, PA-5450
- C. PA-3260, PA-5410, PA-850, PA-460
- D. PA-7050, PA-1420, VM-Series, CN-Series

**Answer:** C

#### Explanation:

The Advanced Routing Engine (ARE) is supported on Palo Alto Networks firewalls that utilize the PAN-OS 11.0+ software and have the required hardware architecture. The supported models include PA-3200 Series, PA-5400 Series, PA-800 Series, and PA-400 Series. These models provide enhanced routing capabilities, including BGP, OSPF, and more complex routing policies.

PA-3260 and PA-5410 are part of the PA-3200 and PA-5400 Series, which are known to support ARE.

PA-850 and PA-460 are within the PA-800 and PA-400 Series, which also support ARE

#### NEW QUESTION 28

In a hybrid cloud deployment, what is the primary function of Ansible in managing Palo Alto Networks NGFWs?

- A. It provides a web interface for managing NGFW hardware clusters.
- B. It enables centralized log collection and correlation for NGFWs.
- C. It facilitates dynamic updates to NGFW threat databases.
- D. It automates NGFW policy updates and configurations through playbooks.

**Answer:** D

#### Explanation:

In a hybrid cloud deployment, Ansible is primarily used for automating configurations and policy updates on Palo Alto Networks Next-Generation Firewalls (NGFWs). Through the use of playbooks, Ansible can automate the process of deploying security policies, updating configurations, and managing the firewall's state, which enhances efficiency and consistency across multiple NGFWs in a large or hybrid cloud environment.

### NEW QUESTION 30

Which configuration in the LACP tab will enable pre-negotiation for an Aggregate Ethernet (AE) interface on a Palo Alto Networks high availability (HA) active/passive pair?

- A. Set Transmission Rate to ??fast.??
- B. Set passive link state to ??Auto.??
- C. Set ??Enable in HA Passive State.??
- D. Set LACP mode to ??Active.??

**Answer: C**

#### Explanation:

In a High Availability (HA) active/passive pair configuration, when setting up an Aggregate Ethernet (AE) interface, enabling the "Enable in HA Passive State" option allows the interface to participate in LACP (Link Aggregation Control Protocol) even when the system is in the passive state. This ensures that the pre-negotiation of the LACP link occurs, allowing the link aggregation to be ready as soon as the firewall becomes active.

### NEW QUESTION 31

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: CD**

#### Explanation:

Separate rules must be created for each direction: Palo Alto Networks firewalls enforce security policies based on traffic direction. To allow bidirectional communication through the IPSec tunnel, two separate rules are required - one for incoming and one for outgoing traffic.

IKE negotiation and IPSec/ESP packets are denied by default: Palo Alto Networks firewalls use an interzone default deny policy, meaning that unless an explicit policy allows IKE (UDP 500/4500) and ESP (protocol 50) traffic, the firewall will block these packets, preventing tunnel establishment. Therefore, administrators must create explicit rules permitting IKE and IPSec/ESP traffic to the firewall's external interface.

### NEW QUESTION 36

How does a Palo Alto Networks NGFW respond when the preemptive hold time is set to 0 minutes during configuration of route monitoring?

- A. It does not accept the configuration.
- B. It accepts the configuration but throws a warning message.
- C. It removes the static route because 0 is a NULL value
- D. It reinstalls the route into the routing information base (RIB) as soon as the path comes up.

**Answer: D**

#### Explanation:

When the preemptive hold time is set to 0 minutes in route monitoring, the firewall is configured to immediately reinstall the route into the Routing Information Base (RIB) as soon as the monitored path comes up. This essentially means that the firewall will not wait for any predefined hold time before reestablishing the route once the monitoring condition is met, ensuring a faster recovery of the route.

### NEW QUESTION 41

To maintain security efficacy of its public cloud resources by using native tools, a company purchases Cloud NGFW credits to replicate the Panorama, PA-Series, and VM-Series devices used in physical data centers. Resources exist on AWS and Azure:

The AWS deployment is architected with AWS Transit Gateway, to which all resources connect

The Azure deployment is architected with each application independently routing traffic The engineer deploying Cloud NGFW in these two cloud environments must account for the following:

Minimize changes to the two cloud environments

Scale to the demands of the applications while using the least amount of compute resources

Allow the company to unify the Security policies across all protected areas Which two implementations will meet these requirements? (Choose two.)

- A. Deploy a VM-Series firewall in AWS in each VPC, create an IPSec tunnel between AWS and Azure, and manage the policy with Panorama.
- B. Deploy Cloud NGFW for Azure in vNET/s, update the vNET/s routing to path traffic through the deployed NGFWs, and manage the policy with Panorama.
- C. Deploy Cloud NGFW for Azure in vWAN, create a vWAN to route all appropriate traffic to the Cloud NGFW attached to the vWAN, and manage the policy with local rules.
- D. Deploy Cloud NGFW for AWS in a centralized Security VPC, update the Transit Gateway to route all appropriate traffic through the Security VPC, and manage the policy with Panorama.

**Answer: BD**

#### Explanation:

To meet the company's requirements - minimizing changes to the cloud environments, optimizing compute resources, and unifying security policies - the best approach is to deploy Cloud NGFW solutions natively for AWS and Azure while managing policies centrally with Panorama.

In Azure, using Cloud NGFW for Azure deployed within vNETs allows traffic to be routed through security appliances efficiently without requiring a complete re-architecture. This approach aligns with Azure's existing routing mechanism while maintaining security.

In AWS, deploying Cloud NGFW for AWS in a centralized Security VPC and integrating it with AWS Transit Gateway enables traffic inspection for all connected VPCs without modifying individual workloads. This method ensures efficient scaling and minimal infrastructure changes while maintaining security consistency.

### NEW QUESTION 42

In a Palo Alto Networks environment, GlobalProtect has been enabled using certificate-based authentication for both users and devices. To ensure proper

validation of certificates, one or more certificate profiles are configured.  
What function do certificate profiles serve in this context?

- A. They store private keys for users and devices, effectively allowing the firewall to issue or reissue certificates if the primary Certificate Authority (CA) becomes unavailable, providing a built-in fallback CA to maintain continuous certificate issuance and authentication.
- B. They define trust anchors (root / intermediate Certificate Authorities (CAs)), specify revocation checks (CRL/OCSP), and map certificate attributes (e.g., CN) for user or device authentication.
- C. They allow the firewall to bypass certificate validation entirely, focusing only on username / password-based authentication.
- D. They provide a one-click mechanism to distribute certificates to all endpoints without relying on external enrollment methods.

**Answer: B**

**Explanation:**

In the context of GlobalProtect with certificate-based authentication, certificate profiles are used to ensure proper validation of the certificates. They perform the following functions: Define trust anchors, which are the root and intermediate Certificate Authorities (CAs) that the firewall trusts to authenticate certificates. Specify revocation checks, such as CRL (Certificate Revocation List) and OCSP (Online Certificate Status Protocol), to ensure that the certificates being used have not been revoked. Map certificate attributes, such as the Common Name (CN), which helps in authenticating users and devices based on their certificates.

**NEW QUESTION 47**

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: CD**

**Explanation:**

To implement post-quantum cryptography (PQC) in VPNs between Palo Alto Networks NGFWs, you would enable the PQ KEM (Post-Quantum Key Encapsulation Mechanism) in the IKE gateway configuration. This enables the firewall to use quantum-resistant encryption for key exchange, which is an essential part of securing communications against the potential future threats posed by quantum computing. By selecting IKE v2 Preferred and enabling the PQ KEM option under Advanced Options, you can add specific Rounds for the post-quantum cryptography process, which will help in implementing quantum-resistant key exchange methods. This option similarly selects IKE v2 and enables PQ KEM while also creating a dedicated IKE Crypto Profile with the necessary Rounds configured for post-quantum cryptography.

**NEW QUESTION 48**

An NGFW engineer is establishing bidirectional connectivity between the accounting virtual system (VSYS) and the marketing VSYS. The traffic needs to transition between zones without leaving the firewall (no external physical connections). The interfaces for each VSYS are assigned to separate virtual routers (VRs), and inter-VR static routes have been configured. An external zone has been created correctly for each VSYS. Security policies have been added to permit the desired traffic between each zone and its respective external zone. However, the desired traffic is still unable to successfully pass from one VSYS to the other in either direction.

Which additional configuration task is required to resolve this issue?

- A. Create a transit VSYS and route all inter-VSYS traffic through it.
- B. Add each VSYS to the list of visible virtual systems of the other VSYS.
- C. Enable the ??allow inter-VSYS traffic?? option in both external zone configurations.
- D. Create Security policies to allow the traffic between the two external zones.

**Answer: B**

**Explanation:**

In Palo Alto Networks firewalls, each virtual system (VSYS) is typically isolated from other VSYSs, meaning that traffic between different VSYSs cannot pass through the firewall by default. In this case, since the interfaces for each VSYS are assigned to separate virtual routers (VRs), and the desired traffic is still not passing between the two VSYSs, the firewall needs to be explicitly configured to allow traffic between them. The required configuration is to add each VSYS to the list of visible virtual systems of the other VSYS. This allows inter-VSYS communication to be enabled, effectively permitting the traffic to pass between the zones of different VSYSs.

**NEW QUESTION 50**

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 building a custom report on a Palo Alto Networks NGFW, you can select detailed logs that provide specific insights into various aspects of firewall activity. The available options for detailed logs typically include:  
Traffic logs: These provide information on the network traffic passing through the firewall. Threat logs: These logs capture data related to identified security threats, such as malware or intrusion attempts.  
Data filtering logs: These logs capture events related to data filtering policies, such as preventing the transfer of sensitive data.  
User-ID logs: These logs associate user identities with the traffic and activities observed on the firewall, enabling user-based policy enforcement.

**NEW QUESTION 55**

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### NGFW-Engineer Practice Exam Features:

- \* NGFW-Engineer Questions and Answers Updated Frequently
- \* NGFW-Engineer Practice Questions Verified by Expert Senior Certified Staff
- \* NGFW-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* NGFW-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
[Order The NGFW-Engineer Practice Test Here](#)