Boost up Your Certification Score

Palo Alto Networks
NGFW-Engineer
Palo Alto Networks Next-Generation Firewall Engineer

For More Information – Visit link below:

https://www.examsboost.com/

Product Version

✓ Up to Date products, reliable and verified.
✓ Questions and Answers in PDF Format.

Latest Version: 6.0

Question: 1

A network engineer is configuring a new Palo Alto Networks firewall to segment a flat corporate network into multiple VLANs. The firewall has an interface, ethernet1/1, intended to carry traffic for VLANs 10 (HR) and 20 (Finance). The HR VLAN requires DHCP services provided by an external server accessible through a Layer 3 interface on the firewall, while the Finance VLAN needs static IP assignment for critical servers. The firewall is also expected to perform inter-VLAN routing for these segments. Which of the following configurations, when applied to ethernet1 11 , correctly prepares it for this scenario, assuming no other relevant interfaces are configured yet?

A. Configure ethernet1/1 as a Layer 3 interface, then create subinterfaces ethernet1/1 .10 and ethernet1/1 .20, assigning them to VLANs 10 and 20 respectively, and configure IP addresses on these subinterfaces.
B. Configure ethernet1/1 as a Layer 2 interface, create VLAN interfaces for VLAN 10 and VLAN 20, and assign ethernet1/1 to a virtual wire with these VLAN interfaces.
C. Configure ethernet1/1 as a Layer 2 interface, add it to a new VLAN object named 'Corporate-VLANs', and then create two subinterfaces on this VLAN object, one for VLAN 10 and one for VLAN 20, each with an IP address.
D. Configure ethernet1/1 as a Layer 2 interface, create two VLAN objects (VLAN-IO and VLAN-20), assign ethernet1/1 as an access port to both VLAN-IO and VLAN-20, and then create Layer 3 VLAN interfaces for routing.
E. Configure ethernet1/1 as a Layer 2 interface, then create a single VLAN interface and add ethernet1/1 as a trunk port to this VLAN interface. Subsequently, create two sub-interfaces on this VLAN interface, one for VLAN 10 and one for VLAN 20, assigning IP addresses.

Answer: A

Explanation:
For inter-VLAN routing and providing DHCP services (which require a Layer 3 presence), the primary interface (ethernet1/1) should be configured as a Layer 3 interface. Subinterfaces (e.g., ethernet1/1 .10 and ethernet1/1 .20) are then created on this Layer 3 interface, each tagged with a specific VLAN ID and assigned its own IP address to serve as the default gateway for that VLAN. This enables the firewall to route traffic between the VLANs and allows for Layer 3 services like DHCP relay to function. Options B, C, D, and E describe incorrect or incomplete configurations for inter-VLAN routing and Layer 3 services in this context.

Question: 2

A Palo Alto Networks firewall is deployed as a transparent bridge in a network segment. Interface ethernet1/2 is configured as a Layer 2 interface and is part of a Bridge Group. The network administrator observes that broadcast traffic from a specific legacy device on ethernet1/2 is flooding other interfaces within the same bridge group, causing performance issues. The administrator needs to prevent this

specific device's broadcast traffic from flooding, without blocking unicast communication or impacting other devices on the same segment. Which of the following configuration steps would be most effective to mitigate this issue, assuming the legacy device's MAC address is 00:0A:95:9D:68:1B?

A. Implement a firewall policy to deny all broadcast traffic sourced from 00:0A:95:9D:68:1B within the Bridge Group.
B. Configure a MAC ACL (Access Control List) on the ethernet1/2 interface to drop packets with source MAC 00:0A:95:9D:68:1B and a broadcast destination MAC.
C. Change ethernet1/2 from a Layer 2 interface to a Layer 3 interface and configure an ARP entry for the legacy device to prevent broadcast lookup.
D. Utilize the 'MAC Limiting' feature on the Bridge Group to cap the number of MAC addresses learned, forcing the firewall to drop excessive broadcasts.
E. On the ethernet1/2 interface configuration, navigate to the Advanced tab and configure a Broadcast Storm Control threshold for the specific MAC address

Answer: B

Explanation:
While Palo Alto Networks firewalls are not typical Layer 2 switches with granular broadcast control features per MAC address, a creative approach using a MAC-based Access Control List (MAC ACL) on the Layer 2 interface is the most direct and effective way to address specific broadcast flooding from a known MAC. By dropping packets with the problematic source MAC and a broadcast destination, the flooding is mitigated without affecting unicast traffic. Options A is incorrect because firewall policies operate at Layer 3/4. Option C would break the transparent bridge functionality. Option D is for general MAC learning limits, not specific broadcast control. Option E does not exist in PAN-OS in the described manner for specific MAC addresses.

Question: 3

A Palo Alto Networks firewall acts as a Layer 2 transparent bridge between two critical network segments (Segment A and Segment B) for a secure data transfer application. Interface ethernet1/3 is connected to Segment A, and ethernet1/4 is connected to Segment B. Both interfaces are part of the same Bridge Group 'BG-DataTransfer'. The application requires strict latency control and minimal jitter. After deployment, the network team reports intermittent latency spikes and packet retransmissions, particularly during periods of high data throughput. Further investigation reveals that the application servers on Segment A occasionally send very large frames (up to 9216 bytes) that are being dropped by the firewall. The application developers confirm that they are using Jumbo Frames. What is the most precise and comprehensive configuration change to ensure seamless Layer 2 operation with Jumbo Frames and optimize for low latency in this transparent bridge scenario?

A. Increase the MTU on both ethernet1/3 and ethernet1/4 interfaces to 9216 bytes. Then, ensure the Bridge Group 'BG-DataTransfer' is configured with an appropriate MTU, and verify that the firewall's physical interfaces support Jumbo Frames.
B. Set the MTU on both ethernet1/3 and ethernet1/4 to 1500 bytes and implement TCP MSS clamping on the Bridge Group to prevent oversized packets from being sent.
C. Configure 'Bypass Mode' on the Bridge Group 'BG-DataTransfer' to ensure that all traffic, including Jumbo Frames, is passed through without inspection, thereby reducing latency.

D. On both ethernet1/3 and ethernet1/4, configure the MTU to 9216 bytes. Additionally, review the firewall's Threat Prevention and Security Policies for any profiles that might be performing deep packet inspection or reassembly, and adjust them to minimize latency for this specific traffic, or create a bypass rule for known Jumbo Frame traffic.
E. The Palo Alto Networks firewall automatically handles Jumbo Frames in Layer 2 mode. The latency issue is likely due to CPU oversubscription. Monitor the firewall's CPU usage and consider upgrading the hardware if necessary.

Answer: A

Explanation:
For Jumbo Frames to traverse a Palo Alto Networks firewall in Layer 2 mode, the MTU on the participating Layer 2 interfaces (ethernet1/3 and ethernet1/4 in this case) MUST be increased to accommodate the larger frame size (e.g., 9216 bytes). While the firewall itself might inherently support Jumbo Frames on its physical ports, the configured interface MTU is paramount. The Bridge Group also needs to have its MTU adjusted if it has a global MTU setting that could override or restrict. Options B and C are incorrect as they either prevent Jumbo Frames or bypass security. Option D is partially correct but the core issue is MTU configuration, not primarily security policy deep inspection. Option E is incorrect because Jumbo Frame support is not automatic without explicit MTU configuration, and while CPU can be a factor, it's not the primary cause of frame drops due to size.

Question: 4

A Palo Alto Networks firewall needs to act as a Layer 2 bridge for a specialized industrial control network that utilizes custom EtherType protocols in addition to standard IP traffic. Interface ethernet1/5 is connected to this network. The operational requirement is to allow all traffic, including these custom Ether Type frames, to pass through the firewall transparently, while still enforcing security policies only on standard IP and ARP traffic. The firewall must log all dropped packets, even those not matching any security policy. Which of the following configuration adjustments are necessary and sufficient to meet these requirements?

A. Configure ethernet1/5 as a Layer 2 interface and add it to a Bridge Group. Create a security policy to permit all IP and ARP traffic. For custom EtherType traffic, a dedicated firewall policy is not needed as Layer 2 interfaces inherently forward non-IP/ARP traffic.
B. Configure ethernet1/5 as a Layer 2 interface and add it to a Bridge Group. Create a security policy with a 'catch-all' rule for IP and ARP traffic. Additionally, create a custom application signature for each EtherType protocol and then create security policies to permit traffic matching these custom applications.
C. Configure ethernet1/5 as a Layer 2 interface and add it to a Bridge Group. Ensure the 'ethernet type forwarding' option is enabled on the Bridge Group or individual interface. Create security policies for IP and ARP traffic. Configure the default inter-zone policy for the bridge group's zones to 'allow all' to explicitly permit custom EtherTypes to pass uninspected.
D. Configure ethernet1/5 as a Layer 2 interface and add it to a Bridge Group. Create security policies to permit IP and ARP traffic. To log all dropped packets, ensure the implicit 'deny' rule at the end of the security policy rulebase has logging enabled. Custom Ether Type traffic will bypass security processing if no explicit rules match them, but will still be forwarded.

E. Configure ethernet1/5 as a Layer 2 interface. Create a Bridge Group and add ethernet1/5. Crucially, in a Layer 2 transparent deployment, non-IP/ARP traffic (including custom Ether Types) is typically forwarded without being subjected to security policies by default, unless specific Packet Filter rules are configured. To log all dropped packets, enable 'Log at Session End' on the implicit deny rule. For any dropped IPIARP packets not matching specific rules, this ensures logging.

Answer: D,E

Explanation:
This question tests the understanding of how Palo Alto Networks firewalls handle non-IP/ARP traffic in Layer 2 mode and logging. Options D and E are the most accurate. Non-IP/ARP Traffic (Custom Ether Types): In Layer 2 (transparent) mode, Palo Alto firewalls, by default, forward non-IP/ARP traffic (like custom Ether Types) without subjecting it to security policies. This is because security policies primarily operate at Layer 3 and above. There's no explicit 'ethernet type forwarding' option to enable/disable for this general behavior, it's inherent to Layer 2 bridging. Creating custom application signatures for Ether Types (Option B) is generally not how you manage forwarding of basic Layer 2 frames; it's for Layer 3+ application identification. Enabling a catch-all policy for custom EtherTypes (Option C) is also not the standard way as they are not subject to standard security policies. Logging All Dropped Packets: To log ALL dropped packets, it's essential to enable logging on the implicit 'deny' rule that exists at the end of every security policy rulebase. This rule catches any traffic (IPIARP) that hasn't been explicitly permitted by preceding rules and drops it. Logging on this rule ensures that you see why packets are being dropped by the firewall's security engine. Therefore, configuring the Layer 2 interface, setting up the Bridge Group, creating specific security policies for IP/ARP, and enabling logging on the implicit deny rule covers all requirements.

Question: 5

A Palo Alto Networks firewall is being configured for a multi-tenant environment. Interface ethernet1/6 is a Layer 2 interface connected to a core switch that carries traffic for several tenants, each segregated by a unique VLAN (e.g., VLAN 100 for Tenant A, VLAN 200 for Tenant B). The security requirement dictates that traffic between tenants (inter-VLAN traffic) must be strictly isolated and pass through a separate, dedicated Layer 3 firewall (not this device), while intra-tenant traffic (traffic within the same VLAN) should be allowed to traverse this Palo Alto firewall for security inspection and logging, specifically for malware detection. This firewall must operate as a transparent Layer 2 device for each tenant's VLAN. Which of the following configurations are necessary to achieve this intricate setup, considering the need for granular intra-VLAN security inspection?

A. Configure ethernet1/6 as a Layer 2 interface. Create a separate Bridge Group for each tenant's VLAN (e.g., BG-TenantA for VLAN 100, BG-TenantB for VLAN 200). Add ethernet1/6 to all these Bridge Groups. Configure security policies between the zones associated with these Bridge Groups to allow intra-tenant traffic and deny inter-tenant traffic.
B. Configure ethernet1 16 as a Layer 2 interface. Create a single Bridge Group and add ethernet1/6 to it. For each tenant's VLAN, create a separate VLAN interface (e.g., VLAN-100, VLAN-200) and assign them to the Bridge Group. Configure a virtual wire between each VLAN interface and a corresponding Layer 3 interface for security enforcement.
C. Configure ethernet1/6 as a Layer 2 interface. Create a separate Bridge Group for each tenant's VLAN (e.g., BG-TenantA for VLAN 100, BG-TenantB for VLAN 200). Add ethernet1/6 to each respective Bridge

Group, ensuring the 'VLAN Tagging' is set for each Bridge Group to match the tenant's VLAN ID. Then, define security zones for each Bridge Group and create security policies within each zone to allow intra-tenant traffic. Inter-tenant traffic will not traverse as separate Bridge Groups do not route between themselves by default.
D. Configure ethernet1/6 as a Layer 2 interface. Create a single Bridge Group. Create a 'VLAN Tag' object for each tenant's VLAN (e.g., VLAN-IOO, VLAN-200). Assign these VLAN Tags to the Bridge Group. Implement security policies using 'Source Zone' and 'Destination Zone' based on the Bridge Group and 'Source VLAN' and 'Destination VLAN' to permit intra-VLAN traffic and deny inter-VLAN traffic.
E. Configure ethernet1 16 as a Layer 2 interface. Create a separate VLAN Interface for each tenant's VLAN (e.g., VLAN-IOO for VLAN 1 00, VLAN-200 for VLAN 200) and assign ethemet1/6 as an access port to each VLAN Interface. Create security zones for each VLAN Interface and configure security policies between these zones to allow intra-VLAN traffic and deny inter-VLAN traffic.

Answer: C

Explanation:
This scenario requires the Palo Alto Networks firewall to operate as a transparent Layer 2 device for each VLAN independently, allowing intra-VLAN traffic through for inspection, while implicitly preventing inter-VLAN routing on this device. The key is using separate Bridge Groups, each configured to specifically handle traffic for a single VLAN ID. Option C correctly identifies this. By creating a separate Bridge Group for each VLAN (e.g., BG-TenantA with VLAN Tag 100, BG-TenantB with VLAN Tag 200) and adding ethernet1/6 to each of these Bridge Groups, the firewall will logically separate the VLAN traffic at Layer 2. Intra-VLAN traffic within BG-TenantA will be inspected by policies within its zone, and similarly for BG-TenantB. Since Bridge Groups do not route between themselves by default, inter-VLAN traffic (e.g., from VLAN 100 to VLAN 200) will not traverse this firewall's security engine, fulfilling the requirement for it to be handled by a separate Layer 3 firewall. Option A is incorrect because adding ethernet1/6 to all Bridge Groups without proper VLAN tagging per Bridge Group would cause confusion. It also incorrectly suggests inter-zone policies to deny inter-tenant traffic when the goal is to prevent it from crossing this firewall entirely at Layer 2. Option B describes a hybrid setup with VLAN interfaces and virtual wires, which is more complex than needed for a pure Layer 2 bridge per VLAN and doesn't explicitly address the inter-VLAN isolation requirement at this device. Option D uses a single Bridge Group with VLAN Tags, but security policies based on 'Source VLAN' and 'Destination VLAN' are typically for Layer 3 inter-VLAN routing scenarios, not for strict Layer 2 isolation where traffic should not even be processed for routing by this firewall. Option E uses VLAN Interfaces, which are Layer 3 constructs. While they can have an associated Layer 2 interface, their primary purpose is Layer 3 routing, which contradicts the requirement for the firewall to act as a transparent Layer 2 device for each tenant's VLAN and explicitly not route between them.

Question: 6

A Palo Alto Networks firewall is providing Layer 2 segmentation within a data center. Interface ethernet1/7 is a Layer 2 interface connected to a virtualized server environment, carrying traffic from multiple virtual machines (VMs) on different VLANs. The security team has mandated that no VM should be able to spoof the MAC address of another VM on the same physical segment. Additionally, if a VM tries to use an unassigned IP address within its VLAN subnet (e.g., an IP not explicitly configured for it), the firewall should drop that traffic without impacting legitimate communication. How can these

requirements be most effectively enforced using PAN-OS Layer 2 features, assuming dynamic MAC learning is enabled?

A. On ethernet1/7, enable 'ARP Inspection' and configure 'Static ARP Entries' for all legitimate VM MAC-to-IP mappings. Also, enable 'MAC Limiting' on the Bridge Group to prevent new MAC addresses from being learned.
B. Configure 'DHCP Snooping' on ethernet1/7 and enable 'Dynamic ARP Inspection' (DAI) on the associated Bridge Group. Set up 'IP Source Guard' policies on the Bridge Group to drop packets with source IPs not learned via DHCP or static entries.
C. Enable 'ARP Inspection' on ethernet1/7 and populate its 'ARP Trust Table' with legitimate MAC-to-IP bindings for each VM. For MAC spoofing, enable 'MAC Security' on ethernet1/7 and define allowed MAC addresses. For unassigned IP addresses, leverage 'IP-MAC Binding' and configure security rules to drop traffic not matching these bindings.
D. Utilize 'MAC-based Forwarding' on ethernet1/7 to restrict traffic to known MAC addresses. Implement 'Dynamic ARP Inspection' (DAI) on the Bridge Group, with a trusted port for the upstream switch and untrusted ports for the VMs, to prevent IP address spoofing for unassigned IPs.
E. On the Bridge Group containing ethernet1/7, enable 'ARP Inspection' and ensure 'Validate ARP' is checked. This helps prevent MAC spoofing. To address unassigned IP addresses, enable TIP-MAC Binding' and either statically configure or dynamically learn the valid IP-to-MAC associations. Then, ensure the security policy drops traffic from unlearned/invalid IP-MAC pairs.

Answer: E

Explanation:
This scenario focuses on Layer 2 security features to combat MAC and IP spoofing within a Layer 2 segment. MAC Spoofing: 'ARP Inspection' on the Bridge Group, specifically with 'Validate ARP' enabled, is the primary mechanism in PAN-OS Layer 2 mode to prevent MAC address spoofing. It validates ARP requests/replies against learned or statically configured ARP entries, dropping invalid ones. This helps ensure that the MAC address advertised in ARP corresponds to the correct IP. Unassigned IP Addresses (IP Spoofing): 'IP-MAC Binding' is the feature used to enforce that specific IP addresses are only used by their legitimate MAC addresses. This can be populated statically or dynamically (e.g., through DHCP snooping if enabled). Once binding is established, the firewall can drop traffic where the source IP address does not match the bound MAC address for that port/VLAN. This effectively prevents a VM from using an IP address that hasn't been assigned to it or is not valid for its MAC. Let's analyze why other options are less optimal or incorrect: Option A: 'MAC Limiting' is about preventing MAC table overflow, not directly preventing MAC spoofing by a specific MAC address. 'Static ARP Entries' are useful but less dynamic for a large VM environment. ARP Inspection itself is good, but the overall solution is incomplete for the IP aspect. Option B: While DHCP Snooping and DAI are related, IP Source Guard is typically a Cisco term, and the direct PAN-OS equivalent for preventing source IP spoofing on Layer 2 is 'IP- MAC Binding' in conjunction with ARP inspection. Option C: 'MAC Security' as described (defining allowed MACs) is more akin to port security on a switch and isn't the primary PAN-OS Layer 2 feature for preventing spoofing through validation . IP-MAC Binding is correct for the IP part, but the MAC spoofing prevention is better handled by ARP Inspection. Option D: 'MAC-based Forwarding' isn't a direct feature for spoofing prevention in the described manner. DAI setup with trusted/untrusted ports is more granular, but the core PAN-OS features for Layer 2 IP/MAC validation are better represented by Option E.

Question: 7

A network engineer is configuring a new Palo Alto Networks firewall to segment an internal network. The firewall needs to connect to an existing Layer 3 switch and route traffic between the 'Servers' VLAN (VLAN ID 100) and the 'Clients' VLAN (VLAN ID 200). The physical interface 'ethernet1/1' on the firewall is connected to a trunk port on the Layer 3 switch. What is the correct PAN-OS CLI configuration to enable routing for both VLANs on 'ethernet1/1' and assign appropriate IP addresses for gateway functionality, assuming the Servers VLAN uses 10.0.100.1/24 and Clients VLAN uses 10.0.200.1/24?

A.
```
set network interface ethernet ethernet1/1 layer3 ip 10.0.100.1/24
set network interface ethernet ethernet1/1 layer3 ip 10.0.200.1/24
commit
```
B.
```
set network interface ethernet ethernet1/1 layer3
set network interface ethernet ethernet1/1 layer3 unit 100 ip 10.0.100.1/24
set network interface ethernet ethernet1/1 layer3 unit 200 ip 10.0.200.1/24
commit
```
C.
```
set network interface ethernet ethernet1/1 layer3
set network interface ethernet ethernet1/1 layer3 unit 100 tag 100 ip 10.0.100.1/24
set network interface ethernet ethernet1/1 layer3 unit 200 tag 200 ip 10.0.200.1/24
commit
```
D.
```
set network interface ethernet ethernet1/1 layer3
set network interface ethernet ethernet1/1 layer3 unit 100 ip 10.0.100.1/24
set network interface ethernet ethernet1/1 layer3 unit 200 ip 10.0.200.1/24
set network interface ethernet ethernet1/1 layer3 unit 100 tag 100
set network interface ethernet ethernet1/1 layer3 unit 200 tag 200
commit
```
E.
```
set network interface ethernet ethernet1/1 layer3
set network interface ethernet ethernet1/1 layer3 unit 100 virtual-router default ip 10.0.100.1/24
set network interface ethernet ethernet1/1 layer3 unit 200 virtual-router default ip 10.0.200.1/24
commit
```

Answer: C

Explanation:
To configure multiple VLANs on a single physical interface for Layer 3 routing, you must create subinterfaces (units) and tag them with the corresponding VLAN IDs. Option C correctly creates Layer 3 subinterfaces (unit 100 and unit 200) on ethernet1/1 , assigns the correct IP addresses, and associates them with the respective VLAN tags (100 and 200). Option B is incorrect because it misses the 'tag'

command which is crucial for VLAN identification. Option A attempts to assign multiple IPs directly to the physical interface, which is not how VLAN routing works. Options D and E have incorrect command syntax or unnecessary commands for this scenario.

Question: 8

A critical application server in the DMZ (172.16.10.100) needs to communicate with an internal database server (192.168.1.50) located behind a Palo Alto Networks firewall. The firewall has 'ethernet1/2' configured as Layer 3 for the DMZ (172.16.10.1/24) and 'ethernet1/3' as Layer 3 for the Internal network (192.168.1.1/24). Due to compliance requirements, all traffic originating from the DMZ towards the Internal network must be sourced from a specific IP address (172.16.10.254) on the firewall's DMZ interface, regardless of the server's actual source IP, while preserving the destination. Which of the following PAN-OS configurations, when combined, would achieve this specific source NAT behavior and ensure connectivity?

A.

```
set rulebase nat rules 'DMZ-to-Internal-Sourcing'
set rulebase nat rules 'DMZ-to-Internal-Sourcing' from zone DMZ to zone Internal
set rulebase nat rules 'DMZ-to-Internal-Sourcing' source-address any
set rulebase nat rules 'DMZ-to-Internal-Sourcing' destination-address 192.168.1.50
set rulebase nat rules 'DMZ-to-Internal-Sourcing' service any
set rulebase nat rules 'DMZ-to-Internal-Sourcing' nat-type ipv4
set rulebase nat rules 'DMZ-to-Internal-Sourcing' source-translation dynamic-ip-and-port translated-address 172.16.10.254
set rulebase nat rules 'DMZ-to-Internal-Sourcing' destination-translation type none
commit
```

B.

```
set rulebase nat rules 'DMZ-to-Internal-Sourcing'
set rulebase nat rules 'DMZ-to-Internal-Sourcing' from zone DMZ to zone Internal
set rulebase nat rules 'DMZ-to-Internal-Sourcing' source-address 172.16.10.100
set rulebase nat rules 'DMZ-to-Internal-Sourcing' destination-address 192.168.1.50
set rulebase nat rules 'DMZ-to-Internal-Sourcing' service any
set rulebase nat rules 'DMZ-to-Internal-Sourcing' nat-type ipv4
set rulebase nat rules 'DMZ-to-Internal-Sourcing' source-translation static-ip translated-address 172.16.10.254
set rulebase nat rules 'DMZ-to-Internal-Sourcing' destination-translation type none
commit
```

C.

```
set rulebase nat rules 'DMZ-to-Internal-Sourcing'
set rulebase nat rules 'DMZ-to-Internal-Sourcing' from zone DMZ to zone Internal
set rulebase nat rules 'DMZ-to-Internal-Sourcing' source-address 172.16.10.100
set rulebase nat rules 'DMZ-to-Internal-Sourcing' destination-address 192.168.1.50
set rulebase nat rules 'DMZ-to-Internal-Sourcing' service any
set rulebase nat rules 'DMZ-to-Internal-Sourcing' nat-type ipv4
set rulebase nat rules 'DMZ-to-Internal-Sourcing' source-translation dynamic-ip-and-port translated-address 172.16.10.254
set rulebase nat rules 'DMZ-to-Internal-Sourcing' destination-translation type none
commit

set security rules 'Allow-DMZ-to-Internal-DB'
set security rules 'Allow-DMZ-to-Internal-DB' from DMZ to Internal
set security rules 'Allow-DMZ-to-Internal-DB' source 172.16.10.254
set security rules 'Allow-DMZ-to-Internal-DB' destination 192.168.1.50
set security rules 'Allow-DMZ-to-Internal-DB' application any service any action allow
commit
```

D.

```
set rulebase nat rules 'DMZ-to-Internal-Sourcing'
set rulebase nat rules 'DMZ-to-Internal-Sourcing' from zone DMZ to zone Internal
set rulebase nat rules 'DMZ-to-Internal-Sourcing' source-address 172.16.10.100
set rulebase nat rules 'DMZ-to-Internal-Sourcing' destination-address 192.168.1.50
set rulebase nat rules 'DMZ-to-Internal-Sourcing' service any
set rulebase nat rules 'DMZ-to-Internal-Sourcing' nat-type ipv4
set rulebase nat rules 'DMZ-to-Internal-Sourcing' source-translation static-ip translated-address 172.16.10.254
set rulebase nat rules 'DMZ-to-Internal-Sourcing' destination-translation type none
commit

set security rules 'Allow-DMZ-to-Internal-DB'
set security rules 'Allow-DMZ-to-Internal-DB' from DMZ to Internal
set security rules 'Allow-DMZ-to-Internal-DB' source 172.16.10.254
set security rules 'Allow-DMZ-to-Internal-DB' destination 192.168.1.50
set security rules 'Allow-DMZ-to-Internal-DB' application any service any action allow
commit
```

E.

```
set network interface ethernet ethernet1/2 layer3 ip 172.16.10.1/24
set network interface ethernet ethernet1/3 layer3 ip 192.168.1.1/24
set zone DMZ network ethernet1/2
set zone Internal network ethernet1/3

set rulebase nat rules 'DMZ-to-Internal-Sourcing'
set rulebase nat rules 'DMZ-to-Internal-Sourcing' from zone DMZ to zone Internal
set rulebase nat rules 'DMZ-to-Internal-Sourcing' source-address 172.16.10.100
set rulebase nat rules 'DMZ-to-Internal-Sourcing' destination-address 192.168.1.50
set rulebase nat rules 'DMZ-to-Internal-Sourcing' service any
set rulebase nat rules 'DMZ-to-Internal-Sourcing' nat-type ipv4
set rulebase nat rules 'DMZ-to-Internal-Sourcing' source-translation dynamic-ip translated-address 172.16.10.254
set rulebase nat rules 'DMZ-to-Internal-Sourcing' destination-translation type none
commit

set security rules 'Allow-DMZ-to-Internal-DB'
set security rules 'Allow-DMZ-to-Internal-DB' from DMZ to Internal
set security rules 'Allow-DMZ-to-Internal-DB' source 172.16.10.254
set security rules 'Allow-DMZ-to-Internal-DB' destination 192.168.1.50
set security rules 'Allow-DMZ-to-Internal-DB' application any service any action allow
commit
```

Answer: D

Explanation:

This scenario requires a specific type of Source NAT and a corresponding security policy. 'Static IP' source translation maps a specific internal source IP to a specific external source IP, without changing the port. 'Dynamic IP' (or 'Dynamic IP and Port') maps to an IP, but the port is also dynamically translated, which is not explicitly required if we just need a specific source IP. However, the critical part is that the security policy must allow traffic from the translated source IP (172.16.10.254), not the original source IP (172.16.10.100), because NAT occurs before security policy evaluation. Option D uses ‘static-ip’ source translation, which is suitable for a fixed source IP, and correctly updates the security policy source to the post-NAT IP (172.16.10.254). Option C uses 'dynamic-ip-and-port , which might change the source port, and while it also updates the security policy, ‘static-ips is more precise for this requirement. Option E uses 'dynamic-ip’ which is effectively 'dynamic-ip-and- port' but the initial interface/zone configurations are already assumed to be in place for the question context.

Question: 9

A large enterprise is migrating its network infrastructure to a new Palo Alto Networks firewall. They have a complex routing policy where certain internal subnets (e.g., 10.10.0.0/16, 10.20.0.0/16) must egress through a primary internet connection (ISP-A) while all other internet-bound traffic (0.0.0.0/0) must use a secondary internet connection (ISP-B) as a default route. Both ISP connections terminate on separate Layer 3 interfaces of the firewall: 'ethernet1/1' for ISP-A and 'ethernet1/2' for ISP-B. Assuming both interfaces are in the 'External' zone and assigned to the 'default' virtual router. How would you configure the routing to achieve this specific routing hierarchy using PAN-OS CLI, considering the need for path monitoring to ensure failover capability for ISP-A traffic?

A.

```
set network virtual-router default routing static-route 'ISP-A-Specific'
set network virtual-router default routing static-route 'ISP-A-Specific' destination 10.10.0.0/16 next-hop ip-address X.X.X.X interface ethernet1/1
set network virtual-router default routing static-route 'ISP-A-Specific' destination 10.20.0.0/16 next-hop ip-address X.X.X.X interface ethernet1/1
set network virtual-router default routing static-route 'ISP-B-Default'
set network virtual-router default routing static-route 'ISP-B-Default' destination 0.0.0.0/0 next-hop ip-address Y.Y.Y.Y interface ethernet1/2
commit
```

B.

```
set network virtual-router default routing static-route 'ISP-A-Specific-1'
set network virtual-router default routing static-route 'ISP-A-Specific-1' destination 10.10.0.0/16 next-hop ip-address X.X.X.X interface ethernet1/1 metric 10
set network virtual-router default routing static-route 'ISP-A-Specific-2'
set network virtual-router default routing static-route 'ISP-A-Specific-2' destination 10.20.0.0/16 next-hop ip-address X.X.X.X interface ethernet1/1 metric 10
set network virtual-router default routing static-route 'ISP-B-Default'
set network virtual-router default routing static-route 'ISP-B-Default' destination 0.0.0.0/0 next-hop ip-address Y.Y.Y.Y interface ethernet1/2 metric 20
commit
```

C.

```
set network virtual-router default routing static-route 'ISP-A-Specific-1'
set network virtual-router default routing static-route 'ISP-A-Specific-1' destination 10.10.0.0/16 next-hop ip-address X.X.X.X interface ethernet1/1 metric 10
set network virtual-router default routing static-route 'ISP-A-Specific-1' path-monitor enable
set network virtual-router default routing static-route 'ISP-A-Specific-1' path-monitor dest Z.Z.Z.Z
set network virtual-router default routing static-route 'ISP-A-Specific-2'
set network virtual-router default routing static-route 'ISP-A-Specific-2' destination 10.20.0.0/16 next-hop ip-address X.X.X.X interface ethernet1/1 metric 10
set network virtual-router default routing static-route 'ISP-A-Specific-2' path-monitor enable
set network virtual-router default routing static-route 'ISP-A-Specific-2' path-monitor dest Z.Z.Z.Z
set network virtual-router default routing static-route 'ISP-B-Default'
set network virtual-router default routing static-route 'ISP-B-Default' destination 0.0.0.0/0 next-hop ip-address Y.Y.Y.Y interface ethernet1/2 metric 20
commit
```

D.

```
set network virtual-router default routing static-route 'ISP-A-Specific-1'
set network virtual-router default routing static-route 'ISP-A-Specific-1' destination 10.10.0.0/16 next-hop ip-address X.X.X.X interface ethernet1/1
set network virtual-router default routing static-route 'ISP-A-Specific-2'
set network virtual-router default routing static-route 'ISP-A-Specific-2' destination 10.20.0.0/16 next-hop ip-address X.X.X.X interface ethernet1/1
set network virtual-router default routing static-route 'ISP-A-Failover'
set network virtual-router default routing static-route 'ISP-A-Failover' destination 10.10.0.0/16 next-hop ip-address Y.Y.Y.Y interface ethernet1/2 metric 20
set network virtual-router default routing static-route 'ISP-A-Failover' destination 10.20.0.0/16 next-hop ip-address Y.Y.Y.Y interface ethernet1/2 metric 20
set network virtual-router default routing static-route 'ISP-B-Default'
set network virtual-router default routing static-route 'ISP-B-Default' destination 0.0.0.0/0 next-hop ip-address Y.Y.Y.Y interface ethernet1/2
commit
```

E.

```
set network virtual-router default routing static-route 'ISP-A-Specific-1'
set network virtual-router default routing static-route 'ISP-A-Specific-1' destination 10.10.0.0/16 next-hop ip-address X.X.X.X interface ethernet1/1 metric 10
set network virtual-router default routing static-route 'ISP-A-Specific-1' path-monitor enable
set network virtual-router default routing static-route 'ISP-A-Specific-1' path-monitor dest Z.Z.Z.Z interval 2 timeout 10 action fail-over
set network virtual-router default routing static-route 'ISP-A-Specific-2'
set network virtual-router default routing static-route 'ISP-A-Specific-2' destination 10.20.0.0/16 next-hop ip-address X.X.X.X interface ethernet1/1 metric 10
set network virtual-router default routing static-route 'ISP-A-Specific-2' path-monitor enable
set network virtual-router default routing static-route 'ISP-A-Specific-2' path-monitor dest Z.Z.Z.Z interval 2 timeout 10 action fail-over
set network virtual-router default routing static-route 'ISP-B-Default'
set network virtual-router default routing static-route 'ISP-B-Default' destination 0.0.0.0/0 next-hop ip-address Y.Y.Y.Y interface ethernet1/2 metric 20
commit
```

Answer: C,E

Explanation:
This scenario requires specific static routes for certain subnets with lower metrics to prioritize them, and a default route for all other traffic. Crucially, path monitoring is needed for the primary ISP-A routes to ensure failover. Both C and E correctly configure the specific routes for 10.10.0.0/16 and 10.20.0.0/16 via ISP-A (ethernet1/1) with a lower metric (10), and a default route via ISP-B (ethernet1/2) with a higher metric (20), ensuring the specific routes are preferred. They also correctly enable path monitoring for the ISP-A specific routes, which is essential for failover. The 'action fail-over' in option E is implicit with path-monitoring for static routes and is not strictly required in the CLI command for basic functionality, but 'interval' and ‘timeout can refine the monitoring. Option C is sufficient to enable the path monitoring with default parameters. The question asks 'how would you configure', implying a functional configuration, and both fulfill the core requirements. Without the explicit ‘action fail-over' option specified in the 'set' command, the default behavior for path monitoring on static routes is to remove the route from the FIB upon failure, thus allowing a higher metric route (if available) to take over, which is the desired failover.

Question: 10

A cybersecurity firm is performing a highly sensitive penetration test. They require absolute segregation of their testing environment from the client's production network, even though both must share a physical firewall interface. The Palo Alto Networks firewall's 'ethernet1/4' is connected to a trunk port on a switch. The penetration test environment uses VLAN 100 (192.168.100.1/24) and must ONLY be able to communicate with the client's 'Staging' network (VLAN 200, 192.168.200.1/24) through the firewall. The client's 'Production' network (VLAN 300, 192.168.300.1/24) also uses 'ethernet1/4' but must have no direct or indirect communication with VLAN 100. What is the most secure and appropriate Layer 3 interface and security zone configuration for this scenario in PAN-OS to enforce the strict segregation?

A.

```
set network interface ethernet ethernet1/4 layer3
set network interface ethernet ethernet1/4 layer3 unit 100 tag 100 ip 192.168.100.1/24
set network interface ethernet ethernet1/4 layer3 unit 200 tag 200 ip 192.168.200.1/24
set network interface ethernet ethernet1/4 layer3 unit 300 tag 300 ip 192.168.300.1/24
set zone PEN_TEST network ethernet1/4.100
set zone STAGING network ethernet1/4.200
set zone PRODUCTION network ethernet1/4.300

set security rules 'Allow-PEN-TEST-to-STAGING'
set security rules 'Allow-PEN-TEST-to-STAGING' from PEN_TEST to STAGING
set security rules 'Allow-PEN-TEST-to-STAGING' source any destination any application any service any action allow
commit
```

B.

```
set network interface ethernet ethernet1/4 layer3
set network interface ethernet ethernet1/4 layer3 unit 100 tag 100 ip 192.168.100.1/24
set network interface ethernet ethernet1/4 layer3 unit 200 tag 200 ip 192.168.200.1/24
set network interface ethernet ethernet1/4 layer3 unit 300 tag 300 ip 192.168.300.1/24
set zone PEN_TEST network ethernet1/4.100
set zone STAGING network ethernet1/4.200
set zone PRODUCTION network ethernet1/4.300

set security rules 'Allow-PEN-TEST-to-STAGING'
set security rules 'Allow-PEN-TEST-to-STAGING' from PEN_TEST to STAGING
set security rules 'Allow-PEN_TEST-to-STAGING' source any destination any application any service any action allow

set security rules 'Deny-PEN-TEST-to-PRODUCTION'
set security rules 'Deny-PEN-TEST-to-PRODUCTION' from PEN_TEST to PRODUCTION
set security rules 'Deny-PEN-TEST-to-PRODUCTION' source any destination any application any service any action deny

set security rules 'Deny-PRODUCTION-to-PEN-TEST'
set security rules 'Deny-PRODUCTION-to-PEN-TEST' from PRODUCTION to PEN_TEST
set security rules 'Deny-PRODUCTION-to-PEN-TEST' source any destination any application any service any action deny
commit
```

C.

```
set network interface ethernet ethernet1/4 layer3
set network interface ethernet ethernet1/4 layer3 unit 100 tag 100 ip 192.168.100.1/24
set network interface ethernet ethernet1/4 layer3 unit 200 tag 200 ip 192.168.200.1/24
set network interface ethernet ethernet1/4 layer3 unit 300 tag 300 ip 192.168.300.1/24
set zone PEN_TEST network ethernet1/4.100
set zone STAGING network ethernet1/4.200
set zone PRODUCTION network ethernet1/4.300

set security rules 'Allow-PEN-TEST-to-STAGING'
set security rules 'Allow-PEN-TEST-to-STAGING' from PEN_TEST to STAGING
set security rules 'Allow-PEN-TEST-to-STAGING' source any destination any application any service any action allow

set security rules 'Default-Deny-Interzone'
set security rules 'Default-Deny-Interzone' from any to any action deny
commit
```

D.

```
set network interface ethernet ethernet1/4 layer3
set network interface ethernet ethernet1/4 layer3 unit 100 tag 100 ip 192.168.100.1/24
set network interface ethernet ethernet1/4 layer3 unit 200 tag 200 ip 192.168.200.1/24
set network interface ethernet ethernet1/4 layer3 unit 300 tag 300 ip 192.168.300.1/24
set zone PEN_TEST network ethernet1/4.100
set zone STAGING network ethernet1/4.200
set zone PRODUCTION network ethernet1/4.300

set security rules 'Allow-PEN-TEST-to-STAGING'
set security rules 'Allow-PEN-TEST-to-STAGING' from PEN_TEST to STAGING
set security rules 'Allow-PEN-TEST-to-STAGING' source any destination any application any service any action allow

set security rules 'Deny-Cross-Communication'
set security rules 'Deny-Cross-Communication' from PEN_TEST to PRODUCTION
set security rules 'Deny-Cross-Communication' source any destination any application any service any action deny
set security rules 'Deny-Cross-Communication' from PRODUCTION to PEN_TEST
set security rules 'Deny-Cross-Communication' source any destination any application any service any action deny
set security rules 'Deny-Cross-Communication' from STAGING to PRODUCTION
set security rules 'Deny-Cross-Communication' source any destination any application any service any action deny
set security rules 'Deny-Cross-Communication' from PRODUCTION to STAGING
set security rules 'Deny-Cross-Communication' source any destination any application any service any action deny
commit
```

E.

```
set network interface ethernet ethernet1/4 layer3
set network interface ethernet ethernet1/4 layer3 unit 100 tag 100 ip 192.168.100.1/24
set network interface ethernet ethernet1/4 layer3 unit 200 tag 200 ip 192.168.200.1/24
set network interface ethernet ethernet1/4 layer3 unit 300 tag 300 ip 192.168.300.1/24
set zone PEN_TEST network ethernet1/4.100
set zone STAGING network ethernet1/4.200
set zone PRODUCTION network ethernet1/4.300

# Enable Inter-Zone Blocking, which implicitly denies all inter-zone traffic unless explicitly allowed.
set zone PEN_TEST zone-protection profile default
set zone STAGING zone-protection profile default
set zone PRODUCTION zone-protection profile default

set security rules 'Allow-PEN-TEST-to-STAGING'
set security rules 'Allow-PEN-TEST-to-STAGING' from PEN_TEST to STAGING
set security rules 'Allow-PEN-TEST-to-STAGING' source any destination any application any service any action allow
commit
```

Answer: E

Explanation:

The most secure and appropriate way to enforce strict segregation on a Palo Alto Networks firewall is to leverage the implicit deny behavior between zones. By default, traffic between security zones is denied unless explicitly allowed by a security policy. Therefore, the strategy should be to define distinct zones for each network (PEN_TEST, STAGING, PRODUCTION) and then only create an ‘allow' policy for the required communication (PEN_TEST to STAGING). All other inter-zone communication, especially between PEN _ TEST and PRODUCTION, will be implicitly denied. Option E accurately reflects this principle by creating the necessary Layer 3 subinterfaces and associating them with distinct security zones. Then, it correctly sets up the single 'Allow-PEN-TEST-to-STAGlNG' rule. The ‘zone-protection profile default’ command doesn't implicitly deny inter-zone traffic; it refers to flood protection and other zone-based protections. However, the core principle of zone-based security and implicit deny is

the key here. The question implies the need for explicit configuration to achieve this, and correctly defining zones and then selectively allowing traffic fulfills this. Options A, B, C, and D either miss the implicit deny principle or add unnecessary explicit deny rules that are redundant or less efficient than relying on the firewall's default behavior.

Thank You for Trying Our Product

For More Information – Visit link below:

https://www.examsboost.com/

15 USD Discount Coupon Code:

G74JA8UF

FEATURES

- ✓ 90 Days Free Updates
- ✓ Money Back Pass Guarantee
- ✓ Instant Download or Email Attachment
- ✓ 24/7 Live Chat Support
- ✓ PDF file could be used at any Platform
- ✓ 50,000 Happy Customer