ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

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: B, D

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.

Question: 2

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 PANOS 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

Question: 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:
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.

Question: 4

Which statement describes the role of Terraform in deploying Palo Alto Networks NGFWs?

A. It acts as a logging service for NGFW performance metrics.
B. It orchestrates real-time traffic inspection for network segments.
C. It provides Infrastructure-as-Code (IaC) to automate NGFW deployment.
D. It manages threat intelligence data synchronization with NGFWs.

Answer: C

Explanation:
Terraform is an Infrastructure-as-Code (IaC) tool that automates the provisioning and management of infrastructure resources, including Palo Alto Networks Next-Generation Firewalls (NGFWs). By using Terraform configuration files, administrators can define and deploy NGFW instances across cloud environments (such as AWS, Azure, and GCP) efficiently and consistently.

Terraform enables:
Automated firewall deployment in cloud environments.
Configuration of security policies and networking settings in a declarative manner.
Scalability and repeatability, reducing manual intervention in firewall provisioning.

Question: 5

By default, which type of traffic is configured by service route configuration to use the management interface?

A. Security zone
B. IPSec tunnel
C. Virtual system (VSYS)
D. Autonomous Digital Experience Manager (ADEM)

Answer: D

Explanation:
By default, the Autonomous Digital Experience Manager (ADEM) traffic is configured to use the management interface in a Palo Alto Networks firewall. The management interface is typically used for management-related traffic, such as monitoring and logging, and it is configured to handle ADEM-related traffic for the optimal performance of digital experience monitoring features.
This default configuration helps ensure that ADEM traffic does not interfere with regular traffic that may traverse other interfaces, such as traffic from security zones or IPSec tunnels.

Question: 6

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.

Question: 7

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:
When configuring a new security zone on a Palo Alto Networks firewall, the two valid zone types are: Tunnel: A Tunnel zone is used for traffic that is associated with a VPN tunnel, such as IPSec tunnels. Traffic passing through a tunnel interface is classified into this zone.
Virtual Wire: A Virtual Wire zone is used when a firewall operates in transparent mode (also known as Layer 2 mode). In this configuration, the firewall can inspect traffic without modifying the IP address structure of the network.

Question: 8

An organization has configured GlobalProtect in a hybrid authentication model using both certificatebased authentication for the pre-logon stage and SAML-based multi-factor authentication (MFA) for user logon.
How does the GlobalProtect agent process the authentication flow on Windows endpoints?

A. The GlobalProtect agent uses the machine certificate to establish a pre-logon tunnel; upon user sign-in, it prompts for SAML-based MFA credentials, ensuring both device and user identities are validated before granting full access.
B. The GlobalProtect agent uses the machine certificate during pre-logon for initial tunnel establishment, and then seamlessly reuses the same machine certificate for user-based

authentication without requiring MFA.
C. Once the machine certificate is validated at pre-logon, the Windows endpoint completes MFA on behalf of the user by passing existing Windows Credential Provider details to the GlobalProtect gateway without prompting the user.
D. GlobalProtect requires the user to log in first for SAML-based MFA before establishing the prelogon tunnel, rendering the pre-logon certificate authentication (CA) flow redundant.

Answer: A

Explanation:
In a hybrid authentication model with both certificate-based authentication for pre-logon and SAML- based multi-factor authentication (MFA) for user logon, the GlobalProtect agent processes the flow as follows:
During the pre-logon stage, the agent uses the machine certificate to authenticate and establish the initial VPN tunnel.
Once the user logs in (after the machine is connected), the agent then triggers SAML-based MFA to ensure the

user is authenticated with multi-factor authentication, validating both the device and the user identity before granting full access.
This method ensures that both the device and user are properly authenticated and validated in the hybrid authentication model.

Question: 9

An NGFW engineer is configuring multiple Panorama-managed firewalls to start sending all logs to Strata Logging Service. The Strata Logging Service instance has been provisioned, the required device certificates have been installed, and Panorama and the firewalls have been successfully onboarded to Strata Logging Service.
Which configuration task must be performed to start sending the logs to Strata Logging Service and continue forwarding them to the Panorama log collectors as well?

A. Modify all active Log Forwarding profiles to select the "Cloud Logging" option in each profile match list in the appropriate device groups.
B. Enable the "Panorama/Cloud Logging" option in the Logging and Reporting Settings section under Device --> Setup --> Management in the appropriate templates.
C. Select the "Enable Duplicate Logging" option in the Cloud Logging section under Device --> Setup --> Management in the appropriate templates.
D. Select the "Enable Cloud Logging" option in the Cloud Logging section under Device --> Setup --> Management in the appropriate templates.

Answer: D

Explanation:
To begin sending logs to Strata Logging Service while continuing to forward them to Panorama log collectors, the necessary configuration is to enable Cloud Logging. This option is configured in the Cloud Logging section under Device → Setup → Management in the appropriate templates. Once enabled, this ensures that logs are directed both to the Strata Logging Service (cloud) and to the Panorama log collectors.

Question: 10

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.

Question: 11

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.

Question: 12

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.

Question: 13

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.

Question: 14

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.

Question: 15

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.

Question: 16

When configuring a Zone Protection profile, in which section (protection type) would an NGFW engineer configure options to protect against activities such as spoofed IP addresses and split handshake session establishment attempts?

A. Flood Protection
B. Protocol Protection
C. Packet-Based Attack Protection
D. Reconnaissance Protection

Answer: B

Explanation:
In the context of a Zone Protection profile, Protocol Protection is the section used to configure protections against activities such as spoofed IP addresses and split handshake session establishment attempts. These types of attacks typically involve manipulating protocol behaviors, such as IP address spoofing or session hijacking, and are mitigated by the Protocol Protection settings.

Question: 17

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:
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.

Question: 18

Which PAN-OS method of mapping users to IP addresses is the most reliable?

A. Port mapping
B. GlobalProtect
C. Syslog
D. Server monitoring

Answer: D

Explanation:
Server monitoring is the most reliable method for mapping users to IP addresses in PAN-OS. This method allows the firewall to monitor specific servers, such as Microsoft Active Directory (AD) or LDAP servers, to dynamically retrieve and update user-to-IP mappings. It provides a more accurate

and up-to-date mapping of users to their associated IP addresses, as it directly queries user databases in real time.

Question: 19

In an active/active high availability (HA) configuration with two PA-Series firewalls, how do the firewalls use the HA3 interface?

A. To forward packets to the HA peer during session setup and asymmetric traffic flow
B. To exchange hellos, heartbeats, HA state information, and management plane synchronization for routing and User-ID information
C. To synchronize sessions, forwarding tables, IPSec security associations, and ARP tables between firewalls in an HA pair
D. To perform session cache synchronization among all HA peers having the same cluster ID

Answer: D

Explanation:
In an active/active HA configuration with two PA-Series firewalls, the HA3 interface is used primarily for the exchange of HA state information between the firewalls. This includes:
Hellos and heartbeats to monitor the status of the HA peer.

Synchronization of management plane data, which includes critical routing and User-ID information.

Question: 20

A PA-Series firewall with all licensable features is being installed. The customer's Security policy requires that users do not directly access websites. Instead, a security device must create the connection, and there must be authentication back to the Active Directory servers for all sessions. Which action meets the requirements in this scenario?

A. Deploy the transparent proxy with Web Cache Communications Protocol (WCCP).
B. Deploy the Next-Generation Firewalls as normal and install the User-ID agent.
C. Deploy the Advanced URL Filtering license and captive portal.
D. Deploy the explicit proxy with Kerberos authentication scheme.

Answer: D

Explanation:
In this scenario, the customer requires that users do not directly access websites and that a security device (the firewall) manages the connection, while also ensuring that there is authentication back to the Active Directory (AD) servers for all sessions. The explicit proxy with Kerberos authentication is the best solution because:
The explicit proxy allows the firewall to intercept user web traffic and manage the connections on behalf of users.
Kerberos authentication ensures that the user's identity is validated against the Active Directory servers before the session is allowed, fulfilling the authentication requirement.

Question: 21

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

A. User Mapping profile
B. Authentication profile
C. Group mapping settings
D. LDAP Server profile

Answer: C

Explanation:
Before a firewall administrator can define policy rules based on users and groups, the Group Mapping settings must be configured. These settings enable the firewall to map users to their respective Active Directory (AD) groups. This mapping allows the firewall to use user and group information to create policy rules based on group membership.

Question: 22

Which statement applies to the relationship between Panorama-pushed Security policy and local firewall Security policy?

A. When a policy match is found in a local firewall policy, if any Panorama shared post-rule is configured, it will still be evaluated.
B. Local firewall rules are evaluated after Panorama pre-rules and before Panorama post-rules.
C. Panorama post-rules can be configured to be evaluated before local firewall policy for the purpose of troubleshooting.
D. The order of policy evaluation can be configured differently in different device groups.

Answer: B

Explanation:
Local firewall rules are evaluated after Panorama pre-rules (those applied before the firewall's local policies) and before Panorama post-rules (those applied after the firewall's local policies). This ensures that the local firewall rules do not override the central Panorama policy and are only applied in the appropriate order within the policy evaluation sequence.

Question: 23

Which networking technology can be configured on Layer 3 interfaces but not on Layer 2 interfaces?

A. DDNS
B. Link Duplex
C. NetFlow
D. LLDP

Answer: C

Explanation:
NetFlow is a Layer 3 (network layer) protocol that collects and monitors IP traffic flows. It is typically configured on Layer 3 interfaces because it relies on IP information for traffic flow analysis, which is not available on Layer 2 interfaces. Layer 2 interfaces handle frames within the local network, and they don't have IP-related details that NetFlow uses to generate traffic statistics.

Question: 24

What is a result of enabling split tunneling in the GlobalProtect portal configuration with the "Both Network Traffic and DNS" option?

A. It specifies when the secondary DNS server is used for resolution to allow access to specific domains that are not managed by the VPN.
B. It allows users to access internal resources when connected locally and external resources when connected remotely using the same FQDN.
C. It allows devices on a local network to access blocked websites by changing which DNS server resolves certain domain names.
D. It specifies which domains are resolved by the VPN-assigned DNS servers and which domains are resolved by the local DNS servers.

Answer: D

Explanation:
When split tunneling is enabled with the "Both Network Traffic and DNS" option in the GlobalProtect portal configuration, it allows the firewall to control which traffic is sent over the VPN tunnel and which is not. Specifically, it determines which domains are resolved by the VPN-assigned DNS servers (for domains requiring VPN access) and which are resolved by local DNS servers (for domains that can be accessed without the VPN tunnel).

Question: 25

According to dynamic updates best practices, what is the recommended threshold value for content updates in a mission- critical network?

A. 8 hours
B. 16 hours
C. 32 hours
D. 48 hours

Answer: A

Explanation:
For a mission-critical network, it is recommended to configure the content update threshold to 8 hours. This ensures that the network is protected with the latest threat intelligence, updates to signatures, and other critical content, minimizing the exposure to newly discovered vulnerabilities and threats.

Regular content updates are crucial in mission-critical environments to ensure the firewall is up-to- date with the latest protections. 8 hours is considered an optimal balance between timely updates and network performance.

Question: 26

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.

Question: 27

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

A. Panorama, syslog, email
B. Syslog, HTTP, NetFlow
C. Panorama, ADEM, syslog
D. SNMP, HTTP, RADIUS

Answer: A

Explanation:
When configuring the Log Forwarding profile on a Palo Alto Networks firewall, the forwarding methods available include:
Panorama: For forwarding logs to a Panorama management system.
Syslog: For forwarding logs to a syslog server.
Email: For sending logs via email.

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.

Question: 29

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:
Palo Alto Networks Next-Generation Firewalls (NGFWs) use SSL/TLS profiles to secure connections for services such as GlobalProtect Gateways and GlobalProtect Portals. These profiles are used to manage the SSL/TLS encryption and decryption for secure communication between the firewall and clients (such as VPN clients for GlobalProtect). This helps ensure the confidentiality and integrity of the data during transmission.

Question: 30

How does a Palo Alto Networks firewall choose the best route when it receives routes for the same destination from different routing protocols?

A. The route that was received first will be entered into the forwarding table, and all subsequent routes will be rejected.
B. It will attempt to load balance the traffic across all routes.
C. It compares the administrative distance and chooses the one with the highest value.
D. It compares the administrative distance and chooses the one with the lowest value.

Answer: D

Explanation:
When a Palo Alto Networks firewall receives routes for the same destination from different routing protocols, it uses the administrative distance (AD) to determine the best route. The administrative distance is a measure of the trustworthiness of a route, with a lower value indicating higher preference. The firewall will choose the route with the lowest administrative distance to populate its forwarding table.

Question: 31

A large enterprise wants to implement certificate-based authentication for both users and devices, using an on-premises Microsoft Active Directory Certificate Services (AD CS) hierarchy as the primary certificate authority (CA). The enterprise also requires Online Certificate Status Protocol (OCSP) checks to ensure efficient revocation status updates and reduce the overhead on its NGFWs. The environment includes multiple Active Directory forests, Panorama management for several geographically dispersed firewalls, GlobalProtect portals and gateways needing distinct certificate profiles for users and devices, and strict Security policies demanding frequent revocation checks with minimal latency.
Which approach best addresses these requirements while maintaining consistent policy enforcement?

A. Deploy self-signed certificates at each site to simplify local certificate validation and reduce dependencies on a centralized CA. Turn off certificate revocation checks for lower overhead, rely on IP-based rules for GlobalProtect authentication, and use a single certificate profile for both users and devices.
B. Distribute the root and intermediate CA certificates via Panorama as shared objects to ensure all firewalls have a consistent trust chain. Configure OCSP responder profiles on each firewall to ofload revocation checks to an internal OCSP server while keeping CRL checks as a fallback. Maintain separate certificate profiles for user and device authentication and use an automated enrollment method – such as Group Policy or SCEP – to deploy certificates to endpoints.
C. Configure each firewall independently to trust the root and intermediate CA certificates. Rely only on

manual CRL checks for certificate revocation, and import both user and device certificates directly into each firewall's local certificate store for authentication.

D. Obtain wildcard certificates from a public CA for both user and device authentication, and configure firewalls to perform CRL polling at the default update interval. Manually install user certificates on endpoints and synchronize firewall certificate stores through frequent manual SSH updates to maintain consistency.

Answer: B

Explanation:

This approach best addresses the enterprise's requirements for certificate-based authentication, OCSP checks, and consistent policy enforcement:

Distributing the root and intermediate CA certificates via Panorama ensures that all firewalls in the enterprise are consistent in their trust chain and can validate certificates properly.

Configuring OCSP responder profiles on each firewall ofloads the revocation checks to an internal OCSP server, which reduces the overhead on the firewalls and ensures fast, real-time certificate status checks.

Using CRL checks as a fallback ensures reliability in case the OCSP responder is unavailable.

Separate certificate profiles for users and devices ensure that the firewall can enforce different security policies based on the type of certificate (user vs. device).

Automated certificate enrollment methods such as Group Policy or SCEP streamline certificate distribution to endpoints, ensuring efficient management of certificates across geographically dispersed firewalls.

Question: 32

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

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.

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 onpremises and cloud environments, providing visibility and control across all clusters.

Question: 33

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.

Question: 34

An engineer at a managed services provider is updating an application that allows its customers to request firewall changes to also manage SD-WAN. The application will be able to make any approved changes directly to devices via API.
What is a requirement for the application to create SD-WAN interfaces?

A. REST API's "sdwanInterfaceprofiles" parameter on a Panorama device
B. REST API's "sdwanInterfaces" parameter on a firewall device
C. XML API's "sdwanprofiles/interfaces" parameter on a Panorama device
D. XML API's "InterfaceProfiles/sdwan" parameter on a firewall device

Answer: B

Explanation:
To create SD-WAN interfaces through an API, the correct approach is to use the REST API's "sdwanInterfaces" parameter on a firewall device. This parameter allows you to configure SD-WAN interfaces directly on the firewall devices via API, ensuring that the required interfaces are set up and managed for SD-WAN functionality.

Question: 35

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:
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.

Question: 36

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.

Question: 37

Without performing a context switch, which set of operations can be performed that will affect the operation of a connected firewall on the Panorama GUI?

A. Restarting the local firewall, running a packet capture, accessing the firewall CLI
B. Modification of local security rules, modification of a Layer 3 interface, modification of the firewall device hostname
C. Modification of pre-security rules, modification of a virtual router, modification of an IKE Gateway Network Profile
D. Modification of post NAT rules, creation of new views on the local firewall ACC tab, creation of local custom reports

Answer: B

Explanation:
In Panorama, without performing a context switch, the administrator can perform local configuration tasks directly on the connected firewall. The following operations can be done:
Modification of local security rules: Security rules can be modified directly on the connected firewall from the Panorama GUI.
Modification of a Layer 3 interface: Changes to the Layer 3 interfaces on the connected firewall can be done from Panorama, without needing to switch to the firewall's local interface.
Modification of the firewall device hostname: The firewall's hostname can be changed via Panorama.

Question: 38

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.

Question: 39

An administrator plans to upgrade a pair of active/passive firewalls to a new PAN-OS release. The environment is highly sensitive, and downtime must be minimized.

What is the recommended upgrade process for minimal disruption in this high availability (HA) scenario?

A. Suspend the active firewall to trigger a failover to the passive firewall. With traffic now running on the former passive unit, upgrade the suspended (now passive) firewall and confirm proper operation. Then fail traffic back and upgrade the remaining firewall.

B. Shut down the currently active firewall and upgrade it ofline, allowing the passive firewall to handle all traffic. Once the active firewall finishes upgrading, bring it back online and rejoin the HA cluster. Finally, upgrade the passive firewall while the newly upgraded unit remains active.

C. Isolate both firewalls from the production environment and upgrade them in a separate, ofline setup. Reconnect them only after validating the new software version, resuming HA functionality once both units are fully upgraded and tested.

D. Push the new PAN-OS version simultaneously to both firewalls, having them upgrade and reboot in parallel. Rely on automated HA reconvergence to restore normal operations without manually failing over traffic.

Answer: A

Explanation:

In an active/passive HA setup, the recommended process for upgrading involves minimizing downtime and ensuring traffic continuity by using the failover process:

Suspend the active firewall: This triggers a failover to the passive unit, making it the active unit.

Upgrade the former passive (now active) unit: With traffic now running on the previously passive unit, upgrade the suspended unit while the active unit continues handling traffic.

Confirm proper operation: Once the upgrade is complete, verify that the upgraded unit is functioning properly.

Fail traffic back: Once the upgraded firewall is confirmed to be working, fail the traffic back to the original active unit and upgrade the remaining firewall.

Question: 40

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 the context of virtual systems (VSYS) on a Palo Alto Networks firewall, the external zone is typically associated with specific interfaces within a VSYS. Zones are fundamental security objects used to define traffic flow between interfaces, and the external zone would be used for interfaces that connect to external networks.
An external zone is associated with an interface within a VSYS of the firewall. This ensures that traffic from specific interfaces can be classified as belonging to the external zone, allowing the firewall to apply appropriate security policies.
The external zone is indeed a security object that is specific to a given VSYS, as each VSYS can have its own set of zones that are isolated from others.

Question: 41

Which zone type allows traffic between zones in different virtual systems (VSYS), without the traffic leaving the firewall?

A. Isolated
B. Transient
C. External
D. Internal

Answer: B

Explanation:
The Transient zone type is used to allow traffic between zones in different virtual systems (VSYS) on a Palo Alto Networks firewall without the traffic leaving the firewall. It provides a way for virtual systems to communicate with each other by acting as a temporary or intermediary zone. Traffic can pass through the firewall between the virtual systems without requiring physical interfaces or leaving the device.

Question: 42

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 firewalls. Rely on per-firewall Security policies to restrict access to out-of-scope user and group information.
B. Establish separate CIE tenants for each business unit, integrating each tenant with the relevant identity sources. 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.

C. Disable redistribution of identity data entirely. Instead, configure each regional firewall to pull user and group details directly from its local identity providers (IdPs).

D. 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.

Question: 43

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: B, D

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.

Question: 44

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 certificates. Rely solely on CRLs for revocation to minimize complexity.
D. 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.

Question: 45

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.

Question: 46

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.

Question: 47

Which CLI command is used to configure the management interface as a DHCP client?

A. set network dhcp interface management
B. set network dhcp type management-interface
C. set deviceconfig system type dhcp-client
D. set deviceconfig management type dhcp-client

Answer: D

Explanation:
To configure the management interface as a DHCP client on a Palo Alto Networks NGFW, the correct CLI command is set deviceconfig management type dhcp-client.
This command configures the management interface to obtain an IP address dynamically using DHCP.

Question: 48

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.

Question: 49

What are the phases of the Palo Alto Networks AI Runtime Security: Network Intercept solution?

A. Scanning, Isolation, Whitelisting, Logging
B. Discovery, Deployment, Detection, Prevention
C. Policy Generation, Discovery, Enforcement, Logging
D. Profiling, Policy Generation, Enforcement, Reporting

Answer: B

Explanation:
The phases of the Palo Alto Networks AI Runtime Security: Network Intercept solution are designed to help identify and protect against potential threats in real time by using AI to detect and prevent malicious activities within the network.

Discovery: Identifying applications, services, and behaviors within the network to understand baseline activity.

Deployment: Implementing the solution into the network and integrating with existing security measures.

Detection: Monitoring traffic and activities to identify abnormal or malicious behavior.

Prevention: Taking action to stop threats once detected, such as blocking malicious traffic or stopping exploit attempts.

Question: 50

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.