PaloAltoNetworks.SecOps-Generalist.v2026-03-23.q105

Exam Code:	SecOps-Generalist
Exam Name:	Palo Alto Networks Security Operations Generalist
Certification Provider:	Palo Alto Networks
Free Question Number:	105
Version:	v2026-03-23
# of views:	108
# of Questions views:	1054
https://www.freecram.net/torrent/PaloAltoNetworks.SecOps-Generalist.v2026-03-23.q105.html	

NEW QUESTION: 1
Differentiate between the packet processing characteristics of the 'slow path' and the 'fast path' in a Palo Alto Networks security platform (Strata/Prisma Access). Select all statements that accurately describe the distinctions.
A. The slow path is primarily responsible for initial session creation and the application of App-ID and policy lookup, utilizing the device's general-purpose CPU(s).
B. The fast path handles the vast majority of traffic volume for established sessions, relying on hardware acceleration (ASICs or FPGAs) for high throughput.
C. Deep packet inspection for security profiles like Threat Prevention, WildFire submission, and Decryption are exclusively performed in the fast path due to performance requirements.
D. Packets entering the fast path undergo a full security policy re-evaluation and App-ID re-identification on every packet to ensure dynamic policy enforcement.
E. If a session on the fast path encounters a specific condition requiring deeper analysis (e.g., a file upload triggering WildFire analysis or encountering a complex attack signature), subsequent packets for that session or the relevant data stream might be temporarily diverted back to the slow path or a dedicated inspection engine before potentially returning to the fast path.
Answer: (SHOW ANSWER)
Understanding the division of labor between the slow path and fast path is crucial for performance troubleshooting and comprehending how the firewall processes traffic. - Option A (Correct): The slow path (CPU path) is indeed where the initial work of session setup occurs, including identifying the application (App-ID), finding the matching security policy rule, determining security profile assignments, and building the session table entry. - Option B (Correct): The fast path (data plane, leveraging ASICs/hardware acceleration) is optimized for forwarding subsequent packets of established sessions at high speed by performing a quick session table lookup. This offloads the bulk of traffic processing from the CPU. - Option C (Incorrect): While performance optimized, many deep inspection tasks like decryption, full file analysis for WildFire, complex signature matching, and applying specific Data Filtering profiles often involve the slow path CPU or

dedicated content inspection engines which are conceptually part of the deeper processing flow, distinct from the simple fast path session lookup and forwarding. The fast path directs the traffic to these engines based on the session setup in the slow path, but the intensive inspection itself isn't purely ASIC- based forwarding. - Option D (Incorrect): The fast path relies on the session state and policy decision made by the slow path during the first packet processing. Packets on the fast path do not undergo a full policy re-evaluation or App-ID re-identification. They are simply forwarded based on the established session parameters. App-ID is a single-pass inspection and re-classification happens dynamically, but the fast path's role is forwarding based on the current session state. - Option E (Correct): This describes a dynamic switching behavior. Even if a session is primarily on the fast path, specific events (like the start of a file transfer, detecting a pattern requiring deeper analysis, or triggering a vulnerability signature) can cause the relevant packets or streams within that session to be diverted to the slow path CPU or specialized inspection engines for thorough examination before allowing the session to continue on the fast path (if deemed safe) or blocking it.

NEW QUESTION: 2
An administrator is configuring SSL Inbound Inspection for an internal web server hosting at 'www.example.com' on a Strata NGFW. The web server uses a certificate issued by a public Certificate Authority (CA). The administrator has successfully imported the private key for 'www.example.com' into the NGFW's Certificate store. Which steps are necessary in the NGFW's configuration to enable inbound decryption for traffic destined to this server?
A. Configure an SSL Forward Proxy rule in the Decryption Policy matching traffic to 'www.example.com' and reference the imported private key.
B. Create an SSL Inbound Inspection rule in the Decryption Policy matching the destination IP/Zone of the internal web server and reference the imported certificate/private key object.
C. Import the public certificate of the web server's signing CA into the NGFW's Trusted Root CA list.
D. Configure a Decryption Exclusion policy rule for traffic destined to 'www.example.com' to ensure it's not accidentally blocked.
E. Enable the 'Decrypt Mirror' option within the Decryption Profile assigned to the relevant Security policy rule.
Answer: (SHOW ANSWER)
To perform SSL Inbound Inspection for a specific internal server, you need to create a Decryption Policy rule that matches the traffic destined for that server and explicitly configure it for Inbound Inspection, referencing the server's private key (which is associated with the imported certificate object). - Option A: This describes configuring SSL Forward Proxy, which is for outbound traffic, not inbound inspection of internal servers. - Option B (Correct): An SSL Inbound Inspection rule in the Decryption policy is the correct mechanism. This rule matches traffic based on source/destination zones and addresses (the internal server's IP/Zone) and specifies 'Inbound Inspection' as the mode, referencing the imported certificate object that contains the private key needed for decryption. - Option C: Importing the signing CA's public certificate is necessary for

the firewall to validate the server's certificate during the handshake, but it is not sufficient for decrypting the traffic itself; the private key is required for decryption. The private key is imported with the server certificate or separately, and the server certificate object is referenced in the decryption rule. - Option D: This would prevent decryption, which is the opposite of the goal. - Option E: 'Decrypt Mirror' is a troubleshooting feature used to send decrypted traffic to an external tool; it doesn't enable decryption itself.

NEW QUESTION: 3
A security administrator is configuring a Security Policy rule on a Palo Alto Networks PA-Series firewall to allow outbound web browsing for the 'Internal-Users' zone to the 'External' zone. The requirement is to apply comprehensive threat prevention, malware detection, and content filtering to this traffic. Which security profiles, considered Cloud-Delivered Security Services (CDSS) or relying on cloud components for full efficacy, should be attached to this Security Policy rule to meet these requirements? (Select all that apply)
A. Threat Prevention profile
B. Antivirus profile
C. WildFire Analysis profile
D. URL Filtering profile
E. File Blocking profile
Answer: (SHOW ANSWER)
Cloud-Delivered Security Services (CDSS) are subscriptions that enhance the security efficacy of Palo Alto Networks platforms by leveraging cloud-based intelligence and analysis. The profiles listed are the key Content-ID security profiles used for deep inspection, many of which heavily rely on cloud lookups and analysis for their full effectiveness: - Option A (Correct): Threat Prevention uses cloud-delivered threat intelligence for IPS and Antispyware. - Option B (Correct): Antivirus uses cloud-delivered malware signatures for real-time scanning. - Option C (Correct): WildFire Analysis submits unknown files to the cloud sandbox for dynamic analysis and verdict determination. - Option D (Correct): URL Filtering queries the cloud-based URL database for categorization and threat intelligence (malicious URLs). - Option E (Correct): File Blocking enforces policy on file types detected via deep inspection, often working in conjunction with Antivirus and WildFire. While some profiles also have on-box components, their full, dynamic, and global intelligence comes from the cloud services. All of these profiles are standard Content-ID security profiles applied to Security Policy rules for comprehensive inspection.

NEW QUESTION: 4
Using the 'No Decrypt' action for specific traffic flows in Palo Alto Networks Strata NGFW or Prisma Access Decryption policy has significant implications for security visibility. When a session matches a 'No Decrypt' rule, which of the following security features or inspection capabilities are typically unavailable or severely limited for that specific encrypted session? (Select all that apply)
A. App-ID identification of the application within the encrypted tunnel.
B. Scanning the file content transferred within the session for malware using WildFire or Antivirus.

C. Applying Threat Prevention signatures (Vulnerability Protection, Antispyware) to detect exploits or command-and-control traffic hidden within the encrypted payload.
D. Enforcing URL Filtering based on the full requested URL path, beyond just the hostname presented in the Server Name Indication (SNI) field.
E. Blocking sessions based on the source or destination IP address matching a high-risk external dynamic list (EDL).
Answer: B,C,D (LEAVE A REPLY)
The purpose of decryption is to gain visibility into the encrypted payload to apply deeper security inspection. When 'No Decrypt' is used, that deeper inspection is lost. - Option A (Incorrect): App-ID can often identify applications even within encrypted traffic by examining the initial handshake (like SNI for HTTPS) and behavioral heuristics, although its accuracy may be reduced compared to decrypted traffic. - Option B (Correct): WildFire and Antivirus scan the file content . If the session is not decrypted, the firewall cannot see or extract the file content to scan it for malware. - Option C (Correct): Threat Prevention signatures operate on the payload data to detect patterns indicative of exploits or malicious communication. Without decryption, the payload remains encrypted and cannot be inspected by these engines. - Option D (Correct): URL Filtering can partially work on encrypted traffic by using the hostname from the SNI field (or the certificate's Common Name if SNI is not used). However, it cannot see the full URL path requested after the connection is established (e.g., '[sensitive_data/upload.php'). Full URL path filtering requires decryption. - Option E (Incorrect): Blocking based on source/destination IP address using EDLs is a network-layer enforcement that occurs regardless of whether the session is encrypted or decrypted. The IP is visible in the packet headers.

NEW QUESTION: 5
Prisma SD-WAN allows administrators to define policies for different categories of applications, such as 'Voice & Video', 'Critical Business Apps', 'Bulk Transfer', and 'Default'. Which type of policy is used to define how traffic matching these application categories should be prioritized, managed, and steered across the available WAN links?
A. Security Policy
B. Qos Policy
C. NAT Policy
D. Path Policy
E. Application Identification Policy
Answer: (SHOW ANSWER)
Prisma SD-WAN's Path Policy (sometimes also referred to as Business Intent Overlay or similar concepts in SD-WAN) is where the application categories are mapped to specific forwarding behaviors and link preferences. You define rules saying 'for Voice & Video traffic, prefer paths with low jitter', 'for Bulk Transfer, use paths with high bandwidth', etc. Option A controls allow/deny/inspect. Option B prioritizes traffic on a link. Option C handles address translation. Option E is part of App-ID, which identifies the application, but doesn't define the pathing behavior.

NEW QUESTION: 6
When configuring a Remote Network in Prisma Access for a branch office, you must specify the local branch subnets that will be sent through the IPSec tunnel to Prisma Access. Why is it important to accurately define these branch-local subnets in the Remote Network configuration?
A. It determines which public IP address range Prisma Access will use to Source NAT outbound internet traffic from the branch.
B. It allows Prisma Access to correctly route traffic from other Prisma Access locations (Mobile Users, other Remote Networks) to the defined branch subnets via the established tunnel.
C. It is used by App-ID to identify applications originating from that branch.
D. It dictates which security profiles (Threat Prevention, URL Filtering) are applied to traffic originating from that branch.
E. It enables Decryption policy for all encrypted traffic originating from those subnets.
Answer: (SHOW ANSWER)
Defining local branch subnets in the Remote Network configuration primarily serves to advertise those subnets into the Prisma Access routing domain. - Option A: Source NAT configuration for internet traffic is typically done in NAT policies, and the public IP used depends on the Prisma Access location and configuration, not the local branch subnets themselves (though the NAT rule matches on those subnets). - Option B (Correct): By defining the local branch subnets, you are essentially telling Prisma Access, "These subnets are behind this Remote Network tunnel." This allows Prisma Access to build its routing table and know that if traffic arrives from a Mobile User or another Remote Network and is destined for an IP within one of those branch subnets, it should be routed down the IPSec tunnel to that specific branch. This is essential for inter-branch and remote user to branch communication. - Option C: App-ID identifies applications based on the traffic stream itself, not based on the source subnet definition in the network configuration. - Option D: Security profiles are applied based on Security Policy rules, which match traffic based on criteria like Source/Destination Zones, User, Application, etc., not directly based on the subnet definition in the Remote Network object (though the zone assigned to the Remote Network is used). - Option E: Decryption policy is configured separately based on matching criteria and actions, not simply by defining subnets in the Remote Network object.

NEW QUESTION: 7
An administrator has configured SSL Forward Proxy decryption for outbound internet traffic on a Palo Alto Networks NGFW They want to exclude a specific application internal-app') running on HTTPS (port 443) from decryption because it uses client-side certificates. The 'internal-app' is hosted externally but accessed by internal users. There is a general 'Decrypt all outbound HTTPS' rule lower in the policy. Which configuration steps are necessary to create the exclusion rule?
A. Create a Decryption policy rule with Action 'No Decrypt', Source Zone 'internal', Destination Zone 'external', Application 'internal-app', and place this rule above the 'Decrypt all outbound HTTPS' rule.

B. Create a Security policy rule with Action 'No Decrypt', Source Zone 'internal', Destination Zone 'external', Application 'internal-app', and place this rule above the 'Decrypt all outbound HTTPS' rule.
C. Edit the 'Decrypt all outbound HTTPS' rule and add the 'internal-app' to its exclusion list within the rule options.
D. Create a custom URL Category for the 'internal-app' domain and add this URL Category to the Decryption Profile used by the 'Decrypt all outbound HTTPS' rule.
E. Remove the 'SSI' service from the 'Decrypt all outbound HTTPS' rule and create a separate rule for 'internal-app' with no decryption.
Answer: (SHOW ANSWER)
Exclusions in Decryption policy are achieved using 'No Decrypt' rules placed strategically. - Option A (Correct): This is the correct method. You create a separate rule in the Decryption Policy that specifically matches the traffic you want to exclude (based on source/destination zones, the specific application, etc.) and set the action to 'No Decrypt'. Placing this rule above the broader 'Decrypt' rule ensures that this specific traffic is evaluated and exempted from decryption before the general decryption rule is encountered. - Option B: 'No Decrypt' is a Decryption Policy action, not a Security Policy action. - Option C: While some policies allow specific exclusions within a rule, the standard and more flexible method for defining broad exceptions based on multiple criteria is through separate 'No Decrypt' rules. - Option D: Decryption Profiles handle error actions and unsupported parameters, not lists of URLs to exclude from decryption policy matching itself. - Option E: Removing 'SSI' from the decrypt rule would prevent decryption for all HTTPS traffic, not just the specific application. Using separate rules for applications is valid in Security Policy but the exclusion itself is configured in the Decryption Policy.

NEW QUESTION: 8
In a Palo Alto Networks NGFW with Advanced DNS Security enabled, where would an administrator configure the policy to specify the action the firewall should take (e.g., sinkhole, block, alert) when a DNS query is classified as malicious by the cloud service?
A. In the Security Policy rule matching the DNS traffic, by selecting a specific action like 'deny'.
B. Within the DNS Security Profile that is attached to the Security Policy rule matching the DNS traffic.
C. In the Decryption Policy rule for DNS traffic.
D. In the WildFire Analysis profile.
E. In the URL Filtering profile for the 'malware' category.
Answer: (SHOW ANSWER)
Actions for detected malicious DNS queries are configured within the DNS Security Profile, which is then applied to Security Policy rules. - Option A: The Security Policy rule defines the overall action for the session (e.g., 'allow' DNS traffic). The specific action upon detection of a malicious query within that allowed traffic is defined in the security profile. - Option B (Correct): The DNS Security Profile is where you configure how the firewall responds to different classifications provided by the Advanced DNS Security cloud service (e.g., 'malware', 'phishing', 'command-

and-control'). You define actions like 'Sinkhole', 'Block', 'Alert', etc., based on these categories. This profile is then attached to the Security Policy rule that permits DNS traffic (UDP/53 or TCP/53). - Option C: Decryption policy is for encrypted traffic, not standard DNS. - Option D: WildFire Analysis profiles are for file analysis. - Option E: URL Filtering profiles are for web access based on URLs, not DNS queries.

NEW QUESTION: 9
When managing a firewall using Panorama, what is the primary benefit of configuring device-specific settings (like interface IP addresses, hostname, system services) within a Template or Template Stack in Panorama, rather than configuring them locally on the firewall itself?
A. Templates allow these settings to be dynamically updated without requiring a commit on the firewall.
B. Templates enable High Availability synchronization of these settings between firewalls.
C. Templates ensure consistency for settings that are common across multiple firewalls (e.g., NTP server, DNS server) while still allowing unique values where needed, and prevent local overrides from disrupting centralized management.
D. Templates allow these settings to be inherited from Shared Policy objects.
E. Configuring in templates significantly reduces the firewall's resource utilization.
Answer: (SHOW ANSWER)
Templates are designed for managing device-specific and network configurations centrally. - Option A: Changes pushed from templates still require a commit on the firewall to take effect. - Option B: HA synchronization handles state and config sync between HA pairs, but templates provide the source of the configuration definition. - Option C (Correct): Templates allow administrators to standardize common device- level settings across many firewalls (e.g., all branch firewalls use the same NTP servers, management profiles). They also allow for variable values (e.g., unique interface IPs per firewall) and enforce that these settings are centrally managed, preventing administrators from making disruptive local overrides outside of Panorama control. This ensures consistency and control. - Option D: Templates manage device/network settings; Shared Policy objects manage security policy and policy objects. - Option E: Centralized management doesn't directly reduce the firewall's operational resource utilization; it simplifies configuration.

NEW QUESTION: 10
A. Threat logs
B. System logs
C. Configuration logs
D. Traffic logs
E. HIP Match logs
Answer: (SHOW ANSWER)
Traffic logs are the fundamental logs generated by the firewall that provide details about every session that hits a policy rule. They include critical information like source/destination IP and

zones, application ID, user ID (if User-ID is enabled), action (allow, deny, drop, reset), bytes transferred, and session duration. This makes them the primary source for analyzing traffic patterns, policy hits, and user activity. Option A focuses on detected threats. Option B tracks system events. Option C logs configuration changes. Option E logs device posture compliance.

NEW QUESTION: 11
An organization is deploying Palo Alto Networks VM-Series firewalls within a public cloud VPC (e.g., AWS, Azure) to secure application tiers. They require High Availability for these firewalls. While Active/Passive HA is supported, they are considering an Active/Active setup using external cloud provider load balancers or routing mechanisms for distributing traffic. Which of the following statements accurately describe aspects or implications of implementing VM-Series HA in public cloud environments, particularly when considering Active/Active configurations? (Select all that apply)
A. Active/Passive HA for VM-Series typically relies on gratuitous ARP and MAC address updates for failover, similar to physical appliances.
B. Implementing Active/Active HA for VM-Series in public cloud often requires external cloud infrastructure (like load balancers or policy-based routing) to distribute incoming sessions across the active firewall instances.
C. Session state synchronization between VM-Series firewalls in an Active/Active configuration is necessary to prevent session disruption if a firewall instance handling a flow fails.
D. Cloud NGFW for AWS/Azure provides native cloud-managed HA, abstracting the underlying HA mechanisms from the user.
E. VM-Series Active/Active HA requires dedicated HA links configured with static IP addresses for control plane and data plane synchronization between the instances.
Answer: (SHOW ANSWER)
HA in virtualized and cloud environments has specific considerations: - Option A (Incorrect): Public cloud networks often restrict or don't support Gratuitous ARP or direct MAC address manipulation for HA failover. VM-Series HA in the cloud typically relies on cloud-specific mechanisms like API calls to update route tables or IP addresses, or external load balancers. - Option B (Correct): Active/Active HA on VM-Series requires an external mechanism (like an AWS Network Load Balancer or Azure Standard Load Balancer, or routing manipulation) to direct incoming traffic to both active firewall instances, distributing the load. - Option C (Correct): In Active/Active HA, multiple firewalls are processing traffic simultaneously. To ensure session continuity if one active instance fails, the session state must be synchronized between the instances. Otherwise, traffic arriving at the remaining active instance for a session previously handled by the failed instance would be seen as a new session, potentially causing disruption. - Option D (Correct): Cloud NGFW for AWS/Azure is a managed service. The cloud provider and Palo Alto Networks handle the underlying HA and scaling mechanisms (often multi-AZ) transparently to the user, who simply consumes the firewall service. - Option E (Incorrect): While physical PA-Series use dedicated HA links, VM-Series in cloud environments typically use

standard virtual network interfaces for HA synchronization traffic, often within a dedicated management or HA subnet/VLAN.

NEW QUESTION: 12
Which Palo Alto Networks Cloud-Delivered Security Services (CDSS) require a firewall to send metadata or copies of suspicious content to a cloud-based analysis or intelligence platform to perform their primary security function? (Select all that apply)
A. App-ID
B. Threat Prevention (specifically threat intelligence feeds)
C. WildFire analysis
D. URL Filtering (specifically URL category lookups)
E. User-ID
Answer: (SHOW ANSWER)
CDSS leverage the cloud for scale, intelligence, and dynamic analysis: - Option A (Incorrect): App-ID identification primarily occurs on the firewall itself using signatures, heuristics, and protocol decoding. While App-ID definitions are updated from the cloud, the core identification process is local. - Option B (Correct): Threat Prevention signatures and dynamic threat intelligence feeds are delivered from the cloud. While enforcement happens on the firewall, the intelligence comes from the cloud service. - Option C (Correct): WildFire's core function is dynamic analysis in a cloud sandbox. Suspicious files and/or session details are sent from the firewall to the WildFire cloud for analysis. - Option D (Correct): URL Filtering relies on a massive, dynamic cloud-based database of URLs and their categories/threat status. The firewall queries this cloud service for real-time lookups. - Option E (Incorrect): User-ID identifies users by mapping IP addresses to usernames, typically by integrating with local or cloud-based identity sources (like AD, LDAP, Okta, etc.) but doesn't involve sending traffic content or metadata to a separate CDSS for the identification itself.

NEW QUESTION: 13
A company is deploying Prisma Access to provide secure internet access and access to internal resources for its branch offices. Each branch office has a router or firewall capable of establishing an IPSec VPN tunnel. Which component of Prisma Access is specifically designed to receive these IPSec VPN connections from branch office locations and provide access to the Prisma Access security capabilities and service connections?
A. Mobile Users Security Processing Nodes
B. Service Connections
C. Remote Networks Security Processing Nodes
D. Cloud Management Console
E. Cortex Data Lake
Answer: C (LEAVE A REPLY)
Prisma Access uses different components to handle different types of connections. Remote Networks are for site-to-site connections (branch offices, headquarters, campuses) using IPSec

tunnels. - Option A: Mobile Users Security Processing Nodes handle connections from individual remote users using GlobalProtect. - Option B: Service Connections represent the tunnels from Prisma Access back to your internal data centers or cloud VPCsNNets. - Option C (Correct): Remote Networks Security Processing Nodes are the dedicated cloud-hosted components of Prisma Access that terminate IPSec tunnels from branch offices and other sites defined as Remote Networks. - Option D: The Cloud Management Console is the management interface. - Option E: Cortex Data Lake is the logging service.

NEW QUESTION: 14
A company uses Prisma Access for mobile users and Remote Networks, with subscriptions for Advanced Threat Prevention, Advanced URL Filtering, WildFire, and Enterprise DLP They need to create a security policy that: - Allows marketing users to access sanctioned social media (e.g., corporate LinkedIn pages) but blocks all other social networking. - Blocks any attempt to download malware (known or unknown). - Prevents the upload of sensitive customer data to any public cloud storage. - Blocks access to known malicious websites (phishing, malware hosting) and C2 domains. Which combination of Security Policy rule elements, CDSS-enabled profiles, and decryption configuration are necessary to achieve these goals? (Select all that apply)
A. Security Policy rule(s) matching source user ('Marketing' group), source zone ('Mobile-Users'/'Remote-Networks'), destination zone ('Public'), with application control for sanctioned/unsanctioned social media App-IDs and specific URL categories.
B. Security Policy rule(s) with WildFire Analysis, Antivirus, and Threat Prevention profiles applied to all traffic allowed to the 'Public' zone to block malware and exploits.
C. Security Policy rule(s) with Data Filtering profile applied, configured to detect sensitive customer data patterns (e.g., PII), matching upload activities (App Functions) to cloud storage applications, and set to a 'block' action.
D. SSL Forward Proxy decryption policy enabled for HTTPS traffic destined for social media, cloud storage, and general internet browsing to allow inspection by App-ID, Content-ID, and Data Filtering.
E. Security Policy rule(s) with Advanced URL Filtering and Advanced DNS Security profiles applied to block access to malicious websites and C2 domains.
Answer: (SHOW ANSWER)
This scenario requires combining multiple CDSS and policy types for comprehensive protection. - Option A (Correct): Security policy rules based on user identity, zones, application App-IDs, and URL categories are needed to allow sanctioned social media and block unsanctioned ones. - Option B (Correct): WildFire, Antivirus, and Threat Prevention profiles (all enhanced by CDSS) are applied to the allow rules to scan for malware and exploits in the allowed traffic. - Option C (Correct): Data Filtering profiles (enhanced by Enterprise DLP CDSS) are configured to detect sensitive data and applied to policy rules that match upload traffic to cloud storage, with a block action for unsanctioned destinations. - Option D (Correct): Decryption is mandatory to inspect encrypted traffic (HTTPS), which is commonly used by social media, cloud storage, and malicious sites/C2, to enable App-ID, Content-ID, and Data Filtering on the actual content. - Option E

(Correct): Advanced URL Filtering and Advanced DNS Security profiles are applied to Security Policy rules (typically outbound to the Public zone) to block access based on malicious URLs and C2 domains at the web and DNS layers, respectively. All these elements work together to provide multi-layered security for various traffic types and threats.

NEW QUESTION: 15

A network administrator is configuring outbound internet access for an internal subnet (192.168.20.0/24) on a Palo Alto Networks Strata NGFW. They are using Dynamic IP and Port (DIPP) Source NAT (SNAT) to translate internal IPs to a single public IP (203.0.113.10) on the firewall's internet-facing interface. The NAT policy rule is configured as follows:

```
NAT Rule Name: Outbound_SNAT Rule Type: ipv4 Original Packet:   Source Zone: Internal   Destination Zone: External
Source Address: 192.168.20.0/24   Destination Address: any   Service: any Translated Packet:   Source Address Translation:
Dynamic IP And Port   Address Type: Interface Address   Interface: ethernet1/1 (Public Interface)
```

After this NAT rule is successfully matched, the firewall proceeds to evaluate Security Policy rules. When creating the Security Policy rule to allow this outbound internet traffic, what combination of Source Address and Destination Address should MOST logically be used in the Security Policy rule to match the traffic flow after the NAT rule is applied and determined?

A. Source Address: 192.168.20.0/24, Destination Address: any

B. Source Address: 203.0.113.10, Destination Address: any

C. Source Address: any, Destination Address: any

D. Source Address: 192.168.20.0/24, Destination Address: 203.0.113.10

E. Source Address: any, Destination Address: 203.0.113.10

Answer: (SHOW ANSWER)

In Palo Alto Networks firewalls, NAT policy is evaluated before Security Policy. However, the Security Policy rule matches based on the packet headers as they are after the NAT rule's effect is determined , or for certain criteria (like Source/Destination Zone, Application, User), based on the properties of the session before the actual packet modification occurs, but after the NAT rule has been identified. The standard practice is that the Security Policy uses the original Source IP and the post-NAT determined Destination IP (which is the original external destination for outbound traffic). Therefore, the Security Policy rule allowing this traffic should match the original source subnet (192.168.20.0/24) and the destination which is 'any' external address. The firewall knows, based on the matched NAT rule, that this traffic originating from 192.168.20.0/24 and going to the External zone will be translated. The Security Policy then evaluates based on the original source and the intended destination. Option B is incorrect because the Security Policy rule's source address field matches the original source IP, not the translated one. Options C, D, and E do not correctly represent the source address based on the scenario's requirements.

NEW QUESTION: 16

A security team wants to harden their network by preventing users from downloading potentially dangerous file types from the internet (e.g., executable files, archive files, batch scripts) while still allowing safe documents like PDFs. They also want to prevent the upload of encrypted or password-protected archive files (like ' -zip' or .rar') to external services, as these cannot be

inspected for malware or sensitive dat a. Which Content-ID feature is specifically used to implement these restrictions based on file type and direction?
A. WildFire analysis profile configured to block unknown file types.
B. Data Filtering profile configured to detect file extensions in the data stream.
C. URL Filtering profile configured to block websites known to host malicious file types.
D. File Blocking profile configured with rules specifying file types and transfer directions (upload/download) to block or alert on.
E. Threat Prevention profile with custom vulnerability signatures matching dangerous file headers.
Answer: (SHOW ANSWER)
The File Blocking profile is the Content-ID component specifically designed to control the transfer of files based on their type and the direction of the transfer (upload or download). Option D accurately describes this functionality. It allows administrators to create granular rules, for instance, blocking .exe' downloads, blocking .zip' uploads (especially if encrypted and thus not inspectable), but allowing .pdf downloads. Option A submits files for analysis but doesn't block based on type. Option B uses data patterns, not file types. Option C blocks sites but not the file types themselves if downloaded from an allowed site. Option E uses signatures for vulnerabilities, not file type control.

NEW QUESTION: 17
A user at a branch office is experiencing poor quality during a video conference call via Zoom. The Prisma SD-WAN ION device at the branch has multiple WAN links. The administrator wants to troubleshoot this specific issue by examining how the Zoom traffic is being treated by the SD-WAN. Which of the following log types or monitoring views within the Prisma SD-WAN Cloud Management Console would provide the MOST relevant information for diagnosing the path and quality issues for this specific call? (Select all that apply)
A. Application Performance Monitoring (APM) data for the 'zoom' application, showing its end-to-end performance metrics over the SD-WAN paths.
B. SD-WAN Flow logs filtered for the user's IP and the destination IP/port of the Zoom call, showing which specific WAN link(s) the traffic traversed and the quality metrics on those links at the time.
C. Path Quality monitoring data showing the real-time and historical latency, jitter, and packet loss for all WAN links at the branch.

D. Threat logs to see if any security events were detected on the Zoom traffic.
E. Traffic logs filtered for the user's IP and the Zoom application, showing the policy rule matched and the action (allow).
Answer: (SHOW ANSWER)
Diagnosing application performance issues over SD-WAN requires focusing on application-specific metrics, flow details, and underlying link quality. - Option A (Correct): APM provides direct insight into the user experience for specific applications, showing performance over the SD-WAN fabric. - Option B (Correct): SD-WAN Flow logs are crucial for seeing the specific path a given application flow (the user's Zoom call) took and the measured quality on that path. This helps determine if the steering policy was applied correctly and if the chosen path had poor quality. - Option C (Correct): Path Quality monitoring provides the overall health of the links. If APM or Flow logs show poor quality on a path, examining the general Path Quality for that link helps understand if it was an isolated incident or a persistent link problem. - Option D: Threat logs are for security detections, not performance issues. - Option E: Traffic logs show policy matches and actions but typically don't include the detailed SD-WAN path selection or performance metrics relevant to quality issues.

NEW QUESTION: 18
An administrator is using the Best Practice Assessment (BPA) feature in AIOps for NGFW to evaluate their firewalls. The BPA generates a score and lists specific findings across various categories. Which category of findings is the BPA PRIMARILY designed to identify?
A. Real-time traffic anomalies and detected threat events.
B. Deviations from Palo Alto Networks recommended security and operational configuration settings.
C. Hardware failures and physical interface status issues.
D. User authentication failures and identity mapping issues.
E. Outdated software versions that are not supported.
Answer: (SHOW ANSWER)
The Best Practice Assessment (BPA) is a tool to evaluate a firewall's configuration against a set of recommended best practices developed by Palo Alto Networks. It checks for deviations from these best practices across various configuration areas (policy, network, device, objects, etc.). Option A describes real-time monitoring and threat detection logs. Option C relates to system health monitoring. Option D relates to User-ID monitoring. Option E relates to system or update status.

NEW QUESTION: 19
Which types of content can typically be submitted to Palo Alto Networks WildFire cloud service for analysis by a Strata NGFW or Prisma Access? (Select all that apply)
A. Executable files (e.g., '.exe', '.dll')
B. Document files (e.g., '.pdf, '.doc', '.xlS, '.ppt)
C. Archive files (e.g., 'zip', '.rar')

D. Scripts (e.g., '.js', '.vbS, '.psl')
E. Encrypted files (e.g., password-protected zips, encrypted documents)
Answer: (SHOW ANSWER)
WildFire supports analysis of a wide variety of file types that are commonly used to deliver malware. - Option A (Correct): Executables and libraries are prime targets for malware. - Option B (Correct): Documents can contain malicious macros or embedded exploits. - Option C (Correct): Archives are often used to package and hide malware; WildFire can unpack many common archive formats for analysis. - Option D (Correct): Scripts are frequently used for malicious purposes (downloaders, execution, reconnaissance). - Option E (Incorrect): WildFire cannot analyze content it cannot decrypt. Password-protected archives or encrypted documents cannot be analyzed in the sandbox unless the password/key is somehow made available or brute-forced (which is not a standard function of WildFire). Such files are often blocked by File Blocking policies precisely because they cannot be inspected.

NEW QUESTION: 20
A. WildFire analysis profile configured to block unknown file types.
B. Data Filtering profile configured to detect file extensions in the data stream.
C. URL Filtering profile configured to block websites known to host malicious file types.
D. File Blocking profile configured with rules specifying file types and transfer directions (upload/download) to block or alert on.
E. Threat Prevention profile with custom vulnerability signatures matching dangerous file headers.
Answer: (SHOW ANSWER)
The File Blocking profile is the Content-ID component specifically designed to control the transfer of files based on their type and the direction of the transfer (upload or download). Option D accurately describes this functionality. It allows administrators to create granular rules, for instance, blocking .exe' downloads, blocking .zip' uploads (especially if encrypted and thus not inspectable), but allowing .pdf downloads. Option A submits files for analysis but doesn't block based on type. Option B uses data patterns, not file types. Option C blocks sites but not the file types themselves if downloaded from an allowed site. Option E uses signatures for vulnerabilities, not file type control.

NEW QUESTION: 21
Which action types are typically available for configuration within the Vulnerability Protection profile on a Palo Alto Networks NGFW to respond to detected exploit attempts? (Select all that apply)
A. Allow
B. Alert
C. Block
D. Reset Server (for server-side exploits)
E. Quarantine the source endpoint
Answer: (SHOW ANSWER)

Vulnerability Protection profile actions define how the firewall responds when an exploit signature is matched. - Option A (Incorrect): 'Allow' is not a typical action for detected exploit attempts; the goal is to prevent the exploitation. - Option B (Correct): 'Alert' generates a log entry and notification without preventing the traffic. Useful for monitoring or testing. - Option C (Correct): 'Block' terminates the session and drops the malicious packets, preventing the exploit from reaching the target. This is a common preventative action. - Option D (Correct): 'Reset Server' (or 'Reset Client', 'Reset Both') injects TCP reset packets into the stream to cleanly terminate the connection. This can be useful for preventing server processes from entering an unstable state after an attempted exploit. - Option E (Incorrect): While quarantining endpoints is a response capability often integrated via platforms like Cortex XDR or network access control (NAC), it is not a direct action within the Vulnerability Protection profile itself on the NGFW.

NEW QUESTION: 22
A security team manages a large fleet of Palo Alto Networks firewalls using Panoram a. They have enabled AIOps for NGFW to improve operational efficiency and security posture. They receive an AIOps alert about high session setup rates on a specific firewall, potentially indicating a performance bottleneck or a network anomaly (like a connection flood). Which of the following are valid actions the team can take or insights they can gain by leveraging the integration between AIOps and Panorama/Cortex Data Lake to investigate and address this alert? (Select all that apply)
A. View historical trends and analyze the rate of new sessions on the affected firewall over time within the AIOps dashboard to determine if the current rate is an anomaly or a consistent pattern.
B. Drill down from the AIOps alert into the detailed Traffic logs for the affected firewall (stored in Cortex Data Lake/Panorama Log Collector) to identify the source IPs, destinations, and applications contributing to the high session setup rate.
C. Receive recommendations from AIOps on potential causes for the high session setup rate, such as short-lived connections or specific application traffic patterns.
D. Automatically apply QOS policies via AIOps to mitigate the impact of high session setup on critical traffic.
E. Identify if the high session setup rate correlates with any specific configuration changes made to the firewall using AIOps' change correlation capabilities.
Answer: (SHOW ANSWER)
AIOps for NGFW analyzes operational data and provides insights, recommendations, and correlation. - Option A (Correct): AIOps tracks key operational metrics like session rates and provides historical trend analysis, allowing administrators to differentiate between temporary spikes and persistent issues. - Option B (Correct): A crucial aspect is integration with logging. AIOps provides context-aware links or drilling capabilities into the relevant logs (in CDL or Panorama) to investigate the details of the events triggering the alert, such as identifying the source/destination of the high session rate traffic. - Option C (Correct): AIOps uses machine learning and analysis to identify potential root causes or contributing factors to observed operational issues, providing actionable recommendations (e.g., optimize policy for short-lived

connections, investigate specific applications). - Option D (Incorrect): While AIOps might recommend applying QOS, it does not automatically implement configuration changes like applying policies. Implementation is done manually via Panorama or the firewall UI. - Option E (Correct): AIOps can correlate operational anomalies or performance changes with recent configuration commits, helping administrators identify if a recent change might be the cause of the issue.

NEW QUESTION: 23
When a remote user connecting via GlobalProtect accesses the public internet through Prisma Access, which security policy flow is evaluated?
A. From the user's local interface zone to the internet destination zone.
B. From the 'Mobile-UserS zone to the 'Public' zone.
C. From the 'Remote-Networks' zone to the 'Public' zone.
D. From the Public' zone to the 'Mobile-Users' zone.
E. From the 'Service-Connection' zone to the 'Public' zone.
Answer: (SHOW ANSWER)
Prisma Access defines specific zones for mobile users and the public internet. - Option A: The user's local interface zone is not relevant to the traffic flow once it's in the Prisma Access cloud. - Option B (Correct): Traffic from mobile users connecting via GlobalProtect originates from the 'Mobile-Users' zone within Prisma Access and is destined for the public internet, represented by the 'Public' zone. Security Policy rules for outbound internet access for mobile users are written from the 'Mobile-Users' zone to the "Public' zone. - Option C: 'Remote- NetworkS zone is for site-to-site VPNs. - Option D: This is the flow for traffic originating from the internet destined for mobile users (less common for standard browsing, more for specific services). - Option E: 'service-Connectiori zone represents internal resources, not the source of mobile user traffic.

NEW QUESTION: 24
In a Prisma SD-WAN deployment using ION devices, an administrator notices that traffic between two internal subnets assigned to the same Security Zone is not appearing in the traffic logs, even though a logging profile is attached to the relevant Security Policy rules. Traffic between these subnets is successfully flowing. What is the MOST likely reason the traffic logs are missing for this intra-zone communication?
A. The Security Policy rule matching this traffic has logging disabled.
B. Intra-zone traffic is implicitly allowed by the 'intra-zone-default' rule and bypasses explicit Security Policy rule evaluation, therefore it is not logged by default security policy logging.
C. The interfaces connected to these subnets are configured in Tap mode instead of Layer 3 mode.
D. User-ID is not enabled on the interfaces, preventing logging of user sessions.
E. A NAT policy rule is incorrectly translating the source or destination IPs, preventing logging.
Answer: (SHOW ANSWER)

This question focuses on the behavior of default zone rules and logging. - Option A: If an explicit rule were matched, a disabled logging profile would prevent logs, but the core issue is whether an explicit rule is matched at all. - Option B (Correct): Traffic between interfaces assigned to the same zone is permitted by the 'intra-zone-default' rule. Crucially, traffic matched by default rules (both intra-zone-default allow and inter-zone-default deny) does not hit the explicit security policy rules table for evaluation or logging unless an explicit policy rule is specifically configured to override the default behavior for intra-zone traffic. Therefore, the traffic is allowed, but doesn't trigger logging associated with explicit policy rules. - Option C: Tap mode is for monitoring, not inline forwarding, and would prevent the traffic from flowing as described. - Option D: While User-ID provides username context in logs, its absence doesn't prevent logging of session details based on IPlapplication/policy match if the traffic hits a logging-enabled rule. - Option E: An incorrect NAT rule might break connectivity, but it wouldn't typically prevent logging if a session was established and matched a logging-enabled security rule.

NEW QUESTION: 25
A network operations team relies on AIOps for NGFW to proactively identify potential performance issues before they impact users. They observe an AIOps alert indicating a high rate of packet drops on a specific interface of a PA-Series firewall. Which specific data points or views available through the AIOps dashboard or its linked components (like Cortex Data Lake) would be MOST helpful in diagnosing the potential root cause of these packet drops? (Select all that apply)
A. Interface statistics showing input/output errors and drop counters on the affected interface over time, visualized in AIOps.
B. System resource utilization (CPU, memory, data plane/management plane load) graphs for the affected firewall at the time of the packet drops.
C. Traffic logs filtered for the affected interface showing the type of traffic and policy action associated with the dropped packets (requires drill-down to CDL/Panorama logs).
D. Performance monitoring metrics related to session setup rate and throughput on the firewall.
E. Configuration history to see if recent changes were made to the affected interface or related policies.
Answer: (SHOW ANSWER)
Diagnosing packet drops requires examining network interface metrics, system resources, traffic logs, performance indicators, and recent changes. AIOps aggregates many of these or links to the source data. - Option A (Correct): Direct interface statistics are crucial for confirming packet drops and potentially identifying the nature of the errors (e.g., input drops due to overload, output errors). AIOps collects and visualizes these. - Option B (Correct): High CPU or data plane load can cause packet drops due to the firewall being overwhelmed. Checking resource utilization is a standard diagnostic step available via AIOps. - Option C (Correct): Traffic logs (in CDL/Panorama) provide details about why traffic is dropped (e.g., denied by policy, hit a specific error). Filtering logs by the affected interface helps correlate drops with specific traffic types or policy enforcement. AIOps facilitates drilling down to these logs. - Option D (Correct): High session setup rate or maximum throughput being reached can indirectly lead to packet drops on

interfaces as the firewall struggles to process traffic. Performance monitoring metrics provide this context. - Option E (Correct): Recent configuration changes (e.g., interface speed/duplex mismatch, new policies causing unexpected load) can cause packet drops. AIOps change correlation helps identify such potential causes.

NEW QUESTION: 26
You are using Panorama to monitor a large number of managed firewalls. You want to create a custom report that shows the top applications consuming the most bandwidth across all managed devices, broken down by Security Zone and User Group. Which log type in Panorama's Monitor tab is the primary source for building this type of report?
A. Threat logs
B. Summary logs
C. Traffic logs
D. URL Filtering logs
E. System logs
Answer: (SHOW ANSWER)
Reports on application usage, bandwidth consumption, user activity, and traffic patterns are built from the detailed session information found in Traffic logs. - Option A: Threat logs are for detected security events. - Option B: Summary logs provide aggregated statistics, but detailed reports broken down by specific criteria like Zone, User Group, and individual Application are best built from the raw session data in Traffic logs. - Option C (Correct): Traffic logs contain the bytes transferred per session, the application ID, the source user/group, and the source/destination zones. This detailed data allows you to aggregate and filter to create reports showing top applications by bandwidth, segmented by user and zone. - Option D: URL Filtering logs focus on web access and categories, not overall application bandwidth for all applications. - Option E: System logs monitor firewall health.

NEW QUESTION: 27
In a PAN-OS SD-WAN deployment, how does the firewall primarily leverage App-ID information when making real-time path selection decisions for application traffic?
A. App-ID is used to encrypt traffic before it is sent over the selected WAN link.
B. App-ID identifies the application, and the Path Selection policy uses this application identity as a matching criterion to apply specific routing rules or performance requirements.
C. App-ID dynamically changes the port and protocol of the application to match the capabilities of the best available WAN link.
D. App-ID directs traffic to the management plane for detailed processing and path selection.
E. App-ID is only used for security policy enforcement (allow/deny), not for path selection.
Answer: (SHOW ANSWER)
App-ID is fundamental to the application-aware capabilities of PAN-OS SD-WAN. - Option A: Encryption is typically handled by IPSec or SSL/TLS tunnels, not App-ID. - Option B (Correct): App-ID identifies the specific application (e.g., Zoom, SQL, SharePoint). The Path Selection

policy rules use these App-IDs as matching criteria. This allows you to define rules like "For Zoom traffic (App-ID 'zoom'), use the path with the lowest jitter," or "For SQL traffic (App-ID 'ms-sql'), use the MPLS path if latency is below 50ms." The real-time path quality metrics are then applied to the links based on the application's needs as defined in the policy. - Option C: App-ID identifies the application; it doesn't modify its port or protocol. - Option D: App-ID identification occurs on the data plane during initial session setup, not the management plane for routing decisions. - Option E: App-ID is used for both security policy (allow/deny/inspect) and for intelligent path selection in SD-WAN.

NEW QUESTION: 28
An organization uses Panorama to manage a large number of distributed PA-Series firewalls. They need to enforce a consistent security policy across groups of similar firewalls (e.g., all branch office firewalls should have the same basic internet access policy). They also need to configure device-specific settings like interface IPs and zones on each firewall. Which two primary concepts within Panorama are used to achieve this separation of shared policy/objects and device-specific configurations?
A. Security Policies and NAT Policies
B. Device Groups and Templates
C. Virtual Systems and Security Zones
D. Shared Policy and Device-Specific Policy
E. Log Collectors and Management Servers
Answer: (SHOW ANSWER)
Panorama uses specific constructs for hierarchical configuration management. - Option A: These are types of policies, but not the containers for shared vs. device-specific settings. - Option B (Correct): Device Groups are used to manage shared security policies and objects that apply to all firewalls within the group. Templates are used to manage shared network and device-specific configurations (interfaces, zones, system settings). Firewalls are assigned to both a Device Group and a Template Stack (a collection of Templates evaluated in order) to receive their full configuration. - Option C: Virtual Systems segment a single firewall into multiple virtual firewalls; Security Zones define trust boundaries on the firewall. These are device-level concepts, not Panorama management constructs for shared vs. unique config. - Option D: While Panorama has shared policy, Device-Specific Policy is applied within the Device Group, and Templates handle the non-policy device config. - Option E: These are components for logging and management, not configuration management hierarchy.

NEW QUESTION: 29
An administrator is using the Palo Alto Networks IoT Security subscription with their NGFW. They need to identify and inventory all previously unknown devices communicating on the internal network, visualize their communication patterns, and assess their security risk posture. Which dashboard or reporting view within the IoT Security portal (or integrated management platform) is designed to provide this comprehensive visibility into the discovered IoT device landscape?

A. Threat logs viewer
B. Traffic logs viewer
C. Device Inventory / Risk Dashboard
D. System logs viewer
E. URL Filtering logs viewer
Answer: (SHOW ANSWER)
The IoT Security solution provides dedicated dashboards for visualizing the discovered device inventory and their associated risks. Option A, B, D, and E are generic log viewers for security events, traffic flows, system events, and web access, respectively. The Device Inventory or Risk Dashboard specifically aggregates information about profiled devices, their types, vulnerabilities, communication patterns, and overall risk score.

NEW QUESTION: 30
A. Granular control based on user identity (e.g., allow Finance users to access Finance app) and application identity (e.g., allow only specific collaboration tools), independent of IP addresses or ports.
B. Consistent policy enforcement for users regardless of their changing IP address (e.g., when moving between locations or getting a new DHCP lease).
C. Improved performance by allowing the firewall to bypass deep packet inspection for trusted users and applications.
D. Reduced administrative overhead by eliminating the need for security zones or NAT policies.
E. Enhanced security posture by allowing policies to be defined based on 'who is doing what', rather than just 'where the traffic is going'.
Answer: (SHOW ANSWER)
User-ID and App-ID are core enablers of next-generation firewall capabilities, moving beyond traditional Layer 3/4 controls. - Option A (Correct): This is a primary advantage. Policy can be tied directly to user groups and specific applications (identified by App-ID), providing much more granular control than IP/port/zone alone. You can say 'Marketing users can use Salesforce, but not Dropbox', regardless of the IPs involved. - Option B (Correct): User-ID maps dynamic IP addresses to static user identities. This ensures that a policy applied to 'jdoe' follows jdoe regardless of which IP address they are currently using (obtained via DHCP at home, a public hotspot, etc.), which is essential for remote users. - Option C (Incorrect): While optimization might occur, the purpose of User-ID and App-ID is to enable more accurate and relevant inspection, not to bypass it. In a Zero Trust model, inspection is applied even to trusted users/apps based on policy. - Option D (Incorrect): User- ID and App-ID enhance security policy rules but do not eliminate the need for zones (which define trust boundaries) or NAT policies (for address translation). They provide additional criteria within the policy framework. - Option E (Correct): This summarizes the shift in security posture. By incorporating User-ID ('who') and App-ID ('what') alongside traditional IP/zone ('where'), policies become more aligned with actual user activities and risks, moving closer to a Zero Trust model based on identity and application.

NEW QUESTION: 31
A large healthcare organization is implementing Palo Alto Networks firewalls for perimeter security. Due to strict regulatory and privacy requirements (like HIPAA in the US, GDPR in Europe), they need to ensure that sensitive patient data transmitted via encrypted channels to approved healthcare providers or cloud services is NOT subjected to SSL Forward Proxy decryption, even though general web browsing is decrypted and inspected. What is the appropriate Decryption Policy action and placement for traffic involving this sensitive data?
A. Configure an SSL Forward Proxy rule with the 'Decrypt' action for the sensitive traffic, but apply a specific Decryption Profile that is configured to bypass inspection.
B. Create a 'No Decrypt' rule in the Decryption Policy matching the sensitive traffic criteria (e.g., source users/groups, destination URL category for healthcare providers) and place this rule above any 'Decrypt' rules that would otherwise match the traffic.
C. Remove HTTPS from the allowed services in the Security Policy rules for sensitive traffic destinations.
D. Configure an SSL Inbound Inspection rule for the sensitive traffic, requiring the server's private key.
E. Apply a URL Filtering profile configured to 'allow' the sensitive destinations within the Security Policy.
Answer: (SHOW ANSWER)
When specific traffic must not be decrypted due to privacy, legal, or technical reasons, the 'No Decrypt' action in the Decryption Policy is used. Option B correctly describes this: a specific rule is created to match the criteria of the sensitive traffic, assigned the 'No Decrypt' action, and crucially, placed above any broader 'Decrypt' rules that might also match this traffic. The firewall processes Decryption policy rules top- down, similar to Security policy. Option A is incorrect; applying 'Decrypt' and then attempting to bypass with a profile is not the standard or explicit way to prevent decryption based on policy matching. Option C is incorrect; removing HTTPS would block the traffic entirely, which is not the goal. Option D is for inspecting inbound traffic to internal servers, not outbound sensitive data transfers. Option E controls access based on URL categories but does not prevent or manage decryption.

NEW QUESTION: 32

An organization is deploying GlobalProtect. They want to implement certificate-based authentication for the GlobalProtect clients to the Gateway, in addition to username/password or multi-factor authentication. This provides an extra layer of trust based on the client device identity Which configuration steps are necessary on the Palo Alto Networks NGFW or Prisma Access Gateway and potentially on the client side to enable this? (Select all that apply)

A. Import a Client CA certificate onto the GlobalProtect Gateway and configure an Authentication Profile to use certificate authentication, referencing this C

B. Issue and deploy unique Client Certificates to each GlobalProtect endpoint that will authenticate via certificate.

C. Configure the GlobalProtect Agent settings to use certificate-based authentication when connecting to the Gateway.

D. Ensure the certificate presented by the Gateway is signed by a CA trusted by the client device.

E. Enable SSL Inbound Inspection on the GlobalProtect Gateway interface.

Answer: (SHOW ANSWER)

Implementing client certificate authentication requires configuration on both the gateway and the client, involving trusted CAS and certificate distribution. - Option A (Correct): The Gateway needs to trust the CA that issued the client certificates. Importing the Client CA (the root or intermediate CA that signed the client certificates) and configuring an Authentication Profile to use certificate authentication referencing this CA enables the gateway to validate client certificates. - Option B (Correct): Each endpoint that will authenticate using a certificate must have a unique client certificate installed and available. - Option C (Correct): The GlobalProtect Agent configuration on the endpoint must be set up to present the client certificate during the authentication process when connecting to the configured gateway. - Option D (Correct): While this option repeats a concept from the previous question, it's relevant here. The client needs to trust the gateway's server certificate for the tunnel to be established securely in the first place, regardless of whether the client is also presenting its own certificate. - Option E (Incorrect): SSL Inbound Inspection is for decrypting incoming traffic destined for internal servers, not for authenticating GlobalProtect clients to the gateway.

NEW QUESTION: 33

An organization needs to implement granular security policies based on user identity and application usage for remote users connecting via Prisma Access. They are leveraging User-ID with SAML integration for authentication and App-ID for application visibility. Which of the following statements accurately describe how User-ID and App-ID work together in this scenario to enable policy enforcement?

(Select all that apply)

A. User-ID maps the remote user's assigned IP address (from the Prisma Access pool) to their username and associated groups, which are then available as matching criteria in Security Policy rules.

B. App-ID identifies the specific application (e.g., 'slack', 'salesforce', 'web-browsing') being used within the remote user's session, independent of the destination port.

C. Security Policy rules combine User-ID information (source user/group) and App-ID information (application) with traditional network criteria (source/destination zone, destination address) to define granular access controls.
D. App-ID identification must occur before User-ID mapping is possible for a given session.
E. Decryption is always required for App-ID to identify applications like HTTPS-based SaaS traffic.
Answer: (SHOW ANSWER)
User-ID and App-ID are complementary technologies for user- and application-aware security. - Option A (Correct): User-ID integrates with identity sources (like SAML providers via CIE or GlobalProtect agent) to obtain the username associated with the IP address that the remote user is assigned by Prisma Access. This mapping is then used in policy. - Option B (Correct): App-ID identifies the application by examining traffic characteristics, protocol decoding, and behavioral analysis, independent of the static port, providing the 'what' of the session. - Option C (Correct): Security Policy rules are the point where User-ID (who), App-ID (what), and traditional Layer 3/4/zone information (where) are combined to create highly specific rules like "Allow Marketing users access to Salesforce App when going from Mobile-Users zone to Public zone." - Option D (Incorrect): App-ID identification and User-ID mapping are often parallel processes during session setup. User-ID maps the source IP to a user; App-ID identifies the application based on the flow characteristics. Neither strictly requires the other to complete first, although both are needed for policies that combine them. - Option E (Incorrect): While decryption significantly enhances App-ID accuracy, especially for distinguishing different applications on the same encrypted port (like various SaaS apps on 443), App-ID can often identify applications using methods like SNI inspection, certificate common names, and behavioral analysis even without full decryption.

NEW QUESTION: 34
Palo Alto Networks performs software updates and maintenance on the underlying Prisma Access infrastructure periodically. Which of the following statements accurately describe how these updates and maintenance activities are designed to affect the availability and security posture of the Prisma Access service for customers? (Select all that apply)
A. Updates are performed on a per-customer basis, requiring manual scheduling by the administrator.
B. Updates are typically performed in a rolling, non-disruptive manner across the global infrastructure to minimize impact on user connectivity and session state.
C. Customers are notified in advance of scheduled maintenance windows for Prisma Access updates.
D. During updates, security inspection capabilities (App-ID, Threat Prevention) are temporarily disabled to ensure connectivity.
E. The administrator is responsible for downloading and installing the new Prisma Access software version via the Cloud Management Console.
Answer: (SHOW ANSWER)

As a cloud service, the vendor (Palo Alto Networks) manages the underlying infrastructure maintenance and updates for Prisma Access, designed for high availability. - Option A: Updates are managed globally by Palo Alto Networks, not scheduled manually by individual customers. - Option B (Correct): Palo Alto Networks employs rolling update strategies across the global infrastructure, updating nodes in clusters or regions sequentially to minimize disruption. The goal is typically non-disruptive updates where existing sessions are maintained or seamlessly failed over. - Option C (Correct): While non-disruptive is the goal, Palo Alto Networks provides advance notification to customers about scheduled maintenance windows and update activities via standard communication channels. - Option Option D (Incorrect): The goal of the updates is to maintain or improve security posture, not disable security inspection during the process. Updates are designed to keep security services active. - Option E: As with dynamic updates, the administrator does not manage the installation of the underlying Prisma Access software itself; this is handled by Palo Alto Networks.

NEW QUESTION: 35
A global company is implementing granular control over SaaS application usage using Palo Alto Networks Strata NGFWs at branch offices and Prisma Access for remote users. They have configured decryption policies to inspect SSL/TLS traffic for sanctioned SaaS applications like Office 365 and Salesforce. However, users accessing unsanctioned shadow IT applications via encrypted channels are still successfully bypassing security controls. Additionally, some legitimate applications are experiencing functionality issues after decryption is enabled. What are potential reasons for these issues and necessary steps to address them?
A. Decryption is not properly configured for all relevant traffic zones, causing some encrypted traffic to pass through uninspected.
B. The applications identified by App-ID are not all being processed by the decryption policy before reaching security profiles.
C. The firewall/Prisma Access might be encountering SSL/TLS protocol versions or cipher suites that are not supported for decryption, leading to decryption failures and fallback to non-decrypted paths (potentially allowing unsanctioned apps).
D. Application functionality issues may arise if the application uses client-side certificates, pinned certificates, or relies on specific SSL/TLS negotiation steps that are disrupted by the decryption proxy.
E. The security policy rules using App-ID are ordered incorrectly, allowing 'allow' rules for 'any' application to match encrypted traffic before the decryption policy is evaluated.
Answer: (SHOW ANSWER)
This scenario highlights common challenges with decrypting encrypted traffic for application layer inspection. Option A is correct because decryption policies must apply to the correct zones and traffic flows; misconfiguration can cause traffic to bypass decryption. Option B is incorrect; App-ID identifies the application regardless of whether it's decrypted or not, although granular enforcement after identification often requires decryption for full Content-ID, Threat Prevention, etc. Option C is correct; the firewall/Prisma Access has limitations on supported SSL/TLS

versions, cipher suites, and key exchange methods. If an application uses unsupported parameters, decryption will fail, and depending on the decryption profile's action for 'decryption errors', the session might be allowed without inspection. Option D is correct; applications using mechanisms like certificate pinning or client authentication can break when a decryption proxy intercepts and re-signs the certificate. Exclusions for such applications are often necessary. Option E is incorrect; Security policy rule evaluation happens after App-ID identification and typically after decryption policy evaluation (if decryption is enabled for the matched rule's traffic). Rule order primarily affects which policy is applied to the identified application, not whether decryption happens or fails beforehand.

NEW QUESTION: 36
A remote user connected to Prisma Access via GlobalProtect attempts to access both a public SaaS application (e.g., Salesforce) and a private application hosted in the corporate data center. Both applications are accessed over HTTPS. How does Prisma Access facilitate and secure access to these two distinct types of applications for the remote user?
A. Public SaaS traffic is routed directly to the internet via the user's local connection, bypassing Prisma Access security inspection.
B. Both public SaaS and private application traffic are tunneled to Prisma Access. Public SaaS traffic is then inspected and routed to the internet via the nearest cloud service edge, while private application traffic is routed through a 'Service Connection' tunnel to the corporate data center for access.
C. Both public SaaS and private application traffic are routed to the corporate data center first, where they are inspected by an on-premises firewall before being forwarded to their destinations.
D. Prisma Access only secures access to private applications; public internet access requires a separate security solution.
E. Access to both application types is determined solely by the user's device posture, as evaluated by GlobalProtect HIP.
Answer: (SHOW ANSWER)
Prisma Access is designed to secure access to both public and private applications for remote users, leveraging its cloud-native architecture. - Option A (Incorrect): A primary goal of Prisma Access for mobile users is to tunnel all relevant traffic through the service for consistent security inspection, including internet-bound traffic to public SaaS. - Option B (Correct): This accurately describes the Prisma Access flow. Traffic destined for the public internet (including SaaS) is sent through the GlobalProtect tunnel to the nearest Prisma Access cloud service edge, inspected by the cloud-based NGFW features, and then routed securely to the internet. Traffic destined for private corporate resources is also sent through the tunnel, but Prisma Access identifies it as private traffic and routes it through the configured 'Service Connection' (an IPSec or GRE tunnel) to the corporate data center or cloud VPC hosting the private application. - Option C (Incorrect): Hairpinning all traffic back to the data center negates the benefits of a cloud-delivered security platform and can introduce latency. Prisma Access routes internet-bound traffic locally from the cloud edge. - Option D (Incorrect): Prisma Access provides comprehensive security for both

public and private application access. - Option E (Incorrect): Device posture (HIP) is a factor in allowing the user to connect and potentially applying policy, but it doesn't determine the routing path taken for public vs. private applications; that's based on destination IP address and Prisma Access routing configuration.

NEW QUESTION: 37
A company is onboarding its remote workforce onto Prisma Access. Users will connect from various locations globally. To secure user traffic and enforce corporate security policies, user endpoints will connect to Prisma Access. Which Palo Alto Networks endpoint software component is typically deployed on users' laptops and mobile devices to establish a secure connection to Prisma Access and provide user and device posture information?
A. Cortex XDR agent
B. Traps endpoint software (legacy name)
C. GlobalProtect agent
D. Xpanse Explorer
E. VM-Series appliance
Answer: (SHOW ANSWER)
GlobalProtect is Palo Alto Networks' secure network access client used by remote users to connect to firewalls (PA-Series, VM-Series, and Prisma Access). It establishes a secure tunnel and can collect user information (User-ID) and device posture (HIP). Option A (Cortex XDR) is for endpoint detection and response, not specifically for network access. Option B is a legacy name for the endpoint protection component, now part of Cortex XDR. Option D (Xpanse Explorer) is for external attack surface management. Option E is a virtual firewall appliance, not endpoint software.

NEW QUESTION: 38
In a Palo Alto Networks Strata NGFW or Prisma Access deployment, configuring interfaces and zones is a prerequisite for policy enforcement. When assigning multiple interfaces (e.g., VLAN subinterfaces, physical Ethernet ports) to a single Security Zone, what are the key implications for traffic flow and security policy application?
A. Traffic between any two interfaces assigned to the same zone is implicitly allowed by the 'intra-zone-default' security rule, bypassing explicit security policy rule evaluation.
B. Traffic between any two interfaces assigned to the same zone is implicitly denied by the 'inter-zone-default' security rule unless explicitly allowed by a policy rule.
C. Explicit security policy rules with the Source Zone and Destination Zone set to the same zone name are required to permit any traffic flow between interfaces within that zone.
D. Security policies cannot be written using zones when multiple interfaces are assigned to the same zone; policies must use interface objects instead.
E. Assigning multiple interfaces to the same zone complicates App-ID identification and reduces the effectiveness of Content-ID inspection for traffic flowing between those interfaces.
Answer: (SHOW ANSWER)

Understanding the default zone behavior is critical. Palo Alto Networks firewalls have built-in default rules: - Intra-zone-default: Allows traffic between interfaces assigned to the same security zone. - Inter-zone-default: Denies traffic between interfaces assigned to different security zones. When multiple interfaces are assigned to a single zone, traffic traversing the firewall between these interfaces is considered 'intra-zone' traffic. Option A correctly states that this traffic is implicitly allowed by the intra-zone-default rule and bypasses explicit security policy evaluation. Option B describes the 'inter-zone-default' rule, which applies between different zones. Option C is incorrect; explicit rules are for inter-zone traffic or overriding the default behavior. Option D is incorrect; policies are written using zones, regardless of how many interfaces are in a zone. Option E is incorrect; the number of interfaces in a zone doesn't inherently complicate App-ID or Content-ID; those functions apply to traffic flows regardless of the specific interface, as long as the policy is matched and decryption (if needed) is performed.

NEW QUESTION: 39
Which of the following statements accurately describes the relationship between Cloud-Delivered Security Services (CDSS) and Security Profiles on Palo Alto Networks NGFWs and Prisma SASE?
A. CDSS are entirely separate cloud services that operate independently of the security profiles configured on the firewall/Prisma Access.
B. Security Profiles are configuration objects on the firewall/Prisma Access where administrators define the desired security actions, and these profiles leverage the intelligence and capabilities provided by the CDSS subscriptions.
C. CDSS subscriptions automatically apply security actions globally without requiring Security Policy or profile configuration.
D. Security Profiles are only used for basic Layer 4 filtering (port/protocol), while CDSS provide advanced inspection.
E. CDSS are physical or virtual appliances deployed alongside the firewall to perform security inspection.
Answer: (SHOW ANSWER)
CDSS subscriptions enhance the efficacy of the security profiles configured on the firewall or Prisma SASE. - Option A: CDSS are cloud services, but they are integrated with and leveraged by the firewall's security profiles. - Option B (Correct): Security Profiles (Threat, URL, WildFire Analysis, etc.) are where the administrator defines the policy (e.g., 'block high-severity threats', 'alert on gambling sites'). These profiles, when subscribed to the relevant CDSS, gain access to the latest threat intelligence, cloud-based analysis engines (WildFire), and dynamic databases (URL Filtering, DNS Security) provided by the CDSS. The firewall enforces the policy defined in the profile using the intelligence from the cloud. - Option C: CDSS provide intelligence and capabilities, but policy actions (allow, block, alert) are defined by the administrator in Security Profiles and applied via Security Policy rules. - Option D: Security Profiles contain configurations for advanced Layer 7 inspection engines (App-ID, Content-ID), not just basic Layer 4 filtering. - Option E: CDSS are cloud-delivered services , not physical or virtual appliances deployed by the

customer (the exception being some on-premises components like WF-500 appliances for specific use cases, but the service itself is cloud-based).

NEW QUESTION: 40

A. Security policies, NAT policies, Decryption policies, and network configuration (interfaces, zones, routing) are configured separately within each virtual system.
B. Shared policy objects (like Address Groups or Security Profiles) created in one virtual system can be directly referenced by policy rules in another virtual system.
C. The default inter-zone-default rule is applied and enforced independently within each virtual system.
D. Traffic flowing between interfaces assigned to different virtual systems is implicitly allowed by default.
E. Panorama can manage multiple virtual systems on a single physical firewall, allowing for centralized policy and object management across vsys.
Answer: (SHOW ANSWER)
Virtual systems provide logical isolation of firewall functions. - Option A (Correct): A primary purpose of vsys is to provide configuration separation. Each vsys has its own distinct set of Security, NAT, and Decryption policies, as well as its own network configuration (interfaces, zones, routing tables). - Option B (Incorrect): Configuration is isolated between vsys. Objects defined within one vsys cannot be directly referenced by policies in another vsys. Shared objects must be defined at the vsys level where they are used or inherited from a Panorama template/device group if managed centrally. - Option C (Correct): Each vsys functions as an independent firewall instance. The default intra-zone-default allow and inter-zone-default deny rules are applied and enforced independently within the context of each vsys's zones. - Option D (Incorrect): Traffic flowing between interfaces assigned to different virtual systems is implicitly denied by default, just like traffic between different zones within a vsys. Explicit inter-vsys policy must be configured in a dedicated inter-vsys zone (if configured) or via a separate firewall/routing if not directly connected. - Option E (Correct): Panorama can manage multiple virtual systems on a single physical or virtual firewall. It allows defining shared policies and objects at higher levels that can be inherited by specific vsys, or managing each vsys as a distinct device group.

NEW QUESTION: 41

An administrator is using AIOps for NGFW to monitor the health, security posture, and performance of their Palo Alto Networks firewalls. They receive an alert from AIOps indicating a potential configuration best practice violation regarding an outdated security zone configuration. Which of the following actions can the administrator typically perform directly within or leverage through the AIOps for NGFW platform to address such a finding?
A. Automatically remediate the configuration violation with a single click from the AIOps dashboard.
B. View detailed information about the specific best practice rule that was violated and the recommended corrective steps.

C. Generate a report summarizing all identified best practice violations across all monitored firewalls.
D. Initiate a configuration commit on the affected firewall directly from the AIOps interface after making changes.
E. Perform real-time packet captures on the affected firewall triggered by the AIOps alert.
Answer: B,C (LEAVE A REPLY)
AIOps for NGFW is primarily a proactive monitoring, analysis, and recommendation engine. While it integrates with management platforms, its core function is providing insights and guidance. - Option A (Incorrect): AIOps for NGFW does not currently support one-click automatic remediation of configuration changes directly from the dashboard. It provides recommendations that the administrator must implement via Panorama or the firewall I-JI. - Option B (Correct): A core function is providing detailed context for findings, including explanations of the best practice rule violated and specific, actionable recommendations for correction. - Option C (Correct): AIOps allows administrators to generate reports on various findings, including configuration best practices, performance bottlenecks, and security risks, across the managed firewall estate. - Option D (Incorrect): Configuration commits are performed on Panorama or the individual firewall, not directly initiated from the AIOps interface. - Option E (Incorrect): Real-time packet captures are troubleshooting tools performed directly on the firewall CLI or UI, not initiated by AIOps alerts.

NEW QUESTION: 42
A company is using Prisma Access for Mobile Users and Remote Networks. They want to apply different levels of security inspection based on the source of the traffic. Traffic from corporate-owned laptops connecting via GlobalProtect should receive full decryption and deep content inspection, while traffic from less-trusted Remote Networks (e.g., guest Wi-Fi at branches) should receive basic threat prevention and URL filtering but may not be fully decrypted. How are Security Profiles and Decryption Policies typically used in conjunction with Security Policy rules in Prisma Access to achieve this tiered security approach? (Select all that apply)
A. Configure separate Security Policy rules for each source type (Mobile Users, Remote Networks), matching the respective source zones.
B. Create different Security Profile Groups, one with comprehensive profiles (Threat, AV, WildFire, URL, File, Data) and another with a subset of profiles (Basic Threat, Basic URL).
C. Create Decryption Policy rules that match the source zone (Mobile Users) and specify the 'Decrypt' action for relevant traffic (like HTTPS), placing them higher than rules for other sources.
D. Apply the comprehensive Security Profile Group to the Security Policy rules matching Mobile IJser traffic.
E. Apply the less comprehensive Security Profile Group to the Security Policy rules matching Remote Network traffic and ensure relevant Decryption Policy rules (e.g., 'No Decrypt' or specific exclusions) are configured for those zones.
Answer: (SHOW ANSWER)
Implementing tiered security in Prisma Access involves segmenting traffic sources by zone, defining different security profiles, and controlling decryption. - Option A (Correct): Policy

evaluation starts by matching traffic to a Security Policy rule. Creating rules based on source zones (Mobile-Users, 'Remote-Networks) is the way to apply different policies to traffic from different origins. - Option B (Correct): Security profiles define the specific inspection settings. Creating different bundles of profiles allows you to apply varying levels of inspection. - Option C (Correct): Decryption is necessary for deep inspection. Decryption Policy rules determine if traffic is decrypted. Rules matching the 'Mobile- Users' zone with a 'Decrypt' action enable full inspection for corporate users. Rules for less trusted zones might specify 'No Decrypt' for certain traffic or have a 'Decrypt' rule placed lower or with more exceptions. - Option D (Correct): Once the Security Policy rule matches the Mobile User traffic (identified by Source Zone 'Mobile-Users'), applying the comprehensive Security Profile Group enforces the desired deep inspection. - Option E (Correct): Similarly, applying the less comprehensive Security Profile Group to the rules matching Remote Network traffic enforces a lower level of inspection. Ensuring Decryption Policies are aligned (e.g., fewer things decrypted, more bypasses, or 'No Decrypt' rules) is necessary because full deep inspection (like Data Filtering or WildFire analysis) requires decryption.

NEW QUESTION: 43
A financial institution is implementing a Palo Alto Networks Strata NGFW to secure its internal network and prevent data exfiltration and malware infections over encrypted channels. They need to inspect all outbound HTTPS traffic from employee workstations to detect sensitive data leaving the network and block access to malicious websites identified via URL filtering and Threat Prevention, even if accessed over SSL/TLS. Which decryption method is required for this use case, and what is its fundamental principle of operation?
A. SSL Forward Proxy decryption, which intercepts the SSL/TLS handshake, presents the client with a certificate signed by the firewall's root CA, and establishes separate encrypted sessions with the client and the server.
B. SSL Inbound Inspection, which requires installing the server's private key on the firewall to decrypt incoming encrypted connections to internal servers.
C. SSL Inbound Inspection with a wildcard certificate, which allows the firewall to decrypt any incoming encrypted connection without needing individual server private keys.
D. SSL Protocol Downgrade, which forces the client and server to use an unencrypted version of the protocol (e.g., HTTP instead of HTTPS).
E. Proxy Automatic Configuration (PAC) file decryption, which redirects encrypted traffic to a transparent proxy for inspection before sending it to the destination.
Answer: (SHOW ANSWER)
The scenario describes the need to inspect outbound encrypted traffic from internal clients (workstations) to external destinations (malicious websites, cloud services for data exfiltration). This is the primary use case for SSL Fomard Proxy decryption. Option A correctly describes the process: the firewall acts as a 'man-in-the-middle' by intercepting the connection attempt, generating a certificate for the requested website on the fly (signed by a root CA trusted by the clients), establishing an encrypted session with the client, and a separate encrypted session with

the actual server. This allows the firewall to see and inspect the unencrypted traffic between these two sessions. Option B describes SSL Inbound Inspection, used for securing traffic to internal servers. Option C is incorrect as wildcard certificates are used for inbound inspection, not outbound forward proxy. Option D is not a standard, secure, or effective decryption method employed by modern firewalls for this purpose; it would break legitimate traffic and is insecure. Option E describes a method for directing traffic, but not the mechanism for performing the SSL/TLS decryption itself, which still relies on a proxy or firewall capability like SSL Forward Proxy.

NEW QUESTION: 44
A company is using Palo Alto Networks Prisma Access for its remote workforce and relies on the Cloud Management Console and Cortex Data Lake (CDL) for monitoring and logging. A security incident involves a remote user potentially downloading a malicious file through a sanctioned SaaS application. Which logging components are involved in capturing the relevant security event data for this incident, and where would an administrator typically view the detailed logs?
A. Logs are generated on the user's endpoint and stored locally for analysis.
B. Prisma Access service edge generates traffic, threat, and other logs and forwards them to Cortex Data Lake.
C. Logs are sent directly from the Prisma Access service edge to the on-premises Panorama appliance for storage and analysis.
D. The administrator views detailed logs and runs reports directly within the Prisma Access Cloud Management Console, which pulls data from Cortex Data Lake.
E. WildFire cloud service generates file download logs and stores them independently from other security event data.
Answer: B,D (LEAVE A REPLY)
Prisma Access, as a SASE offering, integrates cloud-based logging and management. - Option A (Incorrect): While endpoint security (like Cortex XDR) generates endpoint logs, Prisma Access security inspection happens at the cloud service edge, generating network- level logs. - Option B (Correct): Prisma Access service edges (the cloud-hosted firewalls processing user traffic) generate the various log types (traffic, threat, URL, file, etc.) just like a physical NGFW. These logs are automatically streamed to the centralized cloud logging service, Cortex Data Lake (CDL). - Option C (Incorrect): While Prisma Access can integrate with on-premises Panorama for unified management, logs are primarily stored in and accessed via Cortex Data Lake, which is a separate cloud service, rather than being sent directly to an on-premises Panorama (unless specifically configured for a hybrid logging setup, which is less common than using CDL). CDL is the default and scalable logging infrastructure for Prisma Access. - Option D (Correct): The administrator accesses and analyzes the logs stored in Cortex Data Lake through the Prisma Access Cloud Management Console (or potentially via other platforms like Cortex XSIAM that integrate with CDL). The console provides the interface to view, filter, and report on the log data residing in CDL. - Option E (Incorrect): WildFire provides analysis results, which are then recorded in the firewall's Threat logs (specifically as wildfire verdicts) and File logs. WildFire

doesn't independently store detailed logs of every file download; that information is in the traffic and file logs generated by the firewall, with the WildFire verdict referenced within them.

NEW QUESTION: 45
In a Zero Trust environment, granting access to a sensitive application should be based on multiple context factors, not just the user's network segment. A policy is needed to allow only Finance users, on company-issued laptops verified by GlobalProtect Host Information Profile (HIP) to be compliant (e.g., AV updated, disk encrypted), to access the Financial Planning application. This access must be subject to full threat inspection. Which combination of Palo Alto Networks policy elements and features is MOST critical for implementing this granular, context-aware Zero Trust access control?
A. Security Policy rule with Source Zone, Destination Zone, App-ID (Financial Planning App), User-ID (Finance Group), and a specific HIP Profile object in the Source User tab, along with relevant Content-ID profiles.
B. Decryption Policy rule configured for SSL Forward Proxy matching the Financial Planning application traffic and referencing the Forward Trust certificate.
C. NAT Policy rule configured to translate the source IPs of Finance users to a specific IP address for access control.
D. A Service object defined for the Financial Planning application's specific TCP/UDP ports.
E. Threat Prevention profile configured with a block action for critical severity vulnerabilities.
Answer: (SHOW ANSWER)
Implementing granular, context-aware access control in a Zero Trust model requires a security policy that verifies multiple attributes of the connection explicitly before granting access. Option A correctly lists the combination of elements that achieve this using Palo Alto Networks features: - Security Policy Rule: The central point for defining what traffic is allowed or denied. - Source Zone & Destination Zone: Basic zone- based segmentation (part of the network context). - App-ID: Identifies the specific 'Financial Planning Application', ensuring the policy applies only to that application, regardless of port. - User-ID: Identifies the 'Finance Group', ensuring only authorized users are considered. - HIP Profile object in the Source User tab: This is crucial for device posture verification. The HIP object represents the required state of the connecting device (company-issued, compliant based on AV, encryption, etc.), linking the user and device context to the policy. - Content-ID profiles (Threat, URL, WildFire, etc.): Applied to inspect the allowed traffic for threats and data exfiltration, fulfilling the 'Assume Breach' principle. Option B is necessary for inspecting encrypted traffic but doesn't define the access control criteria itself. Option C is a network translation function, not an access control mechanism for user/device context. Option D is a legacy approach focused on ports, not applications, and doesn't include user/device context. Option E is a security profile applied after access is granted, not the mechanism for granting the granular access based on user, device, and app.

NEW QUESTION: 46

An administrator is onboarding a new VM-Series firewall in a public cloud environment (e.g., AWS) and wants to manage it using Strata Cloud Manager (SCM). Unlike physical firewalls, VM-Series often leverage cloud-native capabilities for initial setup. Which method is commonly used for the initial setup and onboarding of a VM-Series firewall into SCM or Panorama in a cloud environment, facilitating Zero Touch Provisioning (ZTP)?

A. Manually configuring the management interface and pointing it to SCM's IP address.

B. Using cloud-init or user data scripts provided during VM launch to bootstrap the firewall with initial configuration and SCM registration details.

C. Connecting a serial console to the virtual machine for manual configuration.

D. Automatically discovering SCM via multicast on the cloud network.

E. Uploading a saved configuration file from the local firewall I-JI.

Answer: (SHOW ANSWER)

Cloud environments offer automation capabilities for VM deployment and configuration. - Option A: While basic connectivity is needed, relying solely on manual configuration after deployment isn't leveraging cloud automation. - Option B (Correct): Cloud platforms like AWS and Azure provide mechanisms (cloud-init for Linux, user data scripts) to inject scripts or configuration data during VM launch. This is commonly used to bootstrap the VM-Series firewall with its management IP, default gateway, DNS, and the information needed to register with SCM or Panorama for ZTP (e.g., authentication key, serial number, management IP of Panorama/SCM). This enables ZTP in the cloud. - Option C: Serial console access is possible but is a manual, legacy method not used for automated ZTP in cloud environments. - Option D: Multicast is generally not supported or used for management discovery in public cloud networks. - Option E: Uploading a saved configuration file is for restoring configuration, not initial onboarding to a management platform.

NEW QUESTION: 47

A large enterprise is migrating some internal applications to a cloud-based Software-as-a-Service (SaaS) model and implementing a SASE architecture leveraging Palo Alto Networks Prisma Access. They are encountering issues with the correct identification and enforcement of policies for a specific custom internal web application that now runs on a standard HTTPS port (443) alongside other legitimate SaaS traffic. The security team needs to ensure this custom application

is identified separately from general 'web-browsing' and enforce specific QOS and security profiles on it.
A. Create a custom application signature using App-ID based on unique characteristics of the application's payload or behavior, then create a security policy rule matching this custom App-ID.
B. Modify the default 'web-browsing' application signature to exclude traffic destined for the specific IP address/FQDN of the custom application.
C. Rely on Content-ID to identify the specific application content and apply policies based on content signatures instead of App-ID.
D. Configure a URL Filtering profile to block access to the custom application's URL, then allow it in a separate rule with the desired profiles.
E. Deploy a separate, dedicated Strata NGFW appliance specifically for this custom application traffic before it reaches Prisma Access.
Answer: (SHOW ANSWER)
Identifying custom or less common applications running on standard ports is a key use case for App-ID's custom application signature capabilities. Option A correctly describes the process: create a custom App-ID signature that looks for unique attributes of the application traffic (like specific HTTP headers, URL patterns, or payload content that identifies it as the custom app), and then use this custom App-ID in security policies to apply granular control and inspection. Option B is incorrect because modifying default signatures is not possible or recommended. Option C is incorrect; Content-ID focuses on threats and sensitive data within applications, not the identification of the application itself. App-ID is required for application identification and policy enforcement. Option D is a workaround using URL filtering but doesn't provide true application-level identification and control based on App-ID. Option E is impractical and defeats the purpose of a unified SASE architecture like Prisma Access.

NEW QUESTION: 48
A security team receives a BPA report via AIOps for NGFW highlighting a 'High' severity finding related to 'Policies Without Log Forwarding'. This finding indicates Security Policy rules configured without a log forwarding profile or with logging disabled, where logging is generally recommended. Which of the following are potential negative impacts of this configuration best practice violation?
(Select all that apply)
A. Reduced visibility into traffic flows matching these specific rules, making it difficult to audit access or investigate security incidents.
B. Failure to record sessions that trigger other security profiles (Threat, URL, etc.) applied by these rules.
C. Inability to utilize AIOps for NGFW's operational insights and reporting features for traffic matching these rules.
D. Increased load on the firewall's data plane due to improper policy configuration.
E. Difficulty in correlating security events (like threats) with the specific traffic session and policy rule that permitted or processed it.

Answer: (SHOW ANSWER)
Logging is fundamental to visibility, monitoring, and incident response. When logging is missing for policy rules, it creates blind spots. - Option A (Correct): The most direct impact is the lack of visibility into the traffic that matches these rules. You won't have records of who accessed what, when, and the result of the session. - Option B (Incorrect): Security profiles like Threat Prevention and URL Filtering generate their own specific logs (Threat logs, URL Filtering logs) when they detect an event, even if the traffic log for the base session is not generated due to policy logging being off. However, correlating these threat/IJRL logs back to the specific traffic flow becomes harder without the traffic log. -Option C (Correct): AIOps relies on logs (primarily traffic logs) for many of its operational and security insights (like application usage, User activity, session trends). If logging is disabled for certain rules, AIOps will not have the necessary data for traffic matching those rules, limiting its effectiveness. - Option D: Lack of logging doesn't typically increase data plane load; it's a control plane function. - Option E (Correct): Security investigations often start with a threat alert and require correlating it back to the originating session and the policy rule that handled it. Without traffic logs for the base session, this correlation becomes very challenging.

NEW QUESTION: 49
A SOC analyst receives an alert about a suspicious IP address attempting multiple login attempts across several endpoints. The analyst wants to automate the process of gathering intelligence on the IP before escalating the case.
Which Cortex XSOAR feature should be used to automate this enrichment process?
Response:
A. A Playbook that queries threat intelligence feeds and correlates IOCs
B. Manually forwarding the alert to another team for verification
C. Running a forensic investigation on each affected endpoint before taking action
D. Manually searching the IP address on different threat intelligence platforms
Answer: (SHOW ANSWER)

NEW QUESTION: 50
A. Source Zone: 'Mobile-UserS, Destination Zone: 'Public'
B. Source Zone: 'Remote-Networks' , Destination Zone: 'Service-Connection'
C. Source Zone: Service-Connection' , Destination Zone: Remote-NetworkS
D. Source Zone: 'Remote-NetworkS , Destination Zone: ' Remote-Networks'
E. Source Zone: 'Public', Destination Zone: 'Service-Connection'
Answer: (SHOW ANSWER)
Prisma Access uses specific zones for different traffic types and connection points. - Mobile-Users zone: Represents individual users connecting via GlobalProtect. - Remote-Networks zone: Represents traffic arriving from site-to-site VPN tunnels (branches, headquarters). - Service-Connection zone: Represents internal corporate resources (data center, cloud VPCs) accessed via tunnels from Prisma Access. - Public zone: Represents the public internet. Traffic from a Remote Network (branch) going to the corporate data center (Service Connection) would

originate from the 'Remote-Networks' zone and be destined for the 'Service-Connection' zone. Option A is for mobile users going to the internet. Option C is for traffic from the data center to the branch. Option D is for inter-branch traffic. Option E is for traffic from the internet to internal resources (though inbound access to Service Connections is less common than outbound from them).

NEW QUESTION: 51
An administrator needs to add a new PA-Series firewall at a remote branch office to their existing Panorama management deployment. The firewall is factory default. What initial configuration step is required on the new firewall itself before it can connect to and be managed by Panorama?
A. Configure the firewall's management interface IP address, subnet mask, default gateway, and DNS server.
B. Apply the full security policy configuration using the local web interface.
C. Install the latest PAN-OS software version and dynamic updates.
D. Establish an IPSec VPN tunnel to the Panorama appliance.
E. Configure Security Zones and assign interfaces to them.
Answer: (SHOW ANSWER)
For a firewall to connect to Panorama, it first needs basic network connectivity to reach the Panorama management interface over the network. This requires configuring its own management port IP settings. Option B, C, D, and E involve configuration that is typically pushed from Panorama after the firewall is connected and managed. The initial step is establishing basic network reachability to Panorama's management

NEW QUESTION: 52
Which log type in Palo Alto Networks Prisma SD-WAN (accessible via the Cloud Management Console/Cortex Data Lake) is specifically generated by the SD-WAN engine and provides visibility into which WAN links a particular application flow traversed, the quality metrics of that path at the time, and if any path changes occurred during the session?
A. Traffic logs
B. Path Monitoring logs
C. SD-WAN Flow logs
D. System logs
E. Tunnel logs
Answer: (SHOW ANSWER)
Prisma SD-WAN introduces specific log types related to the SD-WAN functionality itself. - Option A: Traffic logs show the security policy action and basic session info but don't typically provide detailed path selection information within the log entry itself. - Option B: While there are path monitoring views, 'Path Monitoring logs' as a distinct log type detailing per-flow path traversal isn't the standard term. - Option C (Correct): SD-WAN Flow logs (or similar terminology depending on specific console view/version, but conceptually the 'flow' logs capturing SD-WAN path details) are the logs that capture which path an application flow took across the SD-WAN fabric, including the

real-time path quality metrics (latency, jitter, loss) for that link at the time, and any path changes that occurred during the session lifecycle. This is distinct from standard security- focused traffic logs. - Option D: System logs are for appliance health. - Option E: Tunnel logs show the state of the tunnels (up/down) but not the per-flow path selection decisions.

NEW QUESTION: 53
In a GlobalProtect deployment using a Palo Alto Networks NGFW or Prisma Access, what is the primary role of a GlobalProtect Portal?
A. To terminate the secure tunnel from the GlobalProtect agent.
B. To provide the GlobalProtect agent software and initial client configurations to end-users.
C. To perform deep security inspection (Threat Prevention, URL Filtering) on user traffic.
D. To collect and fomard logs to Cortex Data Lake.
E. To act as the central management point for all GlobalProtect Gateways.
Answer: (SHOW ANSWER)
GlobalProtect architecture separates the Portal and Gateway functions. The Portal is the initial contact point for clients. - Option A: The Gateway terminates the tunnel. - Option B (Correct): The Portal's primary role is to authenticate the user, provide the GlobalProtect agent software installer, and deliver the client configuration (list of available Gateways, connection method, authentication settings, etc.). - Option C: Security inspection is performed by the Gateway. - Option D: Logging is handled by the Gateway and forwarded to CDL/Panorama. - Option E: Panorama or the Cloud Management Console is the central management point for Gateways and Portals.

NEW QUESTION: 54
A security team receives a BPA report via AIOps for NGFW highlighting a 'High' severity finding related to 'Policies Without Log Forwarding'. This finding indicates Security Policy rules configured without a log forwarding profile or with logging disabled, where logging is generally recommended. Which of the following are potential negative impacts of this configuration best practice violation?
(Select all that apply)
A. Increased load on the firewall's data plane due to improper policy configuration.
B. Inability to utilize AIOps for NGFW's operational insights and reporting features for traffic matching these rules.
C. Reduced visibility into traffic flows matching these specific rules, making it difficult to audit access or investigate security incidents.
D. Failure to record sessions that trigger other security profiles (Threat, URL, etc.) applied by these rules.
E. Difficulty in correlating security events (like threats) with the specific traffic session and policy rule that permitted or processed it.
Answer: (SHOW ANSWER)

Logging is fundamental to visibility, monitoring, and incident response. When logging is missing for policy rules, it creates blind spots. - Option A (Correct): The most direct impact is the lack of visibility into the traffic that matches these rules. You won't have records of who accessed what, when, and the result of the session. - Option B (Incorrect): Security profiles like Threat Prevention and URL Filtering generate their own specific logs (Threat logs, URL Filtering logs) when they detect an event, even if the traffic log for the base session is not generated due to policy logging being off. However, correlating these threat/lJRL logs back to the specific traffic flow becomes harder without the traffic log. -Option C (Correct): AIOps relies on logs (primarily traffic logs) for many of its operational and security insights (like application usage, User activity, session trends). If logging is disabled for certain rules, AIOps will not have the necessary data for traffic matching those rules, limiting its effectiveness. - Option D: Lack of logging doesn't typically increase data plane load; it's a control plane function. - Option E (Correct): Security investigations often start with a threat alert and require correlating it back to the originating session and the policy rule that handled it. Without traffic logs for the base session, this correlation becomes very challenging.

NEW QUESTION: 55
An organization wants to protect its users from accessing known malicious websites and command-and-control (C2) infrastructure by preventing the resolution of malicious domain names. They have a Palo Alto Networks NGFW with an Advanced DNS Security subscription. Which key capability provided by Advanced DNS Security enables this protection at the DNS layer?
A. Encrypting all DNS queries to prevent eavesdropping.
B. Analyzing DNS query and response patterns using machine learning to identify malicious domains in real-time.
C. Serving as a local DNS resolver for all internal clients.
D. Blocking DNS traffic based on the source IP address of the querying host.
E. Comparing DNS query domain names against a static blacklist configured manually on the firewall.
Answer: (SHOW ANSWER)
Advanced DNS Security is a cloud-delivered service that uses advanced analytics to identify malicious domains at the DNS layer. Option A describes DNS encryption (DNSSEC or DNS over HTTPS/TLS), which enhances privacy but doesn't inherently detect malicious domains. Option B correctly describes the core of Advanced DNS Security: using machine learning and threat intelligence (often correlated with WildFire, Threat Prevention, etc.) to analyze DNS queries and responses and identify malicious domains in near real-time. Option C is a function of a DNS server, not the security analysis provided. Option D is basic firewall filtering. Option E describes a basic, manual approach that doesn't scale and misses dynamic threats.

NEW QUESTION: 56

A key benefit of using Prisma Access compared to self-managed firewalls (PA-SeriesNM-Series) for remote user and branch security is that the responsibility for performing the underlying software upgrades and patching of the security processing nodes lies primarily with whom?
A. The customer administrator via the Cloud Management Console.
B. The customer administrator via Panorama.
C. The end-user via the GlobalProtect client.
D. Palo Alto Networks as the cloud service provider.
E. A third-party managed security service provider (MSSP).
Answer: (SHOW ANSWER)
Prisma Access is a cloud-delivered security service. A significant advantage of this model is that Palo Alto Networks, as the service provider, is responsible for the ongoing maintenance, including software upgrades and patching, of the underlying security processing nodes and infrastructure. This offloads a major operational burden from the customer's IT team. Options A, B, C, and E are incorrect; these parties are not primarily responsible for upgrading the core Prisma Access infrastructure.

NEW QUESTION: 57
A company wants to implement a Zero Trust policy where access to the internal development code repository application is only allowed for members of the 'DevTeam' Active Directory group if they are connecting from a device identified as a 'Company Laptop' and the device posture is compliant (e.g., antivirus updated, disk encrypted), as verified by GlobalProtect HIP. Which specific Palo Alto Networks features and policy configurations are essential to achieve this granular control on a Strata NGFW or Prisma Access?
A. Configure App-ID to identify the 'development-repo' application and use it in a Security policy rule's 'Application' tab.
B. Ensure User-ID is configured and operational to map user IPs to AD user accounts/groups and use the 'DevTeam' group in the Security policy rule's 'Source User' tab.
C. Configure GlobalProtect with Host Information Profile (HIP) collection and define a HIP Object that represents the 'compliant company laptop' posture, then reference this HIP Object in the Security policy rule's 'Source User' tab.
D. Use Device-ID to identify the device as a 'Company Laptop' and incorporate this Device-ID into the Security policy rule criteria.
E. Create a custom service object for the development repository's port and protocol, and use this service object in the Security policy rule.
Answer: (SHOW ANSWER)
Achieving this granular, context-aware access control requires combining identity (User-ID), application identification (App-ID), and device context (Device-ID/HIP). Let's break down the options: - Option A (Correct): App-ID is essential to identify the specific application traffic ('development-repo') independent of ports, ensuring the policy applies precisely. - Option B (Correct): User-ID is required to identify the user as a member of the 'DevTeam' group, enabling identity-based policy. - Option C (Correct): GlobalProtect HIP is the mechanism to collect device

posture information. Defining a HIP Object for the 'compliant company laptop' posture and referencing it in the Security policy rule's 'Source User' tab (alongside or in conjunction with the User-ID group) allows the firewall to enforce policy based on device compliance. - Option D (Correct): Device-ID provides visibility into the device type (e.g., Windows laptop, iPhone, IoT device). While HIP provides posture, Device-ID identifies the device itself. In this scenario, identifying it as a 'Company Laptop' device type (which Device-ID can often infer from DHCP options, user-agent strings, etc., or via integrated endpoints) is a valid policy criterion, often used in conjunction with or as part of HIP requirements, to ensure the user isn't connecting from a personal phone, for example. - Option E (Incorrect): Using a Service object based on port/protocol is a legacy approach that bypasses the granular application identification provided by App-ID and does not incorporate user or device context.

NEW QUESTION: 58
A branch office using Prisma SD-WAN has a direct internet link. They need to allow guest Wi-Fi users to access the internet, but this guest traffic should be Source NAT'd to a different public IP address range than corporate user traffic to facilitate separate logging and rate limiting by the upstream ISP. The guest network uses a specific VLAN and subnet (172.16.10.0/24). Which Prisma SD-WAN policy type and configuration element is used to define this specific NAT requirement for the guest traffic?
A. A Security Policy rule matching the guest subnet and applying a custom NAT profile.
B. A Path Policy rule matching the guest subnet and directing traffic to a NAT gateway.
C. A NAT Policy rule with the Original Packet Source Zone/Subnet matching the guest network (172.16.10.0/24) and Translated Packet Source Translation configured with a specific static IP or dynamic pool.
D. A QOS Policy rule prioritizing guest traffic and applying a NAT action.
E. Application Override policy configured to classify guest traffic for specific NAT handling.
Answer: C (LEAVE A REPLY)
Defining how specific source traffic (like guest users) is translated when exiting the network is the function of NAT Policy. - Option A: Security Policy determines allow/deny and inspection, not NAT translation rules. - Option B: Path Policy determines which link traffic goes over, not how its address is translated. While traffic might be steered to a link where NAT is performed, the NAT definition itself is separate. - Option C (Correct): NAT Policy is where you configure address translation. You create a rule that matches the 'Original Packet' details (source zone/subnet of the guest network, destination zone/interface like the internet egress). In the 'Translated Packet' section, you configure the Source Address Translation method (Static IP or Dynamic IP/Port) using the specific public IP or pool designated for guest traffic. This ensures only traffic from the guest subnet gets this specific translation. - Option D: QOS Policy prioritizes bandwidth usage; it does not perform NAT. - Option E: Application Override reclassifies traffic for App-ID purposes; it doesn't configure NAT.

NEW QUESTION: 59

An organization uses Prisma Access for mobile users and logs to Cortex Data Lake. A user reports slow performance when accessing a SaaS application. The administrator suspects network latency between the user and the closest Prisma Access location or between Prisma Access and the SaaS provider, or potentially high load on the assigned Prisma Access node. Which log types or monitoring views in Cortex Data Lake or the Cloud Management Console could help diagnose these potential performance bottlenecks? (Select all that apply)
A. GlobalProtect logs, specifically looking for login success/failure and gateway assignment.
B. Performance monitoring views in the Cloud Management Console showing metrics for the user's connected Prisma Access location (e.g., resource utilization, latency to internet/service connections).
C. Traffic logs showing the session details for the SaaS application, including bytes transferred and session duration, correlated with timestamps.
D. Application Performance Monitoring (APM) data (if applicable) showing latency and performance specific to the SaaS application over the Prisma Access path.
E. System logs on the Prisma Access service edge showing daemon restarts or critical errors.
Answer: (SHOW ANSWER)
Troubleshooting performance in a SASE environment involves looking at network path performance, application performance metrics, resource utilization, and session details. - Option A: GlobalProtect logs confirm connection status but don't show performance within the tunnel. - Option B (Correct): Monitoring views showing performance metrics for the Prisma Access location itself provide insight into potential bottlenecks at the cloud edge or connectivity issues from the edge to destinations. - Option C (Correct): Traffic logs, when analyzed for session duration relative to bytes transferred, can indicate slowness (e.g., long duration for small data transfer). While not showing latency directly, they provide session activity context. - Option D (Correct): APM data is specifically designed to measure application performance over the network, showing latency and other quality metrics from the user to the application. - Option E (Correct): System logs can indicate if the underlying Prisma Access node handling the user's traffic is experiencing issues (high CPU, memory pressure, restarts) that would impact performance.

NEW QUESTION: 60
A. Analysis of the evasive executable file in the WildFire sandbox to observe its malicious behavior (e.g., process injection, file modification, network connections).
B. Generation of new signatures (Antivirus, Antispyware, Vulnerability) based on the analysis of the evasive executable, which are then distributed globally.
C. Identification of the suspicious DNS request destination as a newly registered or malicious domain via DNS Security (a related CDSS leveraging WildFire intelligence).
D. Correlation of behavioral indicators from the endpoint (e.g., process creation, registry changes) with network events from the firewall via a unified platform like Cortex XDR (leveraging WildFire verdicts).
E. Real-time blocking of the evasive executable file upon first encounter based on a static hash lookup before submission to the sandbox.

Answer: (SHOW ANSWER)
Advanced WildFire and integrated CDSS provide multi-faceted detection for sophisticated threats. - Option A (Correct): The core of WildFire is dynamic analysis. Executing the file in a sandbox reveals its true behavior, even if it's evasive, allowing detection based on actions rather than just signatures. - Option B (Correct): A key value of WildFire is its feedback loop. When new malware is identified in the sandbox, Palo Alto Networks generates and rapidly distributes new signatures (Antivirus, Threat Prevention) and indicators (URLs, IPs, domains) globally to all subscribers, enabling rapid protection against the newly discovered threat. - Option C (Correct): DNS Security is a CDSS that leverages intelligence, including from WildFire analysis, to identify and block access to malicious or suspicious domains, including newly created C2 domains. WildFire analysis can reveal C2 communication attempts to such domains, feeding this intelligence into DNS Security. - Option D (Correct): Cortex XDR integrates endpoint and network security data. WildFire verdicts and related logs from the firewall, combined with endpoint telemetry (process activity, file changes), enable the correlation needed to detect complex attacks like APTs that involve multiple stages and behaviors. - Option E (Incorrect): Real-time blocking on first encounter is the goal, but if the file is truly unknown and evasive, a static hash lookup (which is for known malware) won't block it. WildFire provides 'inline ML' and rapid analysis results for near real-time prevention of zero-day threats, but blocking on first encounter based purely on hash isn't how zero-day detection works; it's based on analysis after encountering the file.

NEW QUESTION: 61
A company implements strict web access policies using Advanced URL Filtering on their Palo Alto Networks NGFW. They configure a URL Filtering profile to block the 'Social-Networking' category for all users. However, a security analyst notices that some specific social media websites are still being accessed, and the traffic logs show them being categorized as 'none' or a general category like Wveb- services'. What is a possible reason for this miscategorization or bypass of the blocking policy, and how can it be addressed?
A. SSL Decryption is not enabled for HTTPS traffic to these websites, preventing the firewall from seeing the full URL for categorization.
B. The URL Category database on the firewall is outdated and needs to be manually updated.
C. The specific websites in question are new or less common and have not yet been categorized correctly in the cloud database.
D. A custom URL Category needs to be created for the miscategorized websites and set to 'block' in the URL Filtering profile, placed higher than the 'Social-Networking' rule.
E. A Security Policy rule allowing traffic to these specific websites is placed above the rule applying the URL Filtering profile.
Answer: (SHOW ANSWER)
Misclassification or bypass in URL Filtering can occur due to various factors: - Option A (Correct): For HTTPS traffic, the firewall typically sees the hostname via SNI before decryption. However, full URL path categorization and advanced features like real-time analysis require decryption to see the entire request. If decryption is not enabled for these sites, categorization might be based

only on the hostname, potentially leading to a less accurate or 'none' category. - Option Option B (Incorrect): Advanced URL Filtering relies on a cloud-based database, which is dynamically updated, not manually on the firewall (updates happen automatically). - Option C (Correct): Even with Advanced URL Filtering's real-time analysis, new or less common websites might not be immediately or correctly categorized. There's a delay between a site appearing and being fully classified in the cloud database. - Option D (Correct): If specific URLs are consistently miscategorized, creating a custom URL Category for those URLs and explicitly setting the action (e.g., 'block') for that custom category in the URL Filtering profile is a manual override to ensure they are blocked as desired. Custom categories are evaluated before built-in categories. - Option E (Incorrect): A Security Policy rule allowing traffic comes before the IJRL Filtering profile is applied. If an earlier rule allows the traffic without a IJRL Filtering profile, or if the URL Filtering profile applied allows the category, it won't be blocked by a later URL Filtering rule. However, the question implies the traffic hits the policy with the profile but is miscategorized.

NEW QUESTION: 62
A security administrator logging into the AIOps for NGFW dashboard needs a quick overview of the overall health, security posture, and potential operational issues across their fleet of managed firewalls. Which sections or widgets on the AIOps dashboard are designed to provide this high-level summary information?
A. Best Practices Assessment score and findings summary.
B. Operational Status dashboard, showing critical alerts and key performance indicators (KPIs).
C. Security Policy rule usage statistics.
D. Configuration logs viewer.
E. Detailed threat log viewer.
Answer: (SHOW ANSWER)
AIOps dashboards are designed for quick visibility and actionable insights. - Option A (Correct): The Best Practices Assessment score provides a quantitative measure of how well firewalls align with recommended configurations, and the summary highlights key findings (policy, network, device best practices), giving a high-level security posture view. - Option B (Correct): The Operational Status dashboard (or similar section depending on version) provides critical alerts related to device health, resource utilization, licensing, and key performance metrics, offering a snapshot of operational health. - Option C: While usage statistics are available, they are typically

detailed reports, not a primary high-level summary widget. - Option D and E: Log viewers are for detailed investigation, not high-level dashboards.

NEW QUESTION: 63
A network engineer is tasked with deploying a new Prisma SD-WAN ION device at a branch office. After physically installing the device and connecting the necessary cables, the next step is the initial setup process to onboard the device into the Prisma SD-WAN Cloud Management Console. What is the primary method used for the initial bootstrapping and activation of a new ION device?
A. Manually logging into the ION device's local web interface and entering cloud management credentials.
B. Connecting to the ION device via serial console or SSH and running a setup wizard script.
C. Using a Zero Touch Provisioning (ZTP) process that involves connecting the device to the internet and potentially using a USB stick with a configuration file or entering a serial number/one-time key in the cloud console.
D. The ION device automatically discovers the cloud controller on the network via broadcast.
E. Connecting the ION device directly to a Panorama appliance for initial configuration.
Answer: (SHOW ANSWER)
Prisma SD-WAN ION devices are designed for ease of deployment, often leveraging Zero Touch Provisioning (ZTP). Option C describes the ZTP process: the device, upon booting and gaining internet connectivity (often via a temporary link or one of its WAN interfaces), contacts the cloud controller and obtains its initial configuration and management connection details. This might involve manual steps in the cloud console to associate the device serial number with a site or configuration template, or using a preloaded USB key for some initial network parameters. Options A and B describe methods for manual local configuration, which might be used for troubleshooting but not the primary ZTP onboarding. Option D is incorrect; discovery is not typically via local broadcast. Option E is incorrect; ION devices are cloud-managed, not directly by Panorama for initial setup.

NEW QUESTION: 64
An organization is using Panorama to manage its PA-Series firewalls and has integrated Prisma Access logging with Panorama's Log Collector. The security team wants to generate a report that shows all traffic sessions that were denied by any security policy rule across all managed firewalls and Prisma Access nodes, grouped by the denying policy rule name and showing the source user and destination application. Which of the following steps or considerations are necessary to build this comprehensive report in Panorama? (Select all that apply)
A. Ensure that all relevant Security Policy rules on managed firewalls and Prisma Access are configured with logging enabled.
B. Ensure that traffic logs from all managed firewalls and Prisma Access nodes are successfully being forwarded to the Panorama Log Collector.

C. Create a custom report in Panorama's Monitor > Reports tab, filtering for Log Type 'Traffic' and Action 'deny'.
D. Include columns for 'Rule Name', 'Source User', and 'Application' in the custom report definition.
E. Generate the report using System logs, as they contain policy violation details.
Answer: (SHOW ANSWER)
Generating comprehensive reports across multiple devices/services requires data availability and correct reporting configuration. - Option A (Correct): Policy rule logs must be enabled on the individual firewalls/Prisma Access nodes. If a deny rule doesn't have logging enabled, sessions hitting it won't be recorded in the traffic logs. - Option B (Correct): Logs must be successfully collected in Panorama (or CDL if Panorama is forwarding to it). If logs are not forwarded correctly, the central repository won't have the data. - Option C (Correct): You use the 'Traffic' log type because it contains details about allowed/denied sessions, and you filter for the 'deny' action. - Option D (Correct): To see the requested information (rule name, user, application), you must include these fields as columns in the report output. The firewall logs capture this information (assuming User-ID and App-ID were operational). - Option E (Incorrect): System logs are for firewall operational events, not details of denied traffic sessions.

NEW QUESTION: 65
A network administrator notices high CPU utilization and lower than expected throughput on a Palo Alto Networks NGFW during peak hours, despite the total bandwidth usage being well within the hardware capabilities. Reviewing system metrics shows a significant number of new sessions being established per second compared to the overall Mbps throughput. Which configuration or traffic pattern is MOST likely contributing to excessive slow path processing and causing the performance bottleneck?
A. A large volume of long-lived, established HTTP sessions with basic Threat Prevention profiles enabled.
B. Extensive use of Security policies with source/destination NAT configured, primarily for outbound internet traffic.
C. A sudden surge in traffic consisting of many short-lived connections to unique destination IPs/ports, potentially using varied applications or protocols.
D. Heavy traffic consisting mainly of UDP-based video streaming using an established, identified App-I
E. Security policies allowing inter-zone traffic with no security profiles applied.
Answer: (SHOW ANSWER)
High CPU utilization coupled with a high rate of new sessions per second, despite relatively low overall bandwidth, is a strong indicator that the firewall is spending a disproportionate amount of time processing the first packet of many sessions, which occurs on the slow path. The slow path is CPU-intensive because it involves App-ID lookup, policy matching, session creation, NAT/routing decisions, and security profile assignment. - Option A: Long-lived, established sessions are primarily handled by the fast path after the initial setup. While security profiles add

some overhead, the core processing of established flow is hardware-accelerated, not CPU bound for simple forwarding. - Option B: While NAT involves slow path processing for the first packet (or connections requiring dynamic NAT allocation), established sessions with NAT are handled efficiently by the fast path using the created session state. - Option C (Correct): A large volume of short-lived connections, especially if they vary widely in destination and application, means the firewall must process the first packet of each connection individually on the slow path. This puts a heavy load on the CPU for session setup, even if the data transferred within each session is small. This is a classic scenario causing high 'sessions per second' and thus high slow-path CPU load. - Option D: Established UDP sessions, once identified by App-ID and allowed by policy, are also typically handled efficiently by the fast path (or hardware session acceleration), similar to TCP established sessions. - Option E: Policies allowing traffic with no security profiles still require App-ID identification and policy lookup for the first packet, putting it on the slow path for session creation. However, this processing is generally less intensive than processing requiring deep inspection, and the bottleneck described points to the volume of new sessions overwhelming the CPU's ability to perform the initial setup, which is exacerbated by complex policies or varied traffic, but fundamentally driven by the 'new session' rate.

NEW QUESTION: 66
A company needs to provide secure network access for its employees working remotely from various locations. They require a solution that establishes an encrypted tunnel to the corporate network (or a cloud security platform), supports multi-factor authentication, and allows for policy enforcement based on user identity and device compliance. Which Palo Alto Networks product or service is specifically designed to meet these remote access requirements for mobile users?
A. Prisma SD-WAN ION devices
B. PA-Series firewalls deployed as internet edge devices.
C. GlobalProtect (Client, Gateways, Portals)
D. VM-Series firewalls deployed in the cloud.
E. Cloud NGFW for AWS.
Answer: (SHOW ANSWER)
GlobalProtect is Palo Alto Networks' comprehensive remote access solution for mobile users. It consists of the GlobalProtect client software on the endpoint, GlobalProtect Gateways (on NGFWs or Prisma Access) that terminate the encrypted tunnels, and GlobalProtect Portals for client configuration and authentication. It fully supports user identity (User-ID integration), multi-factor authentication, and device posture checking (HIP). Option A is for site-to-site SD-WAN. Options B, D, and E are firewall form factors that can host GlobalProtect Gateways but aren't the solution name itself.

NEW QUESTION: 67
A company is implementing SSL Forward Proxy decryption for outbound internet traffic using a Palo Alto Networks NGFW. After deploying the firewall's Forward Trust Certificate to employee laptops via GPO, users accessing some internal applications and certain external banking

websites report certificate errors or connection failures. Which of the following are potential reasons for these issues and how certificates play a role? (Select all that apply)
A. The internal applications use client-side certificates for authentication, which is disrupted by the firewall's MITM decryption process.
B. The banking websites use certificate pinning, causing the client browser to reject the certificate re-signed by the firewall's Forward Trust CA.
C. The firewall is configured to use the Forward Untrust Certificate for these sites, causing browsers to explicitly warn users.
D. The Forward Trust Certificate was not successfully installed or trusted in the certificate store of the user's device or specific application.
E. The firewall's Decryption policy rule for these sites is set to 'No Decrypt', causing connection failures.
Answer: (SHOW ANSWER)
SSL Forward Proxy acts as a Man-in-the-Middle, and certificate handling is critical for its success and potential issues. - Option A (Correct): Client-side certificates are presented by the client to the server for authentication. The firewall intercepting the connection cannot present the client's private key, breaking this type of authentication. - Option B (Correct): Certificate pinning means the client trusts only a specific certificate (hash or public key) from the server. The firewall presents a different certificate (signed by its CA), which the client rejects. - Option C: The Forward Untrust Certificate is used for sites with certificate errors or unknown status to explicitly warn users or block access, but the primary issue with trusted sites or internal apps is disruption caused by the MITM, not intentionally marking them untrusted. - Option D (Correct): If the firewall's Forward Trust Certificate is not installed and trusted on the client, the client will not trust any certificate signed by it, leading to certificate errors or warnings for sites that are decrypted. - Option E: Setting a rule to 'No Decrypt' would typically bypass decryption for those sites, preventing issues caused by the decryption process, not cause connection failures (unless combined with other policies).

NEW QUESTION: 68
When utilizing Cortex Data Lake (CDL) for centralized logging from various Palo Alto Networks platforms (NGFWs, Prisma Access, Prisma SD-WAN), what is a key advantage compared to using local firewall logging or individual syslog servers at each location?
A. CDL provides unlimited, perpetual log storage for all log types.
B. CDL aggregates logs from all connected devices and services into a single, searchable, and correlatable repository.
C. CDL performs real-time security enforcement based on log analysis.
D. CDL eliminates the need for administrators to configure logging on individual firewalls.
E. CDL can collect logs from any network device, regardless of vendor.
Answer: (SHOW ANSWER)
Centralized logging platforms are designed for scalability, aggregation, and ease of analysis. - Option A: CDL provides scalable storage, but it is typically licensed based on ingest rate and data

retention period, not unlimited and perpetual. - Option B (Correct): The primary advantage of CDL is its ability to receive and store logs from all supported Palo Alto Networks sources in a unified cloud-based repository, enabling administrators to search, filter, report, and correlate events across the entire distributed environment from a single interface (like the Cloud Management Console or Panorama). This is crucial for comprehensive visibility and incident response. - Option C: CDL is a logging and analytics platform; security enforcement actions are performed by the firewalls/Prisma Access/SD-WAN devices based on their policies. - Option D: Administrators still need to configure logging profiles and apply them to policy rules on the firewalls/services to specify which logs are generated and where they are forwarded (to CDL). - Option E: CDL is specifically designed for logs from Palo Alto Networks products.

NEW QUESTION: 69
Prisma SD-WAN leverages application identification for intelligent traffic steering and optimization. How does the combination of App-ID and WAN optimization features in Prisma SD-WAN enhance application performance compared to traditional, port-based WAN optimization solutions?
A. It allows the SD-WAN appliance to apply the same universal optimization techniques (like basic compression) to all traffic equally, simplifying configuration.
B. It enables the SD-WAN appliance to identify the specific application (e.g., SharePoint, Oracle, Zoom) and apply optimization techniques and path selection policies specifically tailored to that application's requirements and sensitivity to latency, jitter, and bandwidth.
C. It ensures that only traffic using standard, well-known ports (like 80, 443) receives optimization, ignoring traffic on non-standard ports.
D. It primarily helps in blocking malicious applications, with optimization being a secondary, unrelated function.
E. It offloads the App-ID processing entirely to the central Panorama appliance, freeing up local SD-WAN appliance resources for optimization.
Answer: (SHOW ANSWER)
The application-aware nature of Palo Alto Networks' platforms, extended to Prisma SD-WAN, is a key differentiator. - Option A (Incorrect): A primary benefit is not applying universal techniques. Different applications benefit from different techniques (VoIP needs low latency/loss paths, file transfer benefits from data reduction). App-ID allows for differentiation. - Option B (Correct): By identifying the application precisely using App-ID (independent of port), Prisma SD-WAN can apply application-specific policies. This means voice/video gets prioritized and steered over low-latency/low-loss paths (Performance sensitive profile), file transfers get data reduction (Bandwidth sensitive profile), and critical business applications get guaranteed bandwidth or preferred paths. This granular, intelligent approach is a major advantage over port-based systems. - Option C (Incorrect): App-ID identifies applications regardless of the port they use, including applications running on non-standard ports or within encrypted tunnels (if decrypted). - Option D (Incorrect): While Prisma SD-WAN integrates security, the primary benefit of combining App-ID with optimization is enhanced application performance and user experience , not primarily blocking

applications. - Option E (Incorrect): App-ID processing occurs on the local NGFW/SD-WAN appliance itself as traffic passes through it; it's fundamental to the real-time processing chain.

NEW QUESTION: 70

A. Configure a URL Filtering profile with the forbidden URL categories set to the 'block' action.
B. Configure a URL Filtering profile with the risky URL categories set to the 'continue' action and customize the 'continue' page message.
C. Configure a Data Filtering profile to detect sensitive data patterns and apply it to a Security Policy rule with 'upload' App Functions for file sharing/webmail, set to a 'block' action.
D. Configure a Threat Prevention profile with signatures for detecting specific sensitive data patterns within HTTP/HTTPS traffic.
E. Apply the configured URL Filtering profile and the configured Data Filtering profile to the relevant Security Policy rules that allow outbound web and application traffic.
Answer: (SHOW ANSWER)
This scenario requires applying policies based on both IJRL category and sensitive data content, with different actions. - Option A (Correct): Blocking URL categories is done in the URL Filtering profile by setting the desired categories to the 'block' action. - Option B (Correct): Providing a warning requires the 'continue' action in the URL Filtering profile for the specific category. The warning message is customizable. - Option C (Correct): Preventing sensitive data upload is the function of the Data Filtering profile. The profile detects the patterns, and the Security Policy rule applying this profile (matching upload activities) is set to 'block' or 'alert' when a match occurs. - Option D (Incorrect): Threat Prevention is for malware/exploits, not sensitive data patterns. Sensitive data detection is done via the Data Filtering profile with the DLP subscription. - Option E (Correct): Once the profiles are configured, they must be applied to the relevant Security Policy rules to enforce the actions on matching traffic. Options A and B handle the URL category actions. Option C handles the sensitive data detection and action. Option E ties the profiles to the traffic flows via security policy.

NEW QUESTION: 71

A large manufacturing facility has deployed numerous IoT devices (sensors, cameras, controllers) on a dedicated network segment.
These devices are known for having weak security controls and often communicate using proprietary or insecure protocols, potentially accessing external cloud services. The security team wants to gain visibility into these devices, identify risky behavior, and enforce granular policies to restrict their communication. Which Palo Alto Networks capability, often leveraging Cloud-Delivered Security Services (CDSS), is specifically designed to provide visibility and security enforcement for previously unmanaged or poorly understood IoT devices?
A. App-ID with custom signatures
B. User-ID with Captive Portal
C. IoT Security subscription
D. Standard Threat Prevention signatures

E. URL Filtering with category blocking

Answer: (SHOW ANSWER)

Securing diverse and often unmanaged IoT devices requires specialized capabilities beyond traditional firewall features. Palo Alto Networks offers a dedicated IoT Security subscription (often tightly integrated with NGFWs/Prisma SASE) that leverages cloud-based machine learning and threat intelligence to profile devices, identify risks, and generate recommended policies. Option A is useful for identifying known applications but struggles with the vast, unknown IoT device landscape. Option B is for user authentication, not device identification or behavior analysis. Option D and E are for general threat and web filtering, less effective at identifying the devices themselves or their specific risky behaviors within proprietary protocols. The IoT Security subscription is the specialized solution for this challenge.

NEW QUESTION: 72

An organization relies heavily on Cortex Data Lake (CDL) for logging and analytics from its Prisma Access deployment. They are integrating CDL with a third-party Security Information and Event Management (SIEM) system for centralized security monitoring and alerting. Which types of logs generated by Prisma Access and stored in CDL are MOST critical for providing comprehensive visibility into user activity, security threats, and policy enforcement for remote users and remote networks? (Select all that apply)

A. Traffic logs (showing allowed/denied sessions with App-ID and User-ID)

B. Threat logs (detailing detected malware, exploits, etc.)

C. URL Filtering logs (recording web access attempts and categories)

D. HIP Match logs (indicating device posture compliance status)

E. Configuration logs (tracking changes to Prisma Access setup)

Answer: (SHOW ANSWER)

For security monitoring and SIEM integration, logs that capture traffic flow, detected threats, user activity, and device compliance are essential. - Option A (Correct): Traffic logs are fundamental, providing records of every session, including which policy ruled it, the application, user, and action taken. This gives baseline visibility into network activity. - Option B (Correct): Threat logs are critical for identifying and investigating security incidents. They contain details about malware detections, exploit attempts, command-and-control traffic, etc. - Option C (Correct): URL Filtering logs show user web browsing activity, which is vital for enforcing acceptable use policies, identifying risky websites, and detecting access to malicious URLs. - Option D (Correct): HIP Match logs provide visibility into the compliance status of connecting devices. This is crucial for Zero Trust implementations where access or policy might depend on device posture. - Option E (Incorrect): Configuration logs track changes to the system itself, which is important for auditing and change management but less critical for real-time security monitoring of user traffic and threats compared to the other log types.

NEW QUESTION: 73

A company is using Prisma Access to provide secure internet access for its remote workforce. They have configured Security Policy rules that leverage User-ID, App-ID, URL Filtering, Threat Prevention, and Decryption for outbound traffic. Users report that access to a newly deployed SaaS application is being blocked by the Prisma Access policy, and traffic logs show the session hitting the default 'deny' rule. Troubleshooting indicates that the required security policy rule intended to allow the application is not being matched. Which of the following are potential reasons why the traffic is not matching the intended 'allow' security policy rule for the SaaS application? (Select all that apply)

A. App-ID is not correctly identifying the new SaaS application, causing the 'Application' field in the policy rule to not match.

B. SSL Forward Proxy decryption is failing for the new SaaS application's traffic, preventing accurate App-ID identification or policy evaluation.

C. User-ID is not successfully mapping the user's IP address to their username or group, preventing the 'Source User' field in the policy rule from matching.

D. A more specific 'deny' rule is placed higher in the policy list and is matching the traffic before it reaches the intended 'allow' rule.

E. The destination IP addresses used by the SaaS application are not included in the 'Public' zone definition.

Answer: (SHOW ANSWER)

If traffic hits the default deny, it means no preceding allow or deny rule matched. Troubleshooting involves checking the criteria of the intended rule and rules above it, and ensuring the firewall has the information needed to evaluate those criteria. - Option A (Correct): If App-Ld doesn't recognize the application, a rule specified with that application's App-ID will not match. This is a common issue with new or custom applications. - Option B (Correct): Decryption failure can impactApp-ID accuracy, especially for distinguishing applications on standard ports like 443. If App-ID relies on seeing content after decryption, and decryption fails, the application might be misidentified or identified as 'unknown', preventing the rule match. - Option C (Correct): If the rule includes a 'Source User' criterion, and User-ID isn't working for that user's session, the rule requiring a specific user or group will not match. The session would likely show 'unknown' user in the logs. - Option D (Correct): Security policy rules are evaluated top-down. A more specific deny rule higher up (e.g., denying access to certain URL categories, source IPs, or applications) could be blocking the traffic before it reaches the intended allow rule. - Option E (Incorrect): The 'Public' zone typically represents the entire internet. Destination IP addresses are evaluated against routing and zones, but the zone definition usually encompasses all public IPs, not requiring specific inclusion of SaaS IPs within the zone itself (though address objects could be used in policies within the zone context).

NEW QUESTION: 74

An administrator configures a new VLAN interface on a Palo Alto Networks Strata NGFW and assigns it to an existing Security Zone named 'VLAN-Zone'. The administrator then attempts to create a Security Policy rule allowing traffic from 'Internal-Users' zone to However, traffic between

these zones fails, and logs show the traffic hitting the implicit 'deny' rule, even though interfaces are correctly configured and IP routing is working. Which configuration aspect related to zones and interfaces was MOST likely overlooked?
A. The interfaces in the 'VLAN-Zone' were configured as Layer 2 interfaces instead of Layer 3 interfaces.
B. The new VLAN interface was not explicitly assigned to the 'VLAN-Zone' during configuration.
C. Security Policy rules are processed top-down, and a broader 'deny' rule above the new rule is blocking the traffic.
D. The 'Internal-Users' zone is configured as a 'Tap' zone, which does not permit traffic forwarding.
E. The Zone Type for 'VI-AN-Zone' was set to 'External' instead of 'Internal'.
Answer: (SHOW ANSWER)
For a security policy rule defined between two zones (e.g., 'Internal-Users' and ' VLAN-Zone') to be evaluated and potentially matched by traffic flowing through the firewall, the interfaces where that traffic enters and exits the firewall must be assigned to the respective source and destination zones specified in the policy rule. If the new VLAN interface intended for the 'VLAN-Zone' was created but not explicitly associated with the 'VI-AN-Zone' object in the configuration, traffic coming in on that interface will not be seen as originating (or destined for, depending on direction) the 'VI-AN-Zone' , and thus will not match the zone-based policy rule. The traffic then proceeds down the rule list and hits the implicit deny. Option A describes an interface mode, but the core issue is the zone assignment itself. Option C is a general policy troubleshooting step but doesn't address the initial problem of the traffic not being associated with the correct zone for policy lookup. Option D describes a specific zone type that wouldn't forward traffic, but the question implies the zone configuration is correct, while the interface assignment might be missing. Option E is irrelevant; the zone name and type are logical labels for policy, not direct blockers like the lack of interface assignment.

NEW QUESTION: 75
A large organization is deploying SSL Forward Proxy decryption across its SASE infrastructure (Palo Alto Networks Prisma Access) for global users accessing the internet. After initial rollout, they encounter several challenges, including users reporting certificate errors on specific websites and internal applications, and some applications failing to function correctly when decryption is enabled. Which of the following are common reasons for these issues and crucial considerations when implementing SSL Forward Proxy?
A. The firewall's Forward Trust Certificate (the root CA used to re-sign certificates) has not been deployed and trusted by all client devices' operating systems or browser trust stores.
B. Some applications utilize security mechanisms like certificate pinning, where the client application is hardcoded to trust only the original server certificate, causing it to reject the certificate re-signed by the firewall.

C. The decryption policy is configured to decrypt traffic to categories or specific URLs that use client-side certificates for authentication, which the firewall's proxy function cannot handle transparently.
D. The firewall is configured to block sessions that encounter decryption errors (e.g., unsupported cipher suites, protocol errors), rather than bypassing decryption for such sessions.
E. The Decryption policy is placed after security policies that allow encrypted traffic, preventing the decryption engine from processing the traffic before it's allowed to pass.
Answer: (SHOW ANSWER)
SSL Forward Proxy decryption introduces a 'man-in-the-middle' which requires careful consideration of various factors: - Option A (Correct): Clients must trust the firewall's root CA (Forward Trust Certificate) that is used to re-sign certificates. If this certificate isn't deployed or trusted on client devices, users will receive certificate warnings/errors in browsers and applications. This is a fundamental requirement. - Option B (Correct): Applications employing certificate pinning (e.g., some banking apps, mobile apps) are designed to prevent Man-in-the-Middle attacks by only trusting a specific server certificate. The firewall's re-signed certificate will be seen as untrusted by these applications, causing connection failures. These applications often require exclusion from decryption. - Option C (Correct): Applications using client-side certificates for authentication (where the client presents a certificate to the server) are typically incompatible with SSL Forward Proxy. The firewall intercepts the flow, but doesn't possess the user's private key to present the client certificate to the server, breaking authentication. Traffic to sites requiring client-side certificates must generally be excluded from decryption. - Option D (Correct): The Decryption profile action for 'Decryption Errors' is critical. If set to 'Block', any issue encountered during the SSL/TLS negotiation or decryption attempt (like unsupported ciphers, protocol violations, or errors) will result in the session being blocked, causing application failures. Setting it to 'No Decryption' (bypass) for errors allows the session to proceed without inspection but prevents the block. - Option E (Incorrect): Policy evaluation order is crucial, but the Decryption policy is evaluated independently from the Security policy (or concurrently in modern flows). Decryption is determined based on the Decryption policy rules and Decryption profile before the Security policy applies security inspection after the traffic state (decrypted or not) is known. A policy allowing encrypted traffic before a decryption policy wouldn't prevent decryption; rather, the flow determines if decryption applies based on decryption rules first, then the security policy is applied to the flow (whether decrypted or not). However, placing the decryption exclusion rule after an inclusion rule in the decryption policy could cause issues, but the general order of Security vs. Decryption policy evaluation is not the cause described.

NEW QUESTION: 76
An organization is deploying GlobalProtect to secure access for its remote workforce. They want to ensure users authenticate using Azure AD via SAML and that access is only granted if the user's device passes a Host Information Profile (HIP) check verifying antivirus status and disk encryption. Which components of the GlobalProtect configuration on the Palo Alto Networks

NGFW or Prisma Access are involved in implementing this secure access process? (Select all that apply)

A. GlobalProtect Portal configuration, defining authentication methods and agent configurations.

B. GlobalProtect Gateway configuration, defining the tunnel settings, authentication profiles, and HIP requirements.

C. Authentication Profile configured to integrate with Azure AD (e.g., via SAML) and an Authentication Sequence referencing this profile.

D. Host Information Profile (HIP) objects and profiles defining the required endpoint compliance criteria.

E. Security Policy rules matching the user and HIP profile to allow access to specific resources.

Answer: A,B,C,D (LEAVE A REPLY)

GlobalProtect setup involves multiple configuration points for authentication, tunnel establishment, and posture checking. - Option A (Correct): The GlobalProtect Portal is where users initially connect to obtain their agent configuration and list of available Gateways. It handles primary authentication and policy retrieval. - Option B (Correct): The GlobalProtect Gateway terminates the secure tunnel from the client. It enforces authentication (referencing Authentication Profiles), defines tunnel settings, and applies HIP requirements based on configured profiles. - Option C (Correct): Authentication Profiles and Sequences are configured to integrate with external identity providers like Azure AD using protocols like SAML, allowing the firewall/Prisma Access to authenticate users and obtain group membership. - Option D (Correct): HIP Objects define individual compliance checks (like AV status, disk encryption). HIP Profiles combine these objects to define an overall compliance state. These are configured on the firewall/Prisma Access. - Option E (Incorrect): Security Policy rules grant access after the user has successfully connected via the gateway and passed checks. The policy rule doesn't configure the GlobalProtect access process itself.

NEW QUESTION: 77

An administrator is using Panorama to manage multiple PA-Series firewalls. They have created a shared address object named 'Sensitive-Servers' that contains the IP addresses of critical internal servers. They want to use this shared object in security policy rules for different Device Groups. What is the primary benefit of using a shared address object in Panorama compared to creating the same address object locally on each managed firewall?

A. It reduces the load on individual firewalls by offloading address resolution to Panorama.
B. It ensures consistency of the 'Sensitive-Servers' definition across all firewalls that use the shared object in their policies.
C. It automatically updates the address object on the firewalls whenever the IP addresses of the sensitive servers change dynamically.
D. It allows the 'Sensitive-Servers' object to be used in NAT policies, which is not possible with local address objects.
E. It enables High Availability synchronization for the address object between managed firewall pairs.
Answer: (SHOW ANSWER)
Shared objects in Panorama are a key feature for centralized management and consistency. - Option A: Panorama is for management and logging; it doesn't participate in the real-time forwarding or address resolution performed by the firewall data plane. - Option B (Correct): The primary benefit of shared objects is ensuring uniformity. By defining 'Sensitive-Servers' once in Panorama and referencing it in policies across multiple Device Groups/firewalls, you guarantee that all firewalls using that object have the exact same definition. If the IP addresses change, you update the object once in Panorama, push the configuration, and it's consistently updated everywhere, preventing configuration drift. - Option C: Shared objects themselves don't automatically handle dynamic IP changes (unless populated by sources like EDLs referenced within the object). This benefit is about consistent configuration , not dynamic updates based on network changes. - Option D: Both shared and local address objects can be used in NAT policies. - Option E: HA synchronization happens directly between firewalls in an HA pair and synchronizes session state, NAT sessions, and policy/configuration (which includes objects), but the benefit of the shared object in Panorama is the centralized consistency of the definition before it's pushed and potentially synchronized via HA.

NEW QUESTION: 78
A company is upgrading a pair of PA-5220 firewalls in an Active/Passive HA configuration to a new PAN-OS version. They have reviewed the release notes and determined the correct upgrade path. Which is the recommended sequence of steps to perform the PAN-OS software upgrade on the HA pair to minimize downtime and disruption? (Assume the new image has been downloaded to both firewalls).
A. Upgrade the Active firewall first, then perform a failover, then upgrade the now Passive firewall.
B. Suspend the Passive firewall from the HA state, upgrade the Suspended (originally Passive) firewall, make it Active, suspend the originally Active firewall, and upgrade it.
C. Upgrade the Passive firewall first, then perform a manual failover to make it Active, then upgrade the originally Active (now Passive) firewall.
D. Upgrade both firewalls simultaneously to reduce the overall upgrade window.
E. Install the new PAN-OS version on both firewalls first, then reboot the Active firewall, wait for it to come back up, and then reboot the Passive firewall.

Answer: (SHOW ANSWER)
The standard and recommended method for upgrading an Active/Passive HA pair is to upgrade the Passive unit first to maintain redundancy during the process. - Option A: Upgrading the Active firewall first leaves the network vulnerable during the upgrade and subsequent failover, as there's no ready Passive unit. - Option B: While suspending the Passive is a valid troubleshooting step, the most common and recommended sequence for an upgrade is to start with the Passive unit. - Option C (Correct): This is the recommended sequence. Upgrade the Passive firewall first (download and install the new PAN-OS image). Once it's successfully upgraded and ready, perform a manual failover. The originally Passive unit (now running the new version) becomes Active and starts processing traffic. Then, upgrade the originally Active unit (which is now Passive). This ensures one firewall is always active and processing traffic throughout most of the upgrade process, minimizing downtime. -Option D: Upgrading simultaneously introduces significant downtime as both firewalls are unavailable. - Option E: Installing the image is separate from rebooting to run the new version. While you do install first, rebooting the Active unit before the Passive unit is upgraded and ready to take over causes an outage.

NEW QUESTION: 79
An administrator runs a BPA report on a recently deployed Palo Alto Networks VM-Series firewall in a cloud VPC. The report highlights a 'Medium' severity finding under the 'Network Settings' category titled 'Interfaces with Default Profile Settings'. What does this finding likely indicate, and what is the recommended best practice it refers to?
A. The firewall interfaces are not assigned to any Security Zone, violating the zone-based policy model.
B. The firewall interfaces are using default Link Monitoring or Path Monitoring profiles instead of custom profiles tailored to the specific network links.
C. The firewall interfaces are configured with default MTU or duplex settings that may not be optimal for the network environment.
D. The firewall interfaces are configured without User-ID or Device-ID collection enabled, limiting visibility.
E. The firewall interfaces are using default security zone names ('trust', 'untrust') instead of custom, descriptive names.
Answer: (SHOW ANSWER)
The finding 'Interfaces with Default Profile Settings', especially under Network Settings and related to profiles, typically refers to operational monitoring profiles. - Option A: Interfaces not assigned to a zone would likely trigger a different, more severe finding. - Option B (Correct): This finding usually indicates that the default Link Monitoring or Path Monitoring profiles (which have generic probe settings) are applied to WAN interfaces, instead of custom profiles where probe settings (interval, threshold, destination) are tuned for the specific characteristics of the actual links. This can lead to inaccurate link state detection or sub-optimal SD-WAN performance. The best practice is to create custom monitoring profiles. - Option C: While MTU/duplex settings are part of interface configuration, the 'Default Profile Settings' finding points to the monitoring profiles

specifically. - Option D: User-ID/Device-ID are features applied to zones/interfaces, but this finding is about profile settings, specifically monitoring profiles. - Option E: Using default zone names is a naming convention issue, not typically flagged as a 'Default Profile Settings' violation.

NEW QUESTION: 80

A. The Security policy rule allowing SSH traffic must have a WildFire analysis profile configured.
B. The SSH Proxy decryption feature must be enabled and successfully decrypting the session.
C. The SSH client and server must be configured to explicitly allow file transfers (like SCP or SFTP) on standard SSH port 22.
D. The firewall must have the root CA certificate used to sign the server's SSH host key installed as a Trusted Root CA.
E. The session must be using SSH protocol version 1, as later versions are not inspectable.
Answer: (SHOW ANSWER)
To inspect the content and activities happening inside an encrypted SSH tunnel (like file transfers or command execution which could trigger threat signatures), the firewall must be able to decrypt the tunnel. This is the function of the SSH Proxy feature. Once decrypted, App-ID can identify activities like 'file-transfer' within the SSH session, and Content-ID/Threat Prevention engines can scan the data stream for threats. Option A is necessary for detecting malware if the traffic is decrypted, but decryption is the prerequisite. Option C describes how file transfers happen over SSH but doesn't explain how the firewall sees them within the encrypted tunnel. Option D is related to validating certificates, which is part of SSL/TLS, not the host key verification process used in SSH Proxy. Option E is incorrect; SSH Proxy is designed for modern, secure SSH protocol versions (like v2); SSHv1 is deprecated and insecure, and less likely to be supported for advanced inspection.

NEW QUESTION: 81

An administrator is configuring a Threat Prevention profile on a Palo Alto Networks NGFW to leverage the Advanced Threat Prevention (ATP) CDSS. Which section within the Threat Prevention profile configuration allows the administrator to define how the firewall should react when a specific severity level of threat signature is matched (e.g., critical, high, medium, low, informational)?
A. Exclusions
B. Rule Details (or Rules tab)
C. Signatures
D. Threat Exceptions
E. Advanced
Answer: (SHOW ANSWER)
Within a Threat Prevention profile, the actions for different threat severities are configured in the 'Rules' tab or 'Rule Details' section, which defines how the firewall should respond (allow, alert, reset, block) when a signature matches at a specific severity level. Option A is for excluding specific signatures. Option C is where you might view or manage individual signatures (less

common in practice). Option D is for creating exceptions for specific threats under certain conditions. Option E contains other settings like packet capture options.

NEW QUESTION: 82
A company uses Palo Alto Networks Prisma Access for its remote workforce. They have a strict policy to prevent the exfiltration of sensitive customer data, specifically documents containing patterns resembling Social Security Numbers (SSNs) or Credit Card Numbers (CCNs). Users should be blocked if they attempt to upload such documents to cloud storage or webmail services. Assuming App-ID correctly identifies the applications and SSL Forward Proxy decryption is successfully enabled for relevant traffic, which Content-ID feature is used to enforce this policy, and what is a key aspect of its configuration?
A. Threat Prevention profile configured with signatures for SSNs and CCNs, which scans the decrypted data stream.
B. URL Filtering profile configured to block access to all cloud storage and webmail categories.
C. File Blocking profile configured to block document file types (like .doc, .pdf) being uploaded to the internet.
D. Data Filtering profile configured with specific patterns (regex or built-in) for SSNs and CCNs, applied to relevant security policy rules with an action like 'block' or
E. Antivirus profile configured to detect data patterns associated with sensitive information.
Answer: (SHOW ANSWER)
Preventing sensitive data loss based on pattern matching within application traffic is the specific function of the Data Filtering profile (part of Content-ID). Option D correctly identifies this feature and a key aspect of its configuration: defining the patterns to look for (using regular expressions or built-in data identifiers) and specifying the action (block, alert, etc.) when a match is found within the traffic flow that the Data Filtering profile is applied to via a security policy. Option A is incorrect; Threat Prevention signatures are primarily for exploits and malware, not data patterns. Option B is too blunt; it blocks access entirely rather than inspecting the content being transferred. Option C blocks file types, not specific content within files. Option E is incorrect; Antivirus profiles scan for malware signatures, not sensitive data patterns.

NEW QUESTION: 83
A company has deployed Prisma SD-WAN with ION devices at its branch offices. They need to control and secure traffic flowing not only from internal users to the internet and data center but also between internal segments within the branch itself (e.g., preventing devices on the IoT VLAN from initiating connections to the Corporate VLAN, except for specific management traffic). Which of the following are valid approaches using Prisma SD-WAN's zone-based firewall capabilities to achieve this internal segmentation and security within the branch? (Select all that apply)
A. Assign each internal segment (Corporate VLAN, IoT VLAN) to a distinct Security Zone on the ION device.
B. Create Security Policy rules with Source Zone being one internal zone and Destination Zone being another internal zone (e.g., Source Zone 'IoT', Destination Zone 'Corporate').

C. Configure the inter-zone-default security rule to 'allow' instead of 'deny' to permit all traffic between internal zones by default.
D. Apply appropriate security profiles (Threat Prevention, Antivirus, etc.) to the Security Policy rules controlling traffic between internal zones.
E. Rely solely on access control lists (ACLs) configured on the local switches to control traffic between VLANs, bypassing the ION's zone-based firewall.
Answer: (SHOW ANSWER)
Securing traffic between internal segments (east-west traffic) within a branch is a key use case for the zone-based firewall on the ION. - Option A (Correct): The foundational step is to define distinct Security Zones for each internal segment that needs to be separated and controlled. This establishes the trust boundaries. - Option B (Correct): To control traffic flow between these internal zones, you must create explicit Security Policy rules that specify the source zone and destination zone as the respective internal zones. These rules dictate what applications/services are allowed or denied between those segments. - Option C (Incorrect): The default inter-zone-default rule is 'deny'. Changing this to 'allow' would defeat the purpose of segmentation and allow all traffic between different zones by default, which is highly insecure. - Option D (Correct): For hardening, even trusted-looking internal traffic can carry threats (e.g., lateral movement of malware). Applying security profiles (Threat Prevention, Antivirus, Data Filtering, etc.) to the allow rules between internal zones provides deep inspection and protection against threats propagating laterally. - Option E (Incorrect): Relying solely on basic ACLs on switches provides only limited L3/L4 filtering and completely bypasses the App-ID, User-ID, and advanced Content-ID inspection capabilities of the ION's zone-based NGFW, which are necessary for modern security.

NEW QUESTION: 84
During the ZTP process for a Prisma SD-WAN ION device, after the device successfully connects to the cloud controller, what is the primary configuration information that the device downloads to become fully operational within the SD-WAN fabric and managed by the cloud console?
A. The full security policy set (Security, NAT, Decryption policies) defined for the site.
B. The latest PAN-OS software image.
C. Device-specific configuration including interface settings, zones, WAN link details, local subnets, and initial connectivity parameters for tunnels to other sites/services.
D. Dynamic content updates (App-ID, Threat, URL).
E. User-ID agent software.
Answer: (SHOW ANSWER)
ZTP provides the initial device-specific configuration to get the ION online and connected to the fabric. - Option A: While security policies are applied, ZTP typically provides the device-specific network configuration and the framework to receive policies. The full policy set might be pushed subsequently or inherited from templates. - Option B: Software images are downloaded and installed separately, typically before or as part of the ZTP process, but the initial download from the controller is the configuration . - Option C (Correct): The ZTP process delivers the configuration that makes the ION specific to its site: interface assignments and settings, zone

mapping, details about its WAN links (type, bandwidth, ISP), definitions of local subnets behind it, and the parameters needed to establish initial control plane connections and potentially data plane tunnels to other sites (like the controller or other IONs/hubs). - Option D: Dynamic content updates are downloaded after the core configuration and connectivity are established. - Option E: User-ID agent software is installed on domain controllers/servers, not typically on the ION device itself.

NEW QUESTION: 85
An administrator is configuring SSL Inbound Inspection on a Palo Alto Networks NGFW to decrypt incoming HTTPS traffic destined for an internal web server. Which type of certificate, specifically the private key component, must be imported onto the firewall to enable successful decryption of traffic destined for that specific server?
A. The firewall's self-signed root CA certificate.
B. The firewall's self-signed intermediate CA certificate for forward proxy.
C. The public certificate of the external client connecting to the server.
D. The server certificate of the internal web server, including its private key.
E. A wildcard certificate trusted by internal clients.
Answer: (SHOW ANSWER)
SSL Inbound Inspection requires the firewall to decrypt traffic destined for internal servers. This is achieved by having the server's private key, which allows the firewall to decrypt the symmetric session key exchanged during the SSL handshake. Option A and B are for SSL Forward Proxy. Option C is for client authentication, not server-side decryption. Option E is a type of certificate that might be used, but specifically the server's private key associated with the server certificate is required.

NEW QUESTION: 86
A company is using Prisma SASE (Prisma Access) with Enterprise DLP and SaaS Security features. They want to monitor for accidental or malicious sharing of confidential documents (identified by content signatures or keywords) within sanctioned SaaS applications like Microsoft SharePoint Online and Slack. Access to these applications is over HTTPS. What capabilities and configurations are necessary to achieve this monitoring and enforcement within encrypted sanctioned SaaS application traffic? (Select all that apply)
A. SSL Forward Proxy decryption policy configured to decrypt traffic to the domains used by sanctioned SaaS applications (SharePoint, Slack).
B. Data Filtering profiles configured with patterns (e.g., keywords, regex, content identifiers) that define the confidential documents.
C. Security Policy rules matching user traffic to the sanctioned SaaS applications (identified by App-ID), with the Data Filtering profile applied.
D. App-ID successfully identifying specific application functions like 'sharepoint-upload', 'slack-post', or 'slack-file-upload' within the encrypted traffic.

E. WildFire analysis profile configured to scan document files uploaded or shared within the SaaS applications.

Answer: (SHOW ANSWER)

Monitoring sensitive data sharing within encrypted SaaS apps requires decryption, defining data patterns, identifying application actions, and applying policy. - Option A (Correct): Decryption is essential to see the content and specific actions within encrypted SaaS traffic. - Option B (Correct): Data Filtering profiles are used to define what constitutes 'confidential documents' based on content. - Option C (Correct): Security Policy rules specify where the inspection happens. Rules matching the sanctioned SaaS applications and applying the Data Filtering profile ensure content inspection is performed on traffic to/from those apps. - Option D (Correct): App-ID's ability to identify specific application functions (like uploading or posting files/messages) is necessary for creating granular DLP policies (e.g., alert if confidential data is uploaded to SharePoint but maybe just log if it's viewed). - Option E: WildFire scans for malware within files, not sensitive data content. While scanning uploaded files for malware is important, it's not the mechanism for detecting sensitive data patterns.

NEW QUESTION: 87

A large enterprise is modernizing its infrastructure, which includes a traditional on-premises data center, a significant presence in a public cloud (AWS/Azure/GCP), and a growing adoption of Kubernetes for containerized applications. The security architecture mandates next- generation firewall capabilities (App-ID, Content-ID, user/device awareness) at key security inspection points. Match the following Palo Alto Networks NGFW form factors to their MOST appropriate primary deployment scenarios or use cases in this hybrid environment: I. PA-Series II. VM-Series III. CN-Series IV. Cloud NGFW for AWS/Azure Palo Alto Networks security use cases: P. High-performance physical appliance for data center perimeter or core segmentation. Q. Software-based firewall for virtualized environments, private clouds, or public cloud IaaS perimeter/segmentation. R. Kubernetes-native firewall for securing inter-service communication and cluster ingress/egress traffic. S. Managed cloud-native firewall service for protecting public cloud workloads with simplified operations.

A. I-P, II-Q, III-R, IVS

B. I-Q, II-P, III-s, IV-R

C. I-S, II-R, III-Q, IV-P

D. I-P, II-s, III-R, IV-Q

E. I-Q, II-R, III-P, IV-S

Answer: (SHOW ANSWER)

Understanding where each Palo Alto Networks NGFW form factor is best suited is key to designing a comprehensive security architecture. - I. PA-Series (Physical Appliances): These are hardware-based firewalls designed for high throughput and performance, typically deployed at physical perimeters (internet edge) or for high-density segmentation within physical data centers (P). - II. VM-Series (Virtual Appliances): These are software versions running on hypervisors (VMware, KVM, Hyper-V) or in public cloud IaaS environments (AWS EC2, Azure VM, GCP

Compute Engine). They provide flexibility and can be used for virtual data center segmentation, private cloud security, or securing public cloud IaaS environments (Q). - III. CN-Series (Containerized NGFW): Designed specifically for Kubernetes and container environments. They run as containerized workloads and provide security for traffic within the cluster (east-west) and in/out of the cluster (north-south) (R). - IV. Cloud NGFW for AWS/Azure: This is a fully managed cloud-native firewall service offered directly within the public cloud provider's console (AWS Network Firewall integration, Azure Virtual Hub). It provides NGFW capabilities with simplified deployment and management, ideal for protecting public cloud workloads and VPCNNet perimeters (S). Option A correctly matches each form factor to its primary use case.

NEW QUESTION: 88
How does Cortex XSIAM enhance proactive security operations?
Response:
A. By eliminating the need for EDR solutions
B. By enabling AI-powered threat hunting and anomaly detection
C. By automatically blocking all external network traffic
D. By focusing only on known attack signatures
Answer: (SHOW ANSWER)

NEW QUESTION: 89
A security administrator is investigating a potential malware outbreak on the internal network protected by a Palo Alto Networks PA-Series firewall. They need to identify which users are accessing specific malicious URLs or downloading suspicious files. Which log types generated by the firewall are MOST relevant for this investigation, providing visibility into user activity, applications, and detected threats? (Select all that apply)
A. Traffic logs
B. Threat logs
C. URL Filtering logs
D. Configuration logs
E. System logs
Answer: (SHOW ANSWER)
Investigating user activity, application usage, and detected threats relies on specific firewall log types: - Option A (Correct): Traffic logs record details about every session flowing through the firewall that matches a logging-enabled security policy rule. They include source/destination IP/port, zones, application ID, user ID, action (allow/deny/drop), and session duration. This is fundamental for seeing who accessed what application. - Option B (Correct): Threat logs record all detected security threats, including malware, exploits, spyware, and command-and-control activity, based on the applied Threat Prevention, Antivirus, and WildFire profiles. These logs directly indicate malicious activity. - Option C (Correct): URL Filtering logs record details about URL access attempts, including the requested URL, the URL category, the configured action (allow/block/alert), the source user, and the destination IP. This is essential for tracking user

access to specific websites, including known malicious ones. - Option D (Incorrect): Configuration logs track changes made to the firewall's configuration, which is not relevant for investigating traffic-related security incidents. - Option E (Incorrect): System logs record events related to the firewall's operation (e.g., interface status changes, daemon restarts, resource utilization) but not the details of user traffic or detected threats within those flows.

NEW QUESTION: 90

A. Threat Prevention, Antivirus, URL Filtering, File Blocking, Data Filtering
B. Threat Prevention, Antivirus, WildFire Analysis, URL Filtering, File Blocking, Data Filtering
C. Threat Prevention, Antivirus, WildFire Analysis, URL Filtering, File Blocking
D. Threat Prevention, Antivirus, URL Filtering, Data Filtering
E. Only Threat Prevention, Antivirus, and URL Filtering are included by default.
Answer: (SHOW ANSWER)
The standard set of Content-ID security profiles that can be bundled into a Security Profile Group includes all the major inspection engines: Threat Prevention, Antivirus, WildFire Analysis, URL Filtering, File Blocking, and Data Filtering. Option B lists all these profiles.

NEW QUESTION: 91

After successfully downloading and installing a new version of a dynamic update (e.g., App-ID or Threat Prevention) on a Palo Alto Networks NGFW or Prisma Access node, when does the firewall start using the new definitions or signatures?
A. Immediately after the download completes.
B. After the firewall is rebooted.
C. After a configuration commit is performed.
D. Immediately after the installation completes, typically without requiring a reboot or commit for most dynamic updates.
E. After the associated Security Policy rule is modified and committed.
Answer: (SHOW ANSWER)
Dynamic updates are designed to be applied frequently and without disruption. Unlike PAN-OS software upgrades, dynamic updates (App-ID, Threat, URL, WildFire) are typically loaded into the firewall's memory and activated shortly after installation, without requiring a reboot or a configuration commit. This ensures the firewall is using the latest intelligence as quickly as possible. Option A is incorrect; there's an installation step after download. Options B and C describe actions for software upgrades or configuration changes, not dynamic updates. Option E is incorrect; applying updates doesn't require modifying the policy rule itself (unless you want to leverage a new feature enabled by the update, like a new application function).

NEW QUESTION: 92
A security administrator is investigating a user who is suspected of attempting to download malware and access restricted websites using encrypted channels. The Palo Alto Networks NGFW (or Prisma Access) is configured with SSL Forward Proxy decryption, URL Filtering, Antivirus, and WildFire Analysis profiles applied to the relevant security policy rules. Which log types should the administrator examine in Cortex Data Lake or Panorama to gain comprehensive insight into this user's activity and any detected security events?
(Select all that apply)
A. Traffic logs, to see which sessions were allowed or denied, the applications used, and identify sessions related to the user.
B. Decryption logs, to confirm whether SSL decryption was attempted and successful for the user's encrypted traffic.
C. URL Filtering logs, to see which websites the user attempted to access and the categories/actions associated with those sites.
D. Threat logs, to see if any malware, exploit, or other threats were detected within the user's traffic or files.
E. File logs, to see if any files were transferred, their type, and the outcome of Antivirus or WildFire analysis.
Answer: (SHOW ANSWER)
Investigating activity and detected threats over encrypted channels requires looking at multiple interconnected log types: - Option A (Correct): Traffic logs are the starting point, providing the session context (who, what, where, when, allowed/denied). - Option B (Correct): Since the investigation involves encrypted channels, checking Decryption logs is crucial to confirm if decryption was attempted and successful. Decryption logs show status, errors, and policies applied. - Option C (Correct): URL Filtering logs specifically track web access attempts, showing the URLs visited and the policy action (block/allow) based on category or threat feeds. - Option D (Correct): Threat logs record detections from Threat Prevention, Antivirus, and WildFire, directly indicating if malware, exploits, or other threats were found in the traffic payload. - Option E (Correct): File logs provide details about file transfers detected within sessions, including the file type, direction, size, and the results of Antivirus and WildFire scanning for that specific file. This is essential for confirming malware downloads.

NEW QUESTION: 93
When remote users connect to Prisma Access via GlobalProtect, their traffic is directed through the cloud security platform. Which security zone is typically used to represent the source of traffic originating from these connected mobile users in Security Policy rules?

A. The zone assigned to the GlobalProtect Gateway interface.
B. The zone configured for the 'Remote Networks' in Prisma Access.
C. The zone representing the public internet (e.g., 'Public' or 'Internet').
D. A dedicated 'Mobile-Users' zone in Prisma Access.
E. The zone assigned to the user's home network interface.
Answer: (SHOW ANSWER)
Prisma Access assigns traffic from mobile users connecting via GlobalProtect to a specific, dedicated zone for policy enforcement purposes. Option A refers to a zone on a self-managed firewall. Option B is for site-to-site VPNs. Option C is for the destination zone for internet traffic. Option E is the user's local physical interface, not relevant to the traffic flow through Prisma Access. Prisma Access uses the 'Mobile-Users' zone to logically segment traffic originating from connected remote users.

NEW QUESTION: 94
A network administrator is configuring a security policy rule on a Palo Alto Networks Strata NGFW for internal user access to a critical server farm zone. The policy should permit access to specific applications only for authenticated users who belong to certain Active Directory groups. The rule configuration uses User-ID in the 'Source User' field. What happens when a user whose IP address is not currently mapped to a username by User-ID attempts to match this security policy rule?
A. The traffic will automatically be matched by a default 'allow any any' rule hidden from view.
B. The traffic will be explicitly denied because the 'Source User' field is specified, and no matching user is found, effectively defaulting to a deny for unauthenticated/unknown users.
C. The firewall will initiate a Captive Portal authentication request to the user to obtain a mapping before evaluating the policy further.
D. The policy lookup will ignore the 'Source User' field and match the rule based solely on Source Zone, Destination Zone, and Application.
E. The traffic will be processed by the data plane using hardware acceleration without hitting the slow path because identity is not yet determined.
Answer: (SHOW ANSWER)
When a security policy rule includes a 'Source User' (or 'Destination User') criterion, and the firewall does not have a user mapping for the IP address in question, the firewall cannot evaluate the rule based on identity. In Palo Alto Networks policy logic, if a criterion is specified in a rule (like a specific user or group), and the necessary information to evaluate that criterion is missing (like the user mapping), that rule cannot be matched by the traffic. This effectively means that for rules leveraging User-ID/Device-ID, traffic from IPs without the required mapping will not match rules that require that mapping. If there is no broader rule (like an 'any' user rule) below it that allows the traffic, the traffic will eventually hit the default deny policy. By specifying a user/group, you are essentially saying 'only allow these identified users/groups'. Traffic from unknown users will not match this rule and will proceed down the policy list, likely hitting an implicit or explicit deny. Option B is the most accurate description of the typical outcome, as the rule requires a user

identity match that isn't present. Option A is incorrect; there isn't a hidden default allow. Option C would only happen if a separate authentication policy rule or Captive Portal configuration was triggered based on zone or other criteria, not automatically because a security rule with a user field wasn't matched. Option D is incorrect; the firewall does attempt to evaluate all specified criteria. Option E is incorrect; initial session setup and policy lookup occur on the slow path, and identity lookup is part of that process.

NEW QUESTION: 95
A branch office is configured with a Prisma SD-WAN ION device and has two internet links: a primary broadband connection and a secondary LTE link. The organization prioritizes VoIP traffic for business continuity and needs to ensure it uses the best available path based on real-time quality metrics, falling over to the LTE link if the broadband link deteriorates. Which type of Prisma SD-WAN policy is primarily used to define this behavior for VoIP traffic?
A. Security Policy
B. NAT Policy
C. Qos Policy
D. Path Policy
E. Application Override Policy
Answer: (SHOW ANSWER)
Prisma SD-WAN uses different policy types for different functions. Path Policy is specifically designed for dictating how traffic is steered over the available WAN links based on applications, link quality, and business intent. Option A (Security Policy) controls what traffic is allowed/denied and inspected. Option B (NAT Policy) handles address translation. Option C (QOS Policy) prioritizes traffic on a link but doesn't dictate which link to use for a given application flow in the context of SD-WAN path selection. Option E (Application Override) reclassifies traffic but doesn't handle path selection.

NEW QUESTION: 96
An organization is using Palo Alto Networks NGFWs with Enterprise DLP to prevent sensitive data exfiltration. A user attempts to upload a file containing credit card numbers to a cloud storage service via HTTPS. Assuming a Data Filtering profile is configured to detect credit card numbers and the Security Policy rule allows this traffic, what critical step must be successfully completed by the firewall for the Data Filtering inspection to occur and the DLP policy to be enforced on this encrypted traffic?
A. App-ID must identify the traffic as 'web-browsing' or the specific cloud storage application.
B. The firewall must perform SSL Forward Proxy decryption on the HTTPS session.
C. The file type must be allowed by the File Blocking profile.
D. User-ID must identify the user performing the upload.
E. The destination URL must be categorized as 'Cloud Storage' by URL Filtering.
Answer: (SHOW ANSWER)

Data Loss Prevention (DLP) and Data Filtering inspect the content of the traffic stream. If the traffic is encrypted (like HTTPS), the content is not visible to the firewall unless it is decrypted. Option A, C, D, and E are important for policy matching or other security functions, but decryption is the prerequisite for inspecting the sensitive data within the encrypted payload. SSL Forward Proxy decryption is used for outbound encrypted traffic like uploads to cloud storage.

NEW QUESTION: 97
An organization has deployed Palo Alto Networks IoT Security and integrated it with their Strata NGFW. The IoT Security platform has identified a group of 'Smart Thermostats' on the network segment. The security team wants to create a policy on the NGFW to allow these devices to communicate only with their vendor's cloud update server on HTTPS (port 443) and block all other outbound communication. Which type of security policy rule criteria is specifically enabled by the IoT Security integration to represent the group of discovered thermostats?
A. A static Address Group containing the known IP addresses of the thermostats.
B. A dynamic Address Group based on the 'Smart Thermostats' device category provided by the IoT Security subscription.
C. A User-ID mapping for the thermostats to an IoT user group.
D. A custom Application signature for the thermostat's communication protocol.
E. A URL Category created for the vendor's update server domain.
Answer: (SHOW ANSWER)
The IoT Security integration provides dynamic device groups based on the discovered and profiled device inventory. Option A is manual and not dynamic as devices change. Option B correctly identifies the dynamic Address Group concept: the IoT Security cloud service maintains the group membership based on its profiling, and this group object is available for use in NGFW security policies. Option C is incorrect; User-ID is for human users. Option D might identify the application, but not the specific group of devices . Option E identifies the destination, but not the source devices.

NEW QUESTION: 98
A company is deploying a new internal application that uses a standard web server (HTTPS on port 443) but needs specific security policy enforcement (different from general web browsing) and precise visibility into its usage. App-ID currently identifies this traffic as 'web-browsing'. How can an administrator configure the Palo Alto Networks NGFW (Strata/Prisma SASE) to identify this internal application separately and enable granular policy control?
A. Create a custom Service object for port 443 and use it in the Security policy rule instead of the default 'service-https'.
B. Define a custom App-ID signature based on unique characteristics of the application's traffic (e.g., specific HTTP headers, URL patterns), and use this custom App-ID in Security Policy rules.
C. Modify the default 'web-browsing' App-ID signature to exclude traffic to the internal application's IP address.

D. Use a URL Filtering profile to categorize the internal application's URL and apply policy based on that category.
E. Enable SSL Inbound Inspection for the internal application server and rely on Content-ID to differentiate the traffic.
Answer: (SHOW ANSWER)
When App-ID doesn't recognize a custom or specific application, the correct approach for granular identification and policy is to create a custom App-ID signature. Option B correctly describes this process: analyzing the application's traffic for unique patterns and building a custom signature that App-ID can use to identify it separately. Option A uses ports, which is not application-aware. Option C is not possible; built-in App-IDs cannot be directly modified. Option D is for URL categorization, not application identification. Option E is for inspecting content after identification, but doesn't help with the initial App-ID challenge.

NEW QUESTION: 99
A company is using Prisma Access for remote users and wants to enforce a policy where access to file-sharing applications (like Dropbox, Google Drive upload) is restricted to specific user groups, regardless of whether the destination is a sanctioned corporate account or a personal account. All other standard internet browsing should be allowed for everyone. How would this policy be implemented using Prisma Access Security and App-ID?
A. Configure a Security Policy rule with 'Source User' set to the allowed user group, 'Destination Zone' as 'Public', 'Application' set to the file-sharing App-IDs, and 'Action' as 'allow'. Place this rule above a more general 'allow' rule for other web browsing.
B. D Configure a Security Policy rule with 'Source User' set to the user groups that should not have access, 'Destination Zone' as 'Public', 'Application' set to the file- sharing App-IDs, and 'Action' as 'deny'. Place this rule above a general 'allow' rule.
C. Use URL Filtering to block the category 'File Sharing and Storage' for all users except the allowed group.
D. Configure a NAT policy rule to block traffic destined for file-sharing service IPs.
E. Create a custom application signature for file-sharing applications based on port and protocol.
Answer: A,B (LEAVE A REPLY)
Controlling application access based on user identity is a core function of User-ID integrated with Security Policy and App-ID. - Option A (Correct): This is one valid approach. You define an explicit 'allow' rule specifically for the authorized user group, matching the file- sharing App-IDs (like 'dropbox-upload', 'google-drive-upload), and place this rule higher in the policy list. A subsequent, broader rule would allow general internet browsing (e.g., 'web-browsing') for a wider user group (or 'any' user). - Option B (Correct): This is the alternative, equally valid approach often preferred for restricting access. You define an explicit 'deny' rule matching the user groups who should not have access to the file- sharing App-IDs. Placing this deny rule above the general 'allow' rule ensures that prohibited users are blocked before the general browsing rule permits the traffic. Both A and B achieve the desired outcome by using App-ID and User-ID in explicit policy rules placed strategically. - Option C: URL Filtering operates on URL categories. While 'File

Sharing and Storage' is a category, App-ID provides more granular control over the specific application activity (e.g., upload vs. download, authentication). Using App-ID is generally more precise for this type of control. Also, managing exceptions for a group via URL filtering alone can be less straightforward than using user groups in security policy. - Option D: NAT policy handles address translation, not access control based on applications or users. - Option E: App-ID automatically identifies many common file- sharing applications based on more than just port/protocol, making custom signatures usually unnecessary unless dealing with a very uncommon or internal application.

NEW QUESTION: 100

A. App-ID, which identifies applications independent of port, allowing policies to be based on application identity rather than network location.

B. User-ID and Device-ID, which integrate user and device context directly into security policy rules, enabling identity-based access control.

C. Content-ID (Threat Prevention, WildFire, URL Filtering, Data Filtering, File Blocking), which provides deep inspection of allowed traffic for threats and data exfiltration.

D. SSL Decryption (Forward Proxy, Inbound Inspection), which enables visibility into encrypted traffic to apply App-ID and Content-I

E. Security Zones configured for network segmentation based on IP subnets or VLANs.

Answer: (SHOW ANSWER)

Zero Trust moves away from implicit trust based on network location. Palo Alto Networks features enable explicit verification and deep inspection: - Option A (Correct): App-ID allows policies to be based on what the traffic is (the application), verifying the application identity explicitly, moving beyond port-based trust. - Option B (Correct): User-ID and Device-ID verify who is initiating the traffic and what device they are using, allowing policies to be tied directly to user and device identity and posture, a core tenet of explicit verification. - Option C (Correct): Content-ID features embody the 'Assume Breach' principle by inspecting all relevant allowed traffic (not just perceived threats) for malware, exploits, sensitive data, and malicious URLs. This assumes threats can exist within legitimate applications. - Option D (Correct): SSL Decryption is critical because a vast majority of modern threats and data exfiltration attempts occur over encrypted channels. Decryption is necessary to apply App-ID (more accurately) and Content-ID to encrypted traffic, enabling the 'Verify Explicitly' and 'Assume Breach' principles for this traffic. - Option E (Incorrect): While security zones are fundamental for network segmentation and policy structure, they primarily align with a segment- based approach, which is a building block, but less directly representative of the identity-aware, application-aware, content-inspecting principles at the core of modern Zero Trust compared to the other options.

NEW QUESTION: 101

In a scenario where a company wants to allow specific users to access a public SaaS application ('engineering-portal' App-ID) but restrict their access to sensitive functions within that application (e.g., blocking the 'engineering-portal-admin' function), which feature is used in the Security

Policy rule, in conjunction with the base App-ID, to enforce this granular control over application activities?

A. URL Filtering profile with custom URL lists.

B. Data Filtering profile with sensitive data patterns.

C. Application Filters.

D. Service Objects (ports and protocols).

E. Application Function Control within the Security Policy rule's Application tab.

Answer: (SHOW ANSWER)

Palo Alto Networks App-ID often identifies not just the base application but also specific functions within it. The ability to control these functions is built into the Security Policy. - Option A: URL Filtering controls access based on URLs, not specific application functions. - Option B: Data Filtering inspects content. - Option C: Application Filters are for grouping applications, not controlling functions within them. - Option D: Service Objects are port-based and cannot distinguish specific functions within a complex application. - Option E (Correct): Application Function Control (sometimes shown as checkboxes or explicit functions within the Application tab of a Security Policy rule, depending on the App-ID) allows administrators to select which specific functions of an identified application are allowed or denied, providing granular control over application usage.

NEW QUESTION: 102

An organization is using Device-ID and potentially the IoT Security subscription to gain visibility into the diverse endpoints on their network. A security policy needs to allow specific types of devices (e.g., 'Corporate Printers', 'Approved IP Cameras') to access certain network resources while restricting 'Unknown Devices' or 'Personal Devices' from accessing sensitive segments. Which of the following are valid ways to leverage Device-ID and related features in Security Policy rules on a Palo Alto Networks NGFW? (Select all that apply)

A. Using Device-ID categories directly in the 'Source' or 'Destination' tabs of a Security Policy rule (e.g., Source 'Device Category: Corporate Printers').

B. Creating dynamic Address Groups based on Device-ID categories and using these Address Groups in the 'Source Address' or 'Destination Address' fields of a Security Policy rule.

C. Creating HIP Objects that match Device-ID categories and using these HIP Objects in the 'Source User' or 'HIP Profile' tab of a Security Policy rule.

D. Applying different security profiles (Threat, URL, etc.) based on the Device-ID category identified for a session, within the same Security Policy rule.

E. Configuring Authentication Policy rules that require users on specific Device-ID categories to authenticate.

Answer: (SHOW ANSWER)

Device-ID provides identity context about the endpoint, which can be used in various policy types. - Option A (Correct): Device-ID categories (like 'Corporate Printers', 'Unknown Device') are available as direct matching criteria in the 'Source' and 'Destination' tabs of Security Policy rules. - Option B (Correct): Dynamic Address Groups can be created based on Device-ID categories.

These groups automatically include the IP addresses of devices matching the category and can be used in the address fields of Security Policy rules. - Option C (Correct): HIP Objects can be defined to match specific Device-ID categories. These HIP Objects can then be combined into HIP Profiles and used in the 'Source User' or 'HIP Profile' tab of Security Policy rules, often in conjunction with User-ID, to enforce policies based on both user and device type/posture. - Option D (Incorrect): While you apply security profiles to a rule, the specific profiles applied depend on the policy rule matched not dynamically on the Device-ID category within a single rule match. You would use separate rules for different Device-ID categories, each with its own set of security profiles. - Option E (Correct): Authentication Policy rules can be configured to require authentication (e.g., via Captive Portal) for traffic originating from devices matching specific Device-ID categories, providing identity awareness for devices where User-ID agents might not be applicable.

NEW QUESTION: 103
From a customer's perspective, which aspect of managing security posture and feature availability in Prisma Access is directly influenced by the underlying software version running on the security processing nodes?
A. The specific signature content in dynamic updates (Threat, App-ID).
B. The available security features, policy options, and supported protocols.
C. The performance capacity (throughput, sessions/second) of the assigned bandwidth.
D. The geographic location of the service connection to the data center.
E. The number of users concurrently connected via GlobalProtect.
Answer: (SHOW ANSWER)
The software version determines the fundamental capabilities of the platform. - Option A: Dynamic updates provide the latest intelligence but the types of signatures and updates available are determined by the software version. - Option B (Correct): Just like with PAN-OS on self-managed firewalls, major software version upgrades in Prisma Access unlock new features, introduce new policy options, add support for new protocols or decryption standards, and may include performance optimizations or bug fixes to existing features. The software version dictates the capabilities available to the customer. - Option C: Performance capacity is primarily determined by the allocated bandwidth and the underlying hardware/virtual resources provisioned by Palo Alto Networks, not the software version itself. - Option D: Geographic location is a deployment choice. - Option E: The number of users is a factor managed by licensing and bandwidth allocation, not directly by the underlying software version itself.

NEW QUESTION: 104
A company uses Prisma Access for Remote Networks (branch offices). They have configured a Service Connection back to their corporate data center where internal applications reside on a private IP subnet (10.50.1.0/24). Branch office users (on subnet 10.10.10.0/24) need to access these internal applications. Internet-bound traffic from the branch needs to be Source NAT'd to a public IP range assigned to the Prisma Access Remote Network location. Traffic destined for the

data center should not be Source NAT'd. Which NAT policy configurations in Prisma Access are necessary to achieve this? (Select all that apply)
A. ANAT policy rule with Original Packet: Source Zone 'Remote-Networks', Destination Zone 'Public', Translated Packet: Source Address Translation 'Dynamic IP and Port' using the Remote Network's public 12
B. ANAT policy rule with Original Packet: Source Zone 'Remote-Networks', Destination Zone 'Service-Connection' (or zone for the data center), Translated Packet: Source Address Translation 'No NAT'.
C. ANAT policy rule with Original Packet: Source Zone 'Service-Connection', Destination Zone 'Remote-Networks', Translated Packet: Destination Address Translation 'Static IP' to the branch subnet.
D. ANAT policy rule with Original Packet: Source IP 10.10.10.0/24, Destination IP 172.16.1.0/24, Translated Packet: Source Address Translation 'Dynamic IP and port'.
E. Security Policy rules must define the NAT translation needed for each traffic flow.
Answer: (SHOW ANSWER)
NAT policy in Prisma Access, like on Strata NGFWs, handles address translation based on defined rules. The rules match traffic flow (source/destination zone, etc.) and specify the translation action. - Option A (Correct): This rule matches traffic originating from the 'Remote-Networks' zone (the branch offices) destined for the 'Public' zone (the internet). It configures Source NAT using the public IP assigned to the specific Remote Network location in Prisma Access (Dynamic IP and Port is common for outbound user traffic). - Option B (Correct): This rule matches traffic originating from the 'Remote-Networks' zone destined for the 'Service-Connection' zone (representing the data center). By setting the Translated Packet Source Address Translation to 'No NAT', you explicitly tell Prisma Access not to perform SNAT on this internal-bound traffic. This ensures the original private source IPs from the branch are preserved when accessing data center resources, which is typically desired. - Option C: This describes DNAT for traffic originating from the data center towards the branch, which is not the scenario described. - Option D: While you could potentially match based on IP subnets instead of zones, using zones is the standard and recommended approach for policy definition in Palo Alto Networks platforms. More importantly, the desired action for data center traffic is 'No NAT', not Dynamic SNAT. - Option E: Security Policy rules control allow/deny and inspection profiles, but they do not define NAT translations. NAT is configured in a separate NAT Policy.

NEW QUESTION: 105
When reviewing logs and monitoring data in the Prisma SD-WAN Cloud Management Console, what is the significance of the 'Application Health Score' metric?
A. It represents the number of active sessions for a specific application.
B. It is a metric based on the application's performance relative to its defined SLA thresholds or expected quality characteristics (latency, jitter, loss).
C. It indicates the security risk level associated with the application, based on detected threats.
D. It measures the total bandwidth consumed by the application over a given period.

E. It shows the percentage of users accessing the application from a specific branch.

Answer: (SHOW ANSWER)

Application Health Score is a key metric in SD-WAN monitoring, reflecting user experience for specific applications. Option A is session count. Option C relates to security risk (though performance issues can indicate a potential security problem). Option D is bandwidth. Option E is user distribution. The Application Health Score is a composite metric derived from the underlying network performance metrics (latency, jitter, loss) compared to the application's requirements or defined SLA. A high score indicates good performance relative to needs, while a low score indicates poor performance likely impacting user experience.