Palo Alto Networks
CyberSec-Apprentice
Palo Alto Cybersecurity Apprentice

Latest Version: 6.0

Question: 1

A global enterprise has recently implemented a Palo Alto Networks Cortex XSOAR platform to enhance their security operations. They are experiencing a surge in low-level phishing attempts, often involving malicious URLs that bypass initial email gateway filters but are detected by endpoint protection. The security team wants to automate the process of analyzing these URLs, blocking them on the perimeter firewall, and notifying affected users, all while maintaining a comprehensive audit trail. Which of the following XSOAR playbooks components and integrations would be MOST critical to efficiently and securely handle this specific incident type, ensuring rapid containment and accurate reporting?

A. A custom playbook task that directly executes a
```
curl -X POST -H 'Content-Type: application/json' -d '{"url":"malicious_url"}' https://malicious-url-api.com/block
```
command against an external threat intelligence platform to block the URL, followed by an email notification.
B. An XSOAR incident type configured for 'Phishing', a pre-processing rule to extract URLs from incident fields, an 'URL Enrichment' sub-playbook utilizing integrations with Palo Alto Networks WildFire and AutoFocus, and a 'PAN-OS' integration to update URL filtering profiles on the firewall. Email notification via 'Mail Sender' integration.
C. A simple 'Block URL' script that takes the URL as input and directly pushes a configuration change to the firewall using SSH, and a separate script to generate a CSV report of blocked IJRLs for manual distribution.
D. Leveraging only the SIEM (e.g., Splunk Enterprise Security) to create a correlation search for phishing indicators, then manually creating a ticket in a helpdesk system for firewall rule updates and user notification.
E. A basic playbook that identifies phishing emails, then prompts a security analyst for manual approval before initiating any blocking actions or user notifications

Answer: B

Explanation:
Option B provides the most robust and automated solution leveraging XSOAR's strengths. An XSOAR incident type standardizes handling. Pre-processing rules automate data extraction. URL Enrichment with WildFire and AutoFocus provides critical real-time threat intelligence and context, which is paramount for accurate blocking decisions. The PAN-OS integration is the correct and secure way to programmatically update firewall policies, ensuring proper API usage and auditability. The Mail Sender integration handles user notification. Option A is insecure (hardcoding API calls, potentially lacking proper authentication/authorization and audit trail within XSOAR). Option C is less robust, lacks central management, and SSH-based configuration changes are generally less preferred for automation than dedicated API integrations. Option D is a SIEM-only approach, lacking the orchestration and automation capabilities of XSOAR for active response. Option E introduces unnecessary manual intervention for a high-volume, low-risk incident type, defeating the purpose of automation.

Question: 2

Consider a large financial institution utilizing both a SIEM (e.g., IBM QRadar) for log aggregation and correlation, and a SOAR platform (e.g., Swimlane) for incident response automation. An advanced persistent threat (APT) group has successfully gained initial access to a workstation via a novel zero-day exploit, establishing a persistent backdoor and attempting lateral movement. The SIEM detects a highly unusual outbound connection from the compromised workstation to a known C2 server (identified through a newly updated threat intelligence feed). Describe the optimal interaction and data flow between the SIEM and SOAR platforms to rapidly detect, contain, and analyze this sophisticated threat. Which of the following options accurately represents this optimal interplay?

A. The SIEM's correlation rule triggers an alert, which is then manually reviewed by an analyst. The analyst then manually logs into the SOAR platform and initiates a generic 'containment' playbook, providing the compromised IP address. All further investigation is then conducted within the SOAR platform, independent of the SIEM.
B. The SIEM's detection of the C2 connection immediately triggers a custom SOAR playbook via API. This playbook automatically quarantines the affected workstation using EDR integration, blocks the C2 IP on network firewalls, initiates forensic image acquisition, enriches context from ADICMDB, and creates a high- priority incident ticket in the SIEM, pushing all collected artifacts back to the SIEM for centralized logging and long-term analysis. The SOAR then orchestrates a post-incident review workflow.
C. The SOAR platform continuously queries the SIEM for all raw logs, performing its own correlation and anomaly detection. Upon detecting the C2 connection, the SOAR independently executes blocking actions without informing the SIEM, and stores all incident details within its own internal database, treating the SIEM purely as a log source.
D. The SIEM performs the initial detection and containment by automatically disabling the network interface of the compromised host and blocking the C2 IP. It then sends an informational alert to the SOAR, which primarily serves as a reporting tool to generate compliance reports based on the SIEM's actions.
E. Neither the SIEM nor the SOAR can effectively handle a zero-day exploit. The optimal approach is to rely on threat intelligence feeds directly integrated with network IPS/IDS devices for real-time blocking, and then manually investigate any residual alerts.

Answer: B

Explanation:
Option B describes the ideal, tightly integrated, and automated workflow for handling a sophisticated threat. The SIEM excels at large-scale log aggregation and initial correlation, making it perfect for detecting the C2 connection. The critical step is for the SIEM to immediately trigger the SOAR via API/webhook. The SOAR then takes over orchestration: automatic containment (quarantine, firewall blocking), rich context enrichment from other systems (AD, CMDB), automated forensic data collection, and crucially, pushing all collected artifacts and incident status back to the SIEM. This ensures the SIEM remains the central repository for all security data, enabling long-term analysis, historical correlation, and regulatory compliance, while the SOAR handles the rapid, automated response and workflow management. Option A is too manual, negating SOAR's automation benefits. Option C bypasses the SIEM's core function as a centralized security data lake and correlation engine. Option D underutilizes the SOAR's orchestration capabilities and puts too much response logic solely within the SIEM, which is

not its primary strength for complex workflows. Option E ignores the value of both SIEM for detection and SOAR for response against even novel threats, especially when threat intelligence is updated.

Question: 3

A cybersecurity apprentice is tasked with optimizing incident response workflows using a combination of a cloud-native SIEM (e.g., Azure Sentinel) and a SOAR platform. They encounter a recurring challenge where a specific type of 'Brute Force Attack' alert from the SIEM frequently triggers a SOAR playbook that attempts to block the attacker's IP address on the firewall. However, a significant percentage of these IPs are legitimate but misconfigured internal systems or VPN clients, leading to false positives and service disruption. The apprentice needs to refine the SOAR playbook to be more intelligent. Which of the following modifications to the SOAR playbook logic would provide the most effective and safe resolution to this issue, ensuring minimal false positives while still effectively responding to true threats?

A. Add a new step at the beginning of the playbook to unconditionally block any IP address associated with a 'Brute Force Attack' alert for 24 hours, then automatically unblock it, relying on the SIEM's initial detection to be perfectly accurate.
B. Before blocking, the playbook should integrate with the organization's Configuration Management Database (CMDB) and Active Directory (AD). If the source IP is found to belong to an internal asset or a known VPN range, the playbook should escalate to a human analyst for review, othemise, proceed with automated blocking. The playbook should also enrich the IP with geographical data from a threat intelligence feed.
C. Modify the SIEM's correlation rule to only trigger a 'Brute Force Attack' alert if the source IP is from a foreign country, completely ignoring internal or VPN IPs to reduce alert volume.
D. Implement a simple counter. If the same IP address triggers more than 10 'Brute Force Attack' alerts within 5 minutes, then and only then proceed with automated blocking, otherwise, ignore the alert.
E. Remove the automated blocking step from the SOAR playbook entirely and rely solely on manual review and intervention for all 'Brute Force Attack' alerts to avoid any service disruptions.

Answer: B

Explanation:
Option B represents the most intelligent and safest approach. Integrating with CMDB and AD allows the SOAR playbook to gain crucial context about the source IP. If it's an internal or known legitimate external source (like a VPN client), the automated blocking can be bypassed in favor of human review, preventing false positives. Enriching with geographical data further assists the analyst in their decision. This combines automation for clear threats with human oversight for ambiguous cases. Option A is too aggressive and ignores the false positive problem. Option C shifts the problem to the SIEM and might miss legitimate attacks originating from internal misconfigurations. Option D is a simplistic thresholding that might still lead to false positives (e.g., a highly active but legitimate internal scanner) or miss more subtle attacks. Option E completely removes the automation benefit of SOAR, increasing response time and analyst workload.

Question: 4

A newly onboarded Cybersecurity Apprentice is attempting to understand the difference between SIEM and SOAR capabilities in the context of an ongoing ransomware incident. The incident began with a successful phishing attack, leading to malware execution, data encryption, and an extortion note. The security team has a robust SIEM and a newly implemented SOAR platform (e.g., Demisto/Cortex XSOAR). Which of the following statements accurately differentiate their primary roles and how they would jointly contribute to mitigating this specific ransomware incident?

A. The SIEM would primarily focus on ingesting all logs from endpoints, network devices, and email gateways to detect anomalies indicative of ransomware (e.g., mass file renames, suspicious outbound connections). The SOAR would then automatically generate a comprehensive incident report and email it to the CIO.
B. The SIEM's main role is to orchestrate the response, such as isolating infected hosts and blocking C2 servers. The SOAR's primary function is to collect and correlate security events from various sources and generate alerts.
C. The SIEM would aggregate and correlate logs to identify the initial phishing email, the compromised host, and indicators of lateral movement or data exfiltration attempts. The SOAR would then automate response actions like isolating infected machines, terminating malicious processes, rolling back compromised systems from clean backups, and orchestrating communication with legal and PR teams, creating a structured workflow for the entire incident lifecycle.
D. The SOAR is solely responsible for ingesting threat intelligence feeds and applying them to network devices for real-time blocking. The SIEM's role is to provide a user interface for security analysts to manually search logs for indicators of compromise (IOCs).
E. Both SIEM and SOAR are essentially the same product, just marketed differently. Their capabilities fully overlap, and either could be used interchangeably to detect and respond to the ransomware incident.

Answer: C

Explanation:
Option C accurately differentiates and explains the synergistic roles of SIEM and SOAR in a complex incident like ransomware. The SIEM (Security Information and Event Management) is designed for large-scale log aggregation, correlation, and anomaly detection crucial for identifying the initial attack vectors, compromise points, and malicious activity patterns. Once the SIEM generates a high-fidelity alert, the SOAR (Security Orchestration, Automation, and Response) takes over. Its strength lies in automating and orchestrating the complex sequence of response actions (isolation, process termination, data restoration, communication) across disparate security tools, creating a structured and efficient incident response workflow. Option A understates SOAR's active response capabilities. Option B completely reverses the primary roles of SIEM and SOAR. Option D inaccurately describes the core functions of both; SIEM is much more than a search interface, and SOAR orchestrates, it doesn't solely ingest threat intel for direct blocking. Option E is incorrect; SIEM and SOAR have distinct, complementary functionalities.

Question: 5

A Security Operations Center (SOC) is leveraging a new A1-powered alert analysis system. Lately, analysts have observed a significant increase in 'false positive' alerts related to anomalous login attempts, despite the underlying behavior remaining consistent. Upon investigation, it's discovered that

a recent update to the system's threat intelligence feed introduced a new set of indicators of compromise (IOCs) for a previously unseen APT group. How might the A1's alert analysis mechanism, specifically its unsupervised learning components, be contributing to this issue, and what immediate action should the SOC take to mitigate it without disabling the new feed entirely?

A. The unsupervised learning model is over-fitting to the new IOCs, perceiving them as baseline anomalous behavior. The SOC should re-train the model with a filtered dataset excluding the new IOCs temporarily.
B. The new IOCs are significantly different from the A1's established 'normal' baseline, causing high deviation scores. The SOC should adjust the anomaly detection thresholds for login events to be more permissive.
C. The unsupervised learning model is treating the new, legitimate IOCs as novel malicious patterns due to lack of historical context. The SOC should manually 'label' a subset of these new alerts as benign to retrain the model in a supervised manner.
D. The A1's feature engineering component is giving undue weight to the new IOCs. The SOC should implement a rule-based suppression for alerts generated solely based on the new threat intelligence feed.
E. The new IOCs are polluting the A1's internal representation of 'normal' login activity, leading to concept drift. The SOC should implement an adaptive baseline adjustment mechanism that slowly incorporates the new data while monitoring false positive rates.

Answer: E

Explanation:
This scenario highlights 'concept drift' in A1 models used for anomaly detection. When a significant, legitimate shift occurs in the underlying data distribution (like new IOCs that are genuinely different but not necessarily malicious in the current context), the A1's internal model of 'normal' can become outdated. Option E directly addresses this by suggesting an adaptive baseline adjustment, which allows the A1 to gradually incorporate the new data into its understanding of normal behavior, thereby reducing false positives over time without discarding valuable threat intelligence. Option A suggests retraining, which is a broader solution and might discard valuable context. Option B is a blanket threshold adjustment, which could lead to missed true positives. Option C suggests supervised labeling, which is an inefficient and impractical approach for a large volume of new, potentially legitimate data. Option D suggests a rigid rule-based suppression, undermining the A1's core functionality.

Question: 6

A Palo Alto Networks customer is using Cortex XDR with an integrated A1 module for alert correlation and prioritization. The SOC team has noticed that while the A1 effectively identifies high-fidelity attacks, it occasionally misses subtle, multi-stage reconnaissance activities that precede a major breach, even when individual low-severity alerts (e.g., numerous failed login attempts from a new IP, unusual port scans) are generated. The A1's documentation states it uses a 'graph-based neural network' for correlating events. Which of the following, if implemented or adjusted, would most effectively enhance the A1's ability to detect these low-signal, multi-stage reconnaissance patterns?

A. Increase the confidence threshold for generating high-severity alerts. This will force the A1 to be more precise.

B. Lower the individual alert severity thresholds for reconnaissance-related events, making them more likely to be fed into the graph-based correlation engine.
C. Enhance the feature engineering within the A1 by incorporating more contextual data sources, such as external reputation feeds and geo-location data, directly into the graph nodes representing entities (IPs, users, assets).
D. Adjust the weighting parameters within the graph-based neural network to prioritize temporal proximity and sequential dependencies between seemingly disparate low-severity events.
E. Implement an unsupervised clustering algorithm on historical alert data to identify latent patterns of reconnaissance that the supervised models might be overlooking.

Answer: C,D

Explanation:
This question targets the nuances of graph-based A1 for alert correlation. Multi-stage reconnaissance is about connecting low- signal events over time. Option C (Correct) : Enhancing feature engineering with more contextual data directly enriches the nodes and edges of the graph. For a graph-based neural network, richer node attributes (e.g., an IP's reputation, user's typical login times, asset's critical classification) allow the network to form more meaningful connections and detect subtle anomalies in relationships that signify reconnaissance. For example, a series of low-severity events from an IP with a poor reputation score would be more accurately weighted. Option D (Correct) : Graph-based neural networks learn relationships. Explicitly adjusting weighting parameters to prioritize temporal proximity and sequential dependencies (e.g., 'port scan followed by failed login from the same IP within X minutes') directly addresses the 'multi-stage' aspect of reconnaissance. This allows the network to assign higher significance to a chain of events that, individually, might be low severity but collectively indicate malicious intent. Option A is incorrect because increasing thresholds would make it harder, not easier, to detect subtle, low-signal activities. Option B is partially correct in that it feeds more data, but it doesn't address how the A1 correlates this data effectively. Simply lowering thresholds might lead to alert fatigue if the correlation engine isn't optimized to find the relevant chains. Option E suggests an additional unsupervised algorithm, which could be beneficial, but it's an orthogonal improvement rather than a direct enhancement to the existing graph- based neural network's ability to correlate and detect these patterns.

Question: 7

A cybersecurity apprentice is tasked with evaluating the explainability of an A1-driven alert analysis system that uses a deep learning model. The SOC manager is concerned about 'black box' issues, particularly when the system flags a seemingly innocuous event as high-severity. To address this, the apprentice needs to query the system for insights into its decision-making process for a specific alert ID 'ALERT-2023-XYZ'. Which of the following programmatic approaches, if available, would provide the most actionable insights into why the A1 escalated 'ALERT-2023-XYZ'?

A. Retrieve the raw log data associated with 'ALERT-2023-XYZ' and manually review for obvious anomalies that the A1 might have missed.
B. Execute a SHAP (SHapley Additive exPlanations) or LIME (Local Interpretable Model-agnostic Explanations) query on the A1 model for 'ALERT-2023-XYZ' to identify the specific features contributing most to its prediction.

C. Query the system's internal knowledge base for rules or signatures that match 'ALERT-2023-XYZ', as A1 often relies on pre-defined rules.
D. Run a sensitivity analysis on the input features of 'ALERT-2023-XYZ' by perturbing each feature slightly and observing the change in the A1's output confidence score.
E. Access the A1 model's training dataset and search for similar high-severity alerts to understand the patterns it learned.

Answer: B

Explanation:
This question focuses on A1 explainability in the context of deep learning models, which are often 'black boxes'. Option B (Correct) : SHAP and LIME are state-of-the-art model-agnostic explainability techniques specifically designed to explain the predictions of complex machine learning models (like deep learning models) at an individual instance level. They generate feature importance scores, showing which input features contributed positively or negatively to a specific prediction. This directly answers 'why' the A1 escalated 'ALERT-2023-XYZ by highlighting the most influential features. Option A is basic manual review and doesn't explain the A1's internal logic. Option C is incorrect because deep learning models don't typically rely on explicit 'rules' or signatures in the traditional sense; they learn complex patterns. Option D (Sensitivity analysis) can provide some insight but is often less precise and harder to interpret than SHAP/LIME for identifying direct feature contributions to a specific prediction. Option E (Training dataset review) might show learned patterns generally, but it doesn't explain why a specific alert was flagged without a detailed comparison and analysis of feature differences.

Question: 8

Consider a scenario where a Palo Alto Networks SOC uses an A1-powered alert analysis system that employs Reinforcement Learning (RL) for dynamic thresholding and alert routing. The RL agent observes analyst feedback (e.g., marking an alert as true positive/false positive, time to resolve) and adjusts its strategies to optimize alert queue efficiency and detection accuracy. After a major incident, the SOC identifies a critical blind spot: a specific low-volume, high-impact attack vector (e.g., a specific supply chain compromise technique) that the A1 consistently de-prioritizes or fails to correlate, despite relevant low-severity indicators being present in the logs. This leads to a delayed response. Which of the following actions, focusing on the RL aspect, is most likely to rectify this blind spot without significantly degrading overall system performance?

A. Manually hardcode a rule to elevate the severity of all alerts containing keywords related to the identified attack vector, bypassing the RL agent's decision- making.
B. Introduce a 'penalty' into the RL agent's reward function for failing to identify or prioritize alerts related to the specific high-impact attack vector, and provide historical examples of such incidents as 'expert demonstrations' during RL agent re-training.
C. Increase the exploration parameter of the RL agent, allowing it to experiment more with different alert prioritization strategies, even if they initially lead to lower immediate rewards.
D. Retrain the entire RL model from scratch with a much larger and more diverse dataset of attack scenarios, ensuring comprehensive coverage of all known attack types.
E. Reduce the number of states and actions in the RL agent's environment to simplify its learning process, making it easier for it to converge on optimal policies.

Answer: B

Explanation:
This question delves into the application and fine-tuning of Reinforcement Learning in Security Operations. Option B (Correct) : This is the most targeted and effective approach. In RL, the reward function dictates what the agent learns to optimize. By introducing a specific 'penalty' (or a significant positive reward for successful detection/prioritization) for the high-impact attack vector, you are directly guiding the RL agent to prioritize this specific type of event. Providing historical examples as 'expert demonstrations' (a technique often used in Imitation Leaming or Inverse Reinforcement Learning) further accelerates the learning process for these critical but rare scenarios without requiring extensive trial-and-error by the agent. This allows the agent to learn the desired behavior for the blind spot without drastically altering its overall optimal strategy for other alerts. Option A (Hardcoding rules) bypasses the A1's intelligence and adaptability, which defeats the purpose of an RL system. Option C (Increasing exploration) might eventually lead to finding the optimal strategy for the blind spot, but it comes at the cost of immediate performance degradation and increased false positives/misses during the exploration phase, which is undesirable for a critical SOC system. Option D (Retraining from scratch with a larger dataset) is a costly and often unnecessary measure. It's a blunt instrument for a specific blind spot and doesn't guarantee the RL agent will prioritize rare, high-impact events if the reward function isn't specifically tuned for them. Option E (Reducing states/actions) oversimplifies the problem space and could lead to a less optimal overall policy, making the system less effective at handling the complexity of real-world alerts.

Question: 9

A cybersecurity apprentice is tasked with designing a secure and efficient network for a new department within a large enterprise. This department, 'Quantum Innovations,' requires high bandwidth, low latency, and robust segmentation for sensitive R&D dat
a. They also need to integrate a legacy industrial control system (ICS) that communicates via Modbus/ TCP, located in a shielded lab within the same building. The CISO mandates strict isolation of the ICS network from the main corporate LAN and external internet, while still allowing controlled, audited access for specific IT personnel. Which of the following LAN design principles and technologies would be most appropriate to meet these complex requirements, considering best practices for both performance and security?
A. Implementing a flat, single VLAN LAN across the entire Quantum Innovations department to simplify management, relying solely on host-based firewalls for ICS isolation.
B. Deploying a dedicated, physically separate LAN for the ICS, utilizing a Palo Alto Networks NGFW as a highly restrictive, stateful firewall at the perimeter of this isolated network, with specific Security Policy rules for Modbus/TCP traffic- The R&D LAN would use 802. IQ VLANs to segment different research groups.
C. Utilizing a software-defined networking (SDN) overlay network for the entire department, with micro-segmentation policies applied at the hypervisor level for both R&D and ICS systems, bypassing the need for traditional VLANs or physical separatiom
D. Extending the existing corporate wireless LAN (WLAN) into the Quantum Innovations department, with strong WPA3 encryption, and using guest VLANs for the ICS to prevent direct access.
E. Connecting the ICS directly to the main R&D I-AN via a standard Ethernet switch and relying on a network access control (NAC) solution to dynamically assign an isolated VLAN upon device authentication.

Answer: B

Explanation:
Option B is the most appropriate and secure approach- A physically separate LAN for the ICS, coupled with a Palo Alto Networks NGFW acting as a highly restrictive firewall, provides the strongest isolation and granular control over Modbus/TCP traffic, which is critical for industrial control systems- The NGFWs Application-ID and Content-ID capabilities allow for deep packet inspection and policy enforcement specific to Modbus/TCP, preventing unauthorized access or malicious commands- Using 802.1 Q VLANs for the R&D LAN within the same building allows for logical segmentation of different research groups, improving security and performance without the need for additional physical infrastructure. Option A is insecure due to a flat network and over-reliance on host-based firewalls. Option C, while advanced, might be an overkill and more complex for the immediate physical isolation requirement of ICS, and may not offer the same level of physical air- gapping. Option D is highly insecure for sensitive R&D and particularly ICS systems due to the inherent vulnerabilities of wireless networks for critical infrastructure. Option E relies too heavily on dynamic VLAN assignment and might still expose the ICS to some level of the main network before isolation, and lacks the explicit security posture of an NGFW at the perimeter.

Question: 10

Consider a distributed development team at a startup, 'CodeForge', utilizing a local area network for their daily operations. They have an on-premise Git repository server, a Jenkins CI/CD server, and several development workstations. The team frequently pulls large codebases and deploys builds. Recently, they've experienced intermittent network slowdowns during peak development hours, specifically when large Git clones or Jenkins builds are active. A junior cybersecurity apprentice, new to network troubleshooting, observes that a few developers are also streaming high-definition video during these peak times. The current network topology consists of a single unmanaged gigabit switch connecting all devices, with a shared internet uplink.
Which of the following actions, combining network fundamental principles with a focus on potential Palo Alto Networks solutions for future scalability and security, would be the most effective and proactive steps to diagnose and mitigate these issues, and lay a foundation for secure growth?

A. Replace the unmanaged switch with another unmanaged switch, but with more ports, assuming the issue is port exhaustion. Configure static IP addresses for all developer workstations.
B. Implement QOS (Quality of Service) on a new managed switch, prioritizing Git and Jenkins traffic over streaming video. Install a Palo Alto Networks NGFW at the network perimeter to perform App-ID analysis and traffic shaping based on application type, providing visibility into bandwidth hogs. Consider separating development and streaming traffic into different VLANs
C. Advise developers to stop streaming video. Install a basic firewall to block all non-HTTP/HTTPS traffic to prevent bandwidth abuse. Upgrade the internet uplink speed only.
D. Introduce a separate wireless access point for streaming devices, completely isolating them from the wired development network. Implement an Intrusion Prevention System (IPS) for all internal network traffic.
E. Migrate all services (Git, Jenkins) to a cloud provider to offload local network strain. This eliminates the need for local LAN optimization and security considerations.

Answer: B

Explanation:
Option B is the most comprehensive and effective solution. The current unmanaged switch offers no control over traffic, leading to contention. A managed switch with QOS allows prioritizing critical development traffic (Gitl Jenkins) over recreational streaming. Integrating a Palo Alto Networks NGFW provides deep visibility into application usage (App-ID), allowing the apprentice to identify and control bandwidth-intensive applications like streaming video. This also enables traffic shaping to ensure business-critical applications receive priority. Furthermore, segmenting development and streaming traffic into different VLANs provides logical isolation, reducing broadcast domains and improving security This approach directly addresses the current performance issues while building a foundation for future security and scalability. Option A is a naive solution and won't address congestion. Option C is reactive and lacks the granular control needed. Option D isolates streaming but doesn't manage or secure the core development network effectively, and an IPS without App-ID might not be as effective for bandwidth management. Option E is a significant architectural shift that might be an eventual goal but doesn't directly address immediate LAN performance issues or build local network expertise.

Thank You for Trying Our Product

For More Information – Visit link below:

https://www.examsboost.com/

15 USD Discount Coupon Code:

G74JA8UF

FEATURES

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