DUMPSBOSS.

Palo Alto Networks Cybersecurity

Palo Alto Networks Apprentice

Version Demo

Total Demo Questions: 12

Total Premium Questions: 127

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Cybersecurity Fundamentals	33
Topic 2, Network Security Fundamentals	51
Topic 3, Cloud Security Fundamentals	22
Topic 4, Security Operations Fundamentals	21
Total	127

QUESTION NO: 1

Which OSI layer is used to determine how long communications are open between two devices?

A. Transport

B. Application

C. Session

D. Network

ANSWER: C

Explanation:

The Session layer is responsible for establishing, managing, and terminating communication sessions between two devices or applications. A session represents an ongoing exchange of information, and this layer provides the conceptual OSI function for controlling the dialog: when communication begins, remains active, is synchronized, and ends. Therefore, determining how long communications stay open is a Session-layer responsibility. Session management can include maintaining dialog state, coordinating which endpoint transmits at a given time, and adding synchronization points that allow a communication process to recover from an interruption. These functions are especially relevant to security because an active session can represent an authenticated user or an established application interaction; managing session lifetime helps limit exposure from abandoned, hijacked, or unnecessarily persistent sessions. Although many modern TCP/IP implementations combine Session-layer functions into application protocols, middleware, or operating-system services rather than implementing a distinct Session-layer protocol, the OSI reference model assigns communication-duration and dialog-management responsibilities to the Session layer. See the OSI model overview from ISO/IEC 7498-1 and Cisco's OSI model explanation.

QUESTION NO: 2

What are two endpoint security implementation methods? (Choose two.)

A. Installing an anti-malware agent onto a user device

B. Deploying a firewall to prevent traffic from reaching an end user

C. Enforcing security policies on north-south traffic between users and the internet

D. Downloading software onto a laptop to prevent spyware

ANSWER: A D

Explanation:

Installing an anti-malware agent onto a user device is an endpoint security implementation because the protective software runs directly on the host. It can inspect files, processes, memory activity, and device behavior to identify and prevent malicious activity where users access data and applications. This approach protects the endpoint regardless of the particular network through which the device connects.

Downloading software onto a laptop to prevent spyware is also endpoint security. Anti-spyware software is installed locally and is intended to detect, block, or remove software that monitors a user, captures information, or otherwise compromises the device. Both methods apply security controls at the device layer, making them appropriate examples of endpoint protection. Palo Alto Networks delivers endpoint protection capabilities through Cortex XDR, which combines endpoint agents with prevention and detection features to help secure user devices. See the Palo Alto Networks Cortex XDR overview and Palo Alto Networks' endpoint security overview for additional context on endpoint-focused controls.

QUESTION NO: 3

What are two areas in which AI can help Security Operations Center (SOC) teams with alerts? (Choose two.)

A. Vulnerability patching

B. Alert triage

C. SOC resource management

D. Incident response

ANSWER: B D

Explanation:

Alert triage is a core area where AI helps SOC teams handle alerts more effectively. AI can rapidly correlate related telemetry, enrich alerts with relevant context, summarize large volumes of event data, and help prioritize the alerts that warrant an analyst's attention. This reduces time spent on repetitive investigation tasks and supports faster, more consistent assessment of potential security events.

Incident response is also strengthened by AI because analysts need actionable context after an alert has been identified as significant. AI-assisted capabilities can summarize an incident timeline, identify relevant entities and indicators, map observed activity to known adversary behaviors, recommend investigation or containment actions, and help analysts document and execute response workflows. Palo Alto Networks describes its Cortex platform as using AI and automation to support security operations, including alert investigation and response. These capabilities are intended to augment analyst decision-making and accelerate response while keeping people in control of consequential actions. See Palo Alto Networks Cortex XSIAM and Palo Alto Networks Cortex XDR.

QUESTION NO: 4

What are two functions of VPN gateways? (Choose two.)

A. Certificate refresh

B. Site-to-Site connectivity

C. Remote access

D. URL filtering

ANSWER: B C

Explanation:

VPN gateways provide secure connectivity over untrusted networks by creating encrypted VPN tunnels and enforcing the policies that govern traffic entering and leaving those tunnels. **Site-to-Site connectivity** is a core VPN gateway function: an IPsec VPN gateway establishes a tunnel between two network endpoints, allowing protected communication between separate networks such as a branch office and a data center. The gateway encapsulates traffic for the remote protected network, encrypts it for transit, and decrypts traffic received through the tunnel before forwarding it internally. Palo Alto Networks documents this model in its IPsec VPN guidance: IPSec VPN.

Remote access is also a VPN gateway function. A remote-access gateway accepts secure connections from individual users outside the organization, authenticates them, applies access controls, and provides access to permitted internal applications and resources. In Palo Alto Networks GlobalProtect deployments, gateways are the components that provide this secure access and enforce security policy for connected users. The role and operation of these gateways are described in the GlobalProtect components overview. Together, these functions protect both inter-network communications and user-to-network access while traffic crosses public networks.

QUESTION NO: 5

What is the primary responsibility of the cloud provider in the cloud shared responsibility model?

A. Configuring application-level security settings

B. Securing underlying physical servers and network infrastructure

C. Providing end-user training on application usage

D. Monitoring and managing user access and permissions

ANSWER: B

Explanation:

Securing underlying physical servers and network infrastructure is the cloud provider's primary responsibility in the shared responsibility model. This is commonly described as *security of the cloud*: the provider operates and protects the data centers, physical facilities, hardware, networking components, virtualization layer, and foundational services that make its cloud platform available. Because customers do not own or operate this underlying environment, they rely on the provider to maintain its physical protection, resilience, maintenance, and core infrastructure security.

The exact boundary of provider responsibility varies according to the service model. For example, a SaaS provider assumes responsibility for more of the technology stack than an IaaS provider. However, securing the physical infrastructure and the core platform used to deliver cloud services remains a foundational provider obligation across cloud service models. The customer's responsibilities generally concern the resources, data, identities, configurations, and workloads they deploy or use within that platform. Palo Alto Networks describes the shared responsibility model as an allocation of cloud-security duties between the cloud provider and customer: Shared Responsibility Model. For further cloud-security context, see Palo Alto Networks Cloud Security.

QUESTION NO: 6

Syslog would be used for which activity?

A. Transferring log events within networks

B. Securing endpoints via runtime protection

C. Securing logs collected from endpoints

D. Connecting to a system remotely

ANSWER: A

Explanation:

Syslog is used for **transferring log events within networks**. It is a standardized logging protocol and message format through which devices and software—including firewalls, routers, servers, operating systems, and applications—send event records to a centralized log collector, log management platform, or SIEM. These events can include authentication activity, configuration changes, traffic information, system errors, and security alerts. Centralizing this data helps security and IT operations teams monitor infrastructure, investigate incidents, correlate events across sources, and maintain records needed for operational troubleshooting and compliance. In Palo Alto Networks environments, firewalls can forward traffic, threat, system, configuration, and other logs to external syslog servers, allowing those events to be retained and analyzed alongside logs from other technologies. Syslog's role is the delivery of event messages; protection, storage controls, and analysis can then be provided by the receiving logging or security platform. Palo Alto Networks documents syslog forwarding as a method for forwarding firewall logs to an external syslog server: Configure Syslog Monitoring. The underlying Syslog protocol is standardized by the IETF in RFC 5424.

QUESTION NO: 7

Which type of device does a Host-Based Intrusion Detection System (HIDS) monitor?

A. Appliance

B. Computer

C. Switch

D. Router

ANSWER: B

Explanation:

Computer is correct because a Host-Based Intrusion Detection System operates on and monitors an individual host, such as a workstation, laptop, server, or other endpoint. A HIDS collects and analyzes information generated locally by that computer, including operating-system and application logs, authentication activity, system calls, configuration changes, running processes, and file-integrity events. By examining this host-level telemetry, it can identify suspicious behavior or indicators that may not be visible solely from network traffic, such as unauthorized modifications to critical files or unusual local privilege activity. The defining characteristic is that the monitored scope is the individual computing system itself, rather than a network segment. Palo Alto Networks describes host-based intrusion detection as a system that monitors activity occurring within a host and can detect changes to files, logs, and system configuration. This distinction is fundamental to understanding the placement and visibility of intrusion-detection technologies in a cybersecurity architecture. See Palo Alto Networks' overview of intrusion detection systems and its explanation of endpoints.

QUESTION NO: 8

An organization is preparing for a ransomware incident and wants to ensure that critical systems can be restored after an attack. Which two actions most directly support this objective? (Choose two.)

A. Maintaining protected backups of critical data

B. Regularly testing restoration procedures

C. Using multi-factor authentication for administrator access

D. Blocking access to nonbusiness website categories

ANSWER: A B

Explanation:

Maintaining protected backups and **regularly testing restoration procedures** most directly support recovery from ransomware. Backups provide a recoverable copy of data when production files are encrypted or destroyed. Protecting backups from unauthorized modification or deletion helps prevent ransomware from affecting both production data and recovery data.

Restoration testing is equally important because an untested backup may be incomplete, inaccessible, or too slow to restore when an incident occurs. MFA and URL filtering are valuable preventive controls, but they do not by themselves provide the ability to recover encrypted systems and data.

QUESTION NO: 9

What are two components of multi-factor authentication (MFA)? (Choose two.)

A. Something you know

B. Something you observe

C. Something you have

D. Something you create

ANSWER: A C

Explanation:

Multi-factor authentication requires authentication evidence from separate factor categories, so that compromise of one factor alone does not provide access. “Something you know” is a knowledge factor: information the user can recall, such as a password, passphrase, or PIN. “Something you have” is a possession factor: an item or device controlled by the user, such as a hardware security key, smart card, authenticator application, or one-time-password token. When a user supplies both a password and a code generated by an authenticator device, the authentication process combines knowledge and possession factors and is therefore MFA.

These factors are independent categories rather than simply two pieces of the same type of information. Palo Alto Networks supports MFA integrations for GlobalProtect and other access workflows so that identity verification can require an additional factor beyond the user’s primary credentials. This approach reduces the likelihood that a stolen or guessed password alone can be used to gain access. For implementation guidance, see Palo Alto Networks’ GlobalProtect multifactor authentication documentation and the CISA overview of multifactor authentication.

QUESTION NO: 10

What is a function of least privilege in Identity and Access Management (IAM)?

A. Providing temporary access to resources for all users

B. Granting users permission to all resources in the network

C. Ensuring network users have the same access level for simplicity

D. Allowing users the minimum permissions necessary to perform their jobs

ANSWER: D

Explanation:

Allowing users the minimum permissions necessary to perform their jobs is the core function of the principle of least privilege in IAM. It ensures that each human user, service account, application, or device receives access only to the specific systems, data, and actions required for its authorized task. Permissions should be scoped as narrowly as practical and reviewed as responsibilities change. This approach reduces the potential impact of compromised credentials, configuration mistakes, or inappropriate account use: even if an identity is misused, it cannot access resources beyond its assigned role and duties. In practice, organizations implement least privilege with role-based access control, narrowly defined IAM policies, separate administrative accounts, periodic access reviews, and time-limited elevation when a privileged task is genuinely needed. Palo Alto Networks describes least-privilege access as granting users only the access necessary to perform their roles, which helps reduce the attack surface and limit lateral movement. The U.S. National Institute of Standards and Technology also identifies least privilege as restricting privileges to the minimum necessary for assigned duties. See Palo Alto Networks: What Is IAM? and NIST: Least Privilege.

QUESTION NO: 11

A development team wants to identify security issues before a containerized application is deployed. Which two practices support a shift-left security approach? (Choose two.)

A. Scanning third-party dependencies for known vulnerabilities

B. Detecting exposed secrets in source code and build configurations

C. Monitoring workload behavior only after deployment

D. Applying URL filtering to employee web traffic

E. Updating incident response plans after an incident

ANSWER: A B

Explanation:

Scanning third-party dependencies for known vulnerabilities helps identify risk in libraries and packages before they reach production. **Detecting exposed secrets in source code and build configurations** helps prevent credentials, tokens, and keys from being included in application artifacts or repositories.

Runtime monitoring is valuable after deployment, but it is not a predeployment shift-left control. URL filtering governs web access, and incident-response plan updates occur after operational events rather than during application development.

QUESTION NO: 12

What is a cluster in relation to cloud-native security?

A. Portable and self-sufficient unit that packages an application with its dependencies

B. Set of system rules written in a particular programming language

C. Collection of nodes (bare-metal or virtualized machines) that will host application pods

D. Distributed collection of servers that hosts software and is accessible over the internet

ANSWER: C

Explanation:

A collection of nodes (bare-metal or virtualized machines) that will host application pods is the correct definition of a cluster in a cloud-native, Kubernetes-based environment. A Kubernetes cluster provides the shared infrastructure and control plane used to run and manage containerized workloads. Its worker nodes supply the compute capacity on which Pods—the smallest deployable Kubernetes objects—are scheduled and executed. Nodes can be physical servers or virtual machines, depending on where the cluster is deployed. The cluster also supplies the orchestration context for workload scheduling, service networking, scaling, availability, configuration, and identity controls. From a cloud-native security perspective, the cluster is a critical security boundary and management scope: security teams need visibility into its configuration, identities, network exposure, workloads, and runtime activity. Protecting the cluster helps ensure that the applications and Pods it hosts are deployed and operated according to policy. Kubernetes documentation describes a cluster as a set of machines, called nodes, that run containerized applications, with worker nodes hosting Pods. See the Kubernetes Components documentation and Palo Alto Networks' Cloud-Native Security overview.