The SecOps Group CNSP Practice Questions

Shared by Burnett on 26-10-2026

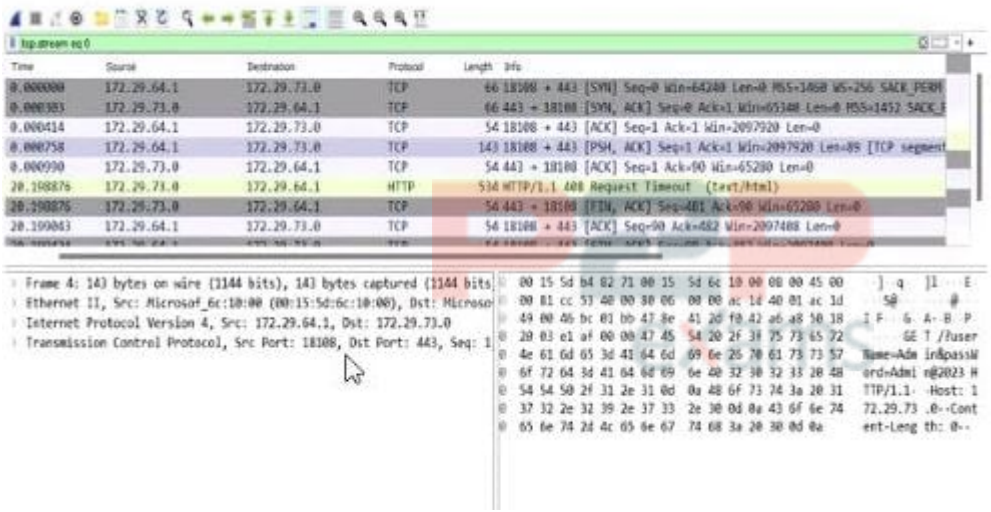
For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

According to the screenshot below, which of the following statements are correct?



Options:

A- The application is running on port 443 and the HTTPS protocol.
B- The credentials have been submitted over the HTTP protocol.
C- The credentials have been submitted over the HTTPS protocol.
D- The application is running on port 80 and the HTTP protocol.

Answer:

A

Explanation:

The screenshot is from Wireshark, a network protocol analyzer, displaying captured network traffic. The relevant columns include the source and destination IP addresses, ports, protocol, and additional information about the packets. Let's break down the details:

Destination Port Analysis: The screenshot shows multiple packets with a destination port of 443 (e.g., in the 'Destination' column, entries like '172.72.61.9:443'). Port 443 is the default port for HTTPS (HTTP Secure), which is HTTP traffic encrypted using SSL/TLS. This indicates that the application is communicating over HTTPS.

Protocol Analysis: The 'Protocol' column lists 'TLSv1.2' for most packets (e.g., frame numbers 2000084, 2000086). TLS (Transport Layer Security) is the cryptographic protocol used by HTTPS to secure HTTP communications. This confirms that the traffic is HTTPS, not plain HTTP.

Packet Details: The 'Info' column provides additional context, such as 'Application Data' for TLS packets, indicating encrypted application-layer data (typical of HTTPS). There are also HTTP packets (e.g., frame 2000088), but these are likely part of the HTTPS session (e.g., HTTP/2 over TLS, as noted by 'HTTP2').

Now, let's evaluate the options:

Option A: 'The application is running on port 443 and the HTTPS protocol.'

This is correct. The destination port 443 and the use of TLSv1.2 confirm that the application is using HTTPS. HTTPS is the standard protocol for secure web communication, and port 443 is its designated port. CNSP documentation emphasizes that HTTPS traffic on port 443 indicates a secure application-layer protocol, often used for web applications handling sensitive data.

Option B: 'The credentials have been submitted over the HTTP protocol.'

This is incorrect. HTTP typically uses port 80, but the screenshot shows traffic on port 443 with TLS, indicating HTTPS. Credentials submitted over this connection would be encrypted via HTTPS, not sent in plaintext over HTTP. CNSP highlights the security risks of HTTP for credential submission due to lack of encryption, which isn't the case here.

Option C: 'The credentials have been submitted over the HTTPS protocol.'

While this statement could be true (since HTTPS is in use, any credentials would likely be submitted securely), the question asks for the 'correct' statement based on the screenshot. The screenshot doesn't explicitly show credential submission (e.g., a POST request with form data); it only shows the protocol and port. Option A is more directly supported by the screenshot as it focuses on the application's protocol and port, not the specific action of credential submission. CNSP notes that HTTPS ensures confidentiality, but this option requires more specific evidence of credentials.

Option D: 'The application is running on port 80 and the HTTP protocol.'

This is incorrect. Port 80 is the default for HTTP, but the screenshot clearly shows port 443 and TLS, indicating HTTPS. CNSP documentation contrasts HTTP (port 80, unencrypted) with HTTPS (port 443, encrypted), making this option invalid.

Conclusion: Option A is the most accurate and comprehensive statement directly supported by the screenshot, confirming the application's use of port 443 and HTTPS. While Option C might be true in a broader context, it's less definitive without explicit evidence of credential submission in the captured packets.

Question 2

Question Type: MultipleChoice

Which of the following services do not encrypt its traffic by default?

Options:

A- DNS
B- SSH
C- FTPS
D- All of these

Answer:

A

Explanation:

Encryption ensures confidentiality and integrity of network traffic. Analyzing defaults:

A . DNS (Domain Name System):

Default: Unencrypted (UDP/TCP 53), per RFC 1035. Queries/responses (e.g., ''google.com 142.250.190.14'') are plaintext.

Modern Options: DNS over HTTPS (DoH, TCP 443) or DNS over TLS (DoT, TCP 853) encrypt, but aren't default in most systems (e.g., pre-2020 Windows).

B . SSH (Secure Shell):

Default: Encrypted (TCP 22), per RFC 4251. Uses asymmetric (e.g., RSA) and symmetric (e.g., AES) crypto for all sessions.

C . FTPS (FTP Secure):

Default: Encrypted (TCP 21 control, dynamic data ports). Extends FTP with SSL/TLS (e.g., RFC 4217), securing file transfers.

Technical Details:

DNS: Plaintext exposes queries to eavesdropping (e.g., ISP snooping) or spoofing (e.g., cache poisoning).

SSH/FTPS: Encryption is baked into their standards; disabling it requires explicit misconfiguration.

Security Implications: Unencrypted DNS risks privacy and integrity (e.g., Kaminsky attack). CNSP likely pushes DoH/DoT adoption.

Why other options are incorrect:

B, C: Encrypt by default.

D: False, as only DNS lacks default encryption.

Real-World Context: The 2013 Snowden leaks exposed DNS monitoring; DoH uptake (e.g.,

Cloudflare 1.1.1.1) counters this.

Question 3

Question Type: MultipleChoice

In the context of the SSH (Secure Shell) public-private key authentication mechanism, which key is uploaded to the server and which key is used by the end-user for authentication?

Options:

A- The public key is uploaded to the server and the private key is used by the end user for authentication.
B- The private key is uploaded to the server and the public key is used by the end user for authentication.

Answer:

A

Explanation:

SSH (Secure Shell), per RFC 4251, uses asymmetric cryptography (e.g., RSA, ECDSA) for secure authentication:

Key Pair:

Public Key: Freely shareable, used to encrypt or verify.

Private Key: Secret, used to decrypt or sign.

Process:

User generates a key pair (e.g., ssh-keygen -t rsa -b 4096).

Public Key is uploaded to the server, appended to ~/.ssh/authorized_keys (e.g., via ssh-copy-id).

Private Key (e.g., ~/.ssh/id_rsa) stays on the user's machine.

Authentication: Client signs a challenge with the private key; server verifies it with the public key.

Technical Details:

Protocol: SSH-2 (RFC 4253) uses a Diffie-Hellman key exchange, then public-key auth.

Files: authorized_keys (server, 0644 perms), private key (client, 0600 perms).

Security: Private key exposure compromises all systems trusting the public key.

Security Implications: CNSP likely stresses key management (e.g., passphrases, rotation) and server-side authorized_keys hardening (e.g., PermitRootLogin no).

Why other options are incorrect:

B: Uploading the private key reverses the model, breaking security---anyone with the server's copy could authenticate as the user. Asymmetric crypto relies on the private key remaining secret.

Real-World Context: GitHub uses SSH public keys for repository access, with private keys on user devices.

Question 4

Question Type: MultipleChoice

Which built-in Windows utility can be used to verify the validity of a Kerberos ticket?

Options:

A- Klist
B- Kerbtray
C- Netsh
D- Kerberos Manager

Answer:

A

Explanation:

Kerberos is the default authentication protocol in Windows Active Directory environments, and tickets are used to prove identity. Verifying ticket validity involves checking their status, expiration, and attributes, which requires a built-in tool available in modern Windows systems.

Why A is correct: Klist is a command-line utility included in Windows (since Vista/2008) that lists cached Kerberos tickets and their details, such as validity period and renewal status. CNSP recognizes it as the standard tool for Kerberos ticket management in security audits.

Why other options are incorrect:

B: Kerbtray is a graphical tool from the Windows Resource Kit, not a built-in utility, and is outdated.

C: Netsh manages network configurations, not Kerberos tickets.

D: 'Kerberos Manager' is not a recognized built-in Windows utility; it's a fictitious name.

Question 5

Question Type: MultipleChoice

What ports does an MSSQL server typically use?

Options:

A- 1433/TCP, 2433/UDP, and 3433/TCP
B- 1433/TCP, 1434/UDP, and 1434/TCP
C- 1433/TCP, 2433/UDP, and 1434/TCP
D- 1533/TCP, 1434/UDP, and 2434/TCP

Answer:

B

Explanation:

Microsoft SQL Server (MSSQL) relies on specific ports for its core services, as defined by Microsoft and registered with IANA:

1433/TCP: The default port for the SQL Server Database Engine. Clients connect here for querying databases (e.g., via ODBC or JDBC). It's a well-known port, making it a frequent target for attacks if exposed.

1434/UDP: Used by the SQL Server Browser Service, which listens for incoming requests and redirects clients to the correct port/instance (especially for named instances). It's critical for discovering dynamic ports when 1433 isn't used.

1434/TCP: Less commonly highlighted but used in some configurations, such as dedicated admin connections (DAC) or when the Browser Service responds over TCP for specific instances. While 1433/TCP is the primary engine port, 1434/TCP can be involved in multi-instance setups.

Technical Details:

Ports can be customized (e.g., via SQL Server Configuration Manager), but these are defaults.

Named instances often use dynamic ports (allocated from the ephemeral range), with the Browser Service (1434/UDP) guiding clients to them.

Firewalls must allow these ports for MSSQL to function externally, posing risks if not secured (e.g., brute-force attacks on 1433/TCP).

Security Implications: CNSP likely covers MSSQL port security, as vulnerabilities like SQL Slammer (2003) exploited 1434/UDP misconfigurations. Hardening includes restricting access, changing defaults, and monitoring traffic.

Why other options are incorrect:

A . 1433/TCP, 2433/UDP, 3433/TCP: 2433/UDP and 3433/TCP are not MSSQL standards; they're likely typos or unrelated ports.

C . 1433/TCP, 2433/UDP, 1434/TCP: 2433/UDP is incorrect; 1434/UDP is the Browser Service port.

D . 1533/TCP, 1434/UDP, 2434/TCP: 1533/TCP and 2434/TCP aren't associated with MSSQL; they deviate from documented defaults.

Real-World Context: Tools like netstat -an | find '1433' on Windows confirm MSSQL's port usage during audits.

Question 6

Question Type: MultipleChoice

What is the response from a closed TCP port which is not behind a firewall?

Options:

A- ICMP message showing Port Unreachable
B- A FIN and an ACK packet
C- A RST and an ACK packet
D- A SYN and an ACK packet

Answer:

C

Explanation:

TCP uses a structured handshake, and its response to a connection attempt on a closed port follows a specific protocol when unobstructed by a firewall.

Why C is correct: A closed TCP port responds with a RST (Reset) and ACK (Acknowledgment) packet to terminate the connection attempt immediately. CNSP highlights this as a key scanning indicator.

Why other options are incorrect:

A: ICMP Port Unreachable is for UDP, not TCP.

B: FIN/ACK is for closing active connections, not rejecting new ones.

D: SYN/ACK indicates an open port during the TCP handshake.

P2P exams

Question 7

Question Type: MultipleChoice

Which one of the following is not an online attack?

Options:

A- Brute force attack
B- Rainbow table attack
C- Password spraying attack
D- Phishing attack

P2P exams

Answer:

B

Explanation:

Online attacks require real-time interaction with a target system (e.g., a login interface), whereas offline attacks occur without direct system interaction, typically after obtaining data like password hashes. A rainbow table attack is an offline method that uses precomputed tables of hash values to reverse-engineer passwords from stolen hash databases, distinguishing it from the other options, which are online.

Why B is correct: Rainbow table attacks are performed offline after an attacker has already

acquired a hash (e.g., from a compromised database). The attacker matches the hash against precomputed tables to find the plaintext password, requiring no interaction with the target system during the attack. CNSP classifies this as an offline password recovery technique.

Why other options are incorrect:

A: Brute force attacks involve repeatedly submitting password guesses to a live system (e.g., via SSH or a web login), making it an online attack.

C: Password spraying attacks test a few common passwords across many accounts on a live system, also an online attack aimed at avoiding lockouts.

D: Phishing attacks trick users into submitting credentials through fake interfaces (e.g., emails or websites), requiring real-time interaction and thus classified as online.

P2P exams

Question 8

Question Type: MultipleChoice

How would you establish a null session to a Windows host from a Windows command prompt?

Options:

A- net use \hostname\c$ '' /u:NULL
B- net use \hostname\c$ '' /u:''
C- net use \hostname\ipc$ '' /u:''
D- net use \hostname\ipc$ '' /u:NULL

P2P exams

Answer:

C

Explanation:

A null session in Windows is an unauthenticated connection to certain administrative shares, historically used for system enumeration. The net use command connects to a share, and the IPC$ (Inter-Process Communication) share is the standard target for null sessions, allowing access without credentials when configured to permit it.

Why C is correct: The command net use \\hostname\ipc$ '' /u:'' specifies the IPC$ share and uses empty strings for the password (first '') and username (/u:''), establishing a null session. This syntax is correct for older Windows systems (e.g., XP or 2003) where null sessions were more permissive, a topic covered in CNSP for legacy system vulnerabilities.

Why other options are incorrect:

A: Targets the c$ share (not typically used for null sessions) and uses /u:NULL, which is invalid syntax; the username must be an empty string ('').

B: Targets c$ instead of ipc$, making it incorrect for null session establishment.

D: Uses ipc$ correctly but specifies /u:NULL, which is not the proper way to denote an empty username.

Question 9

Question Type: MultipleChoice

What will be the subnet mask for 192.168.0.1/18?

Options:

A- 255.255.255.0
B- 255.225.225.0
C- 255.255.192.0
D- 255.225.192.0

Answer:

C

Explanation:

An IP address with a /18 prefix (CIDR notation) indicates 18 network bits in the subnet mask, leaving 14 host bits (32 total bits - 18). For IPv4 (e.g., 192.168.0.1):

Binary Mask: First 18 bits are 1s, rest 0s.

1st octet: 11111111 (255)

2nd octet: 11111111 (255)

3rd octet: 11000000 (192)

4th octet: 00000000 (0)

Decimal: 255.255.192.0

Calculation:

Bits: /18 = 2^14 hosts (16,384), minus 2 (network/broadcast) = 16,382 usable.

Range: 192.168.0.0--192.168.63.255 (3rd octet: 0--63, as 192 = 11000000 covers 6 bits).

Technical Details:

Subnet masks align on octet boundaries or mid-octet (e.g., 192 = 2^7 + 2^6).

Contrast: /24 = 255.255.255.0 (256 hosts), /16 = 255.255.0.0 (65,536 hosts).

Security Implications: Larger subnets (e.g., /18) increase broadcast domains, risking amplification attacks. CNSP likely teaches subnetting for segmentation (e.g., VLANs).

Why other options are incorrect:

A . 255.255.255.0: /24 (8 host bits), not /18.

B . 255.225.225.0: Invalid mask (225 = 11100001, non-contiguous 1s).

D . 255.225.192.0: Invalid (225 breaks binary sequence).

Real-World Context: Subnetting 192.168.0.0/18 isolates departments in enterprise networks.

Question 10

Question Type: MultipleChoice

Which option best statements regarding Authorization and Authentication is true?

Options:

A- Authorization is the process where requests to access a particular resource are granted or denied. Authentication is providing and validating the identity.
B- Authentication is the process where requests to access a particular resource are granted or denied. Authorization is providing and validating identity.
C- Authentication includes the execution rules that determine what functionality and data the user can access. Authentication and Authorization are both the same thing.
D- Authentication controls which processes a person can use and which files they can access, read, or modify. Authentication and authorization typically do not operate together, thus making it impossible to determine who is accessing the information.

Answer:

A

Explanation:

Authentication and Authorization (often abbreviated as AuthN and AuthZ) are foundational pillars of access control in network security:

Authentication (AuthN): Verifies 'who you are' by validating credentials against a trusted source. Examples include passwords, MFA (multi-factor authentication), certificates, or biometrics. It ensures the entity (user, device) is legitimate, typically via protocols like Kerberos or LDAP.

Authorization (AuthZ): Determines 'what you can do' after authentication, enforcing policies on resource access (e.g., read/write permissions, API calls). It relies on mechanisms like Access Control Lists (ACLs), Role-Based Access Control (RBAC), or Attribute-Based Access Control (ABAC).

Option A correctly separates these roles:

Authorization governs access decisions (e.g., 'Can user X read file Y?').

Authentication establishes identity (e.g., 'Is this user X?').

In practice, these processes are sequential: AuthN precedes AuthZ. For example, logging into a VPN authenticates your identity (e.g., via username/password), then authorizes your access to specific subnets based on your role. CNSP likely stresses this distinction for designing secure systems, as conflating them risks privilege escalation or identity spoofing vulnerabilities.

Why other options are incorrect:

B: Reverses the definitions---Authentication doesn't grant/deny access (that's AuthZ), and Authorization doesn't validate identity (that's AuthN). This mix-up could lead to flawed security models.

C: Falsely equates AuthN and AuthZ and attributes access rules to AuthN. They're distinct processes; treating them as identical undermines granular control (e.g., NIST SP 800-53 separates IA-2 for AuthN and AC-3 for AuthZ).

D: Misassigns access control to AuthN and claims they don't interoperate, which is false---they work together in every modern system (e.g., SSO with RBAC). This would render auditing impossible, contradicting security best practices.

Real-World Context: A web server (e.g., Apache) authenticates via HTTP Basic Auth, then authorizes via .htaccess rules---two separate steps.

Question 11

Question Type: MultipleChoice

Which of the following is an example of a SUID program?

Options:

A- /bin/ls
B- /usr/bin/curl
C- /usr/bin/passwd
D- None of the above
Explanation:

The screenshot shows an email labeled 'B' with the subject 'Verify your email address' purportedly from Apple. To determine if this is a phishing email, we need to analyze its content and characteristics against common phishing indicators as outlined in CNSP documentation. Since option A is not provided in the screenshot, we will evaluate email B and infer the context for A.
Analysis of Email B:
Sender and Branding: The email claims to be from 'Apple Support' and includes an Apple logo, which is a common tactic to establish trust. However, phishing emails often impersonate legitimate brands like Apple to deceive users.
Subject and Content: The subject 'Verify your email address' and the body requesting the user to verify their email by clicking a link ('Verify Your Email') are typical of phishing attempts. Legitimate companies like Apple may send verification emails, but the tone and context here raise suspicion.
Link Presence: The email contains a clickable link ('Verify Your Email') that is purportedly for email verification. The screenshot does not show the URL, but phishing emails often include malicious links that lead to fake login pages to steal credentials. CNSP emphasizes that unsolicited requests to click links for verification are a red flag.
Urgency and Vague Instructions: The email includes a statement, 'If you did not make this change or believe an unauthorized person has accessed your account, click here to cancel and secure your account.' This creates a sense of urgency, a common phishing tactic to prompt immediate action without critical thinking.
Generic Greeting: The email starts with 'Dear User,' a generic greeting often used in phishing emails. Legitimate companies like Apple typically personalize emails with the user's name.
Suspicious Elements: The email mentions 'your Apple ID (example@icloud.com),' which is a placeholder rather than a specific email address, further indicating a mass phishing campaign rather than a targeted, legitimate communication.
Phishing Indicators (per CNSP):
CNSP documentation on phishing identification lists several red flags:
Unsolicited requests for verification or account updates.
Generic greetings (e.g., 'Dear User' instead of a personalized name).
Presence of links that may lead to malicious sites (not verifiable in the screenshot but implied).
Urgency or threats (e.g., 'click here to cancel and secure your account').
Impersonation of trusted brands (e.g., Apple).
Email B exhibits multiple indicators: the generic greeting, unsolicited verification request, urgent call to action, and impersonation of Apple.
Option A Context:
Since the screenshot only shows email B, and the correct answer is 'Only B,' we can infer that email A (not shown) does not exhibit phishing characteristics. For example, A might be a

legitimate email from Apple with proper personalization, no suspicious links, or a different context (e.g., a purchase confirmation rather than a verification request).
Evaluation of Options:
1. Only A: Incorrect, as email A is not shown, and the correct answer indicates B as the phishing email.
2. Only B: Correct. Email B shows clear phishing characteristics, such as impersonation, a generic greeting, an unsolicited verification link, and urgency, aligning with CNSP's phishing criteria.
3. Both A and B: Incorrect, as A is implied to be non-phishing based on the correct answer.
4. None of the above: Incorrect, as B is a phishing email.
Conclusion: Email B is a phishing email due to its impersonation of Apple, generic greeting, unsolicited verification request with a link, and use of urgency to prompt action. Since A is not shown but implied to be non-phishing, the correct answer is 'Only B.'

Answer:

C

Explanation:

In Linux/Unix, the SUID (Set User ID) bit allows a program to execute with the owner's permissions, typically root, rather than the caller's. It's denoted by an s in the user execute field (e.g., -rwsr-xr-x). Common SUID programs perform privileged tasks requiring temporary elevation.

Analysis:

C . /usr/bin/passwd:

Purpose: Updates user passwords in /etc/shadow (root-owned, 0600 perms).

Permissions: Typically -rwsr-xr-x, owned by root. The SUID bit lets non-root users modify shadow securely.

Command: ls -l /usr/bin/passwd confirms SUID (s in user execute).

A . /bin/ls:

Purpose: Lists directory contents, no privileged access needed.

Permissions: -rwxr-xr-x (no SUID). Runs as the calling user.

B . /usr/bin/curl:

Purpose: Transfers data over HTTP/FTP, no root privileges required by default.

Permissions: -rwxr-xr-x (no SUID).

Technical Details:

SUID Bit: Set via chmod u+s <file> or chmod 4755.

Security: SUID binaries are audited (e.g., find / -perm -u=s) due to escalation risks if writable or poorly coded (e.g., buffer overflows).

Security Implications: CNSP likely highlights SUID as an attack vector (e.g., CVE-1996-0095 exploited passwd flaws). Hardening removes unnecessary SUID bits.

Why other options are incorrect:

A, B: Lack SUID; no privileged operations.

D: Incorrect, as /usr/bin/passwd is a SUID example.

Real-World Context: SUID on /bin/su or /usr/bin/sudo similarly enables privilege escalation, often targeted in exploits.

P2P exams

Question 12

Question Type: MultipleChoice

What RID is given to an Administrator account on a Microsoft Windows machine?

Options:

A- 0
B- 500
C- 501
D- 100

Answer:

B

P2P exams

Explanation:

In Windows, security principals (users, groups) are identified by a Security Identifier (SID), formatted as S-1--<domain>-<RID>. The RID (Relative Identifier) is the final component, unique within a domain or machine. For local accounts:

RID 500: Assigned to the built-in Administrator account on every Windows machine (e.g., S-1-5-21-<machine>-500).

Created during OS install, with full system privileges.

Disabled by default in newer Windows versions (e.g., 10/11) unless explicitly enabled.

RID 501: Guest account (e.g., S-1-5-21-<machine>-501), limited access.

Technical Details:

Stored in SAM (C:\Windows\System32\config\SAM).

Enumeration: Tools like wmic useraccount or net user reveal RIDs.

Domain Context: Domain Admins use RID 512, but the question specifies a local machine.

Security Implications: RID 500 is a prime target for brute-forcing or pass-the-hash attacks (e.g., Mimikatz). CNSP likely advises renaming/disabling it (e.g., via GPO).

Why other options are incorrect:

A . 0: Reserved (e.g., Null SID, S-1-0-0), not a user RID.

C . 501: Guest, not Administrator.

D . 100: Invalid; local user RIDs start at 1000 (e.g., custom accounts).

Real-World Context: Post-compromise, attackers query RID 500 (e.g., net user Administrator) for privilege escalation.

Question 13

Question Type: MultipleChoice

Which option best is not a DDoS attack?

Options:

A- SYN Flood
B- NTP Amplification
C- UDP Flood
D- Brute Force

Answer:

D

Explanation:

DDoS (Distributed Denial of Service) attacks aim to overwhelm a target's resources with

excessive traffic, disrupting availability, whereas other attack types target different goals.

Why D is correct: Brute force attacks focus on guessing credentials (e.g., passwords) to gain unauthorized access, not on denying service. CNSP classifies it as an authentication attack, not a DDoS method.

Why other options are incorrect:

A: SYN Flood exhausts TCP connection resources, a classic DDoS attack.

B: NTP Amplification leverages amplified responses to flood targets, a DDoS technique.

C: UDP Flood overwhelms a system with UDP packets, another DDoS method.

P2P exams

P2P exams

To Get Premium Files for CNSP Visit

https://www.p2pexams.com/products/cnsp

For More Free Questions Visit

https://www.p2pexams.com/the-secops-group/pdf/cnsp

