CERTCRACK PRACTICE EXAMS

CompTIA Cybersecurity Analyst (CySA+) (CS0-004)

Free sample — 30 questions from the full CertCrack CySA+ (CS0-004) practice bank, with detailed explanations.

Full bank: 750 original practice questions + 10 performance-based questions
Domain 1: Security Operations (34%) · Domain 2: Vulnerability Management (26%)
Domain 3: Incident Response and Management (24%) · Domain 4: Reporting and Communication (16%)

Copyright & Disclaimer

CompTIA Cybersecurity Analyst (CySA+) CS0-004 — Free Sample (30 Questions)

Published by CertCrack · certcrack.com

Copyright © 2026 CertCrack. All rights reserved.

All questions in this book are 100% original practice material written by CertCrack. They are not real exam questions and are not affiliated with or endorsed by CompTIA, Inc.

Unofficial practice material — not affiliated with or endorsed by CompTIA, Inc.

The performance-based questions (PBQs) use simplified, fictional exhibits written for this book. Tool-output samples shown anywhere in this book are illustrative examples, not captures of real product screens.

This free sample contains 20 questions drawn from the full bank plus 10 of the new tool-focused questions. The complete bank — 750 questions and 10 performance-based questions — is available at certcrack.com and on Udemy.

Domain 1: Security Operations (34%)

Q1. [cysa-001] A ransomware outbreak encrypted hosts in the finance VLAN but did not reach the HR or engineering VLANs, even though all three share the same core switches. Which architecture concept MOST limited the blast radius?

A. A single flat Layer 2 domain for all hosts
B. Network segmentation with ACLs between VLANs
C. 802.1X authentication on access ports
D. An intrusion prevention system in detection mode

> Correct answer: B Explanation: Correct: segmentation with inter-VLAN ACLs confines traffic so malware cannot move laterally into other segments. A flat Layer 2 domain would have let it spread freely. 802.1X controls who connects, not traffic between already-connected segments. An IPS in detection mode only logs and does not block.

Q2. [cysa-010] A SOAR playbook is triggered by a phishing alert: it isolates the host, blocks the sender domain at the email gateway, and opens a ticket. Which efficiency benefit does this automation PRIMARILY provide?

A. Reduced mean time to respond through consistent, immediate containment
B. Elimination of the need for any analyst review of future alerts
C. Guaranteed prevention of all future phishing attempts against the company
D. Replacement of the SIEM as the primary detection platform

> Correct answer: A Explanation: Correct: SOAR's primary gain is faster, consistent containment (lower MTTR) for known patterns. Automation does not eliminate analyst review for judgment calls, cannot guarantee prevention, and complements rather than replaces the SIEM.

Q3. [cysa-025] EXHIBIT — SIEM excerpt (times UTC, same workstation WKSTN-1147): 10:02:11 Outlook opened attachment Invoice.docm; macro execution allowed by user 10:02:58 powershell.exe -enc <base64 blob> spawned from WINWORD.EXE 10:03:40 powershell.exe connected to 185.220.101.47:443 (no prior DNS query) 10:04:05 rundll32.exe loaded C:\Users\jdoe\AppData\Local\Temp\upd.dll 10:05:12 credential access alert: LSASS memory read by upd.dll process Based on this sequence, which conclusion is BEST supported?

A. The host is performing a routine Office update
B. The host is compromised: a malicious macro led to code execution, C2 contact, and credential access
C. The alerts are false positives from antivirus scanning
D. The user intentionally installed a software update

> Correct answer: B Explanation: Correct: the chain — macro execution, encoded PowerShell from Word, bare-IP C2 contact, DLL sideloading, then LSASS access — is a coherent compromise sequence. Routine Office updates do not use encoded PowerShell from Word or read LSASS. Antivirus scanning does not spawn encoded PowerShell or contact bare IPs. No legitimate update behaves this way.

Q4. [cysa-050] According to the pyramid of pain, which indicator type is HARDEST for an adversary to change?

A. File hashes, which change with every recompiled binary
B. TTPs (tactics, techniques, procedures)
C. IP addresses, which attackers rotate frequently
D. Domain names, which are cheap for attackers to register

> Correct answer: B Explanation: Correct: the pyramid of pain ranks attacker costs; TTPs are hardest to change because they require retooling behavior. File hashes change with every recompile, IP addresses rotate freely, and domain names are cheap to register.

Q5. [cysa-077] An analyst needs to determine how malware entered a host and what it did after execution. Which data source is MOST valuable?

A. SIEM aggregation
B. Firewall logs
C. EDR telemetry
D. Vulnerability scanner output

> Correct answer: C Explanation: Correct: EDR provides endpoint process, file, and network telemetry needed for root-cause analysis. A firewall shows network flows only. A SIEM aggregates but does not generate endpoint telemetry. A vulnerability scanner finds weaknesses, not attack activity.

Q6. [cysa-100] Match each identity attack to its BEST description: 1. MFA fatigue 2. Kerberoasting 3. Session token theft 4. Pass-the-ticket

A. 1-D, 2-B, 3-A, 4-C
B. 1-A, 2-C, 3-B, 4-D
C. 1-C, 2-A, 3-D, 4-B
D. 1-B, 2-D, 3-C, 4-A

> Correct answer: C Explanation: Correct: the correct mapping is 1-C, 2-A, 3-D, 4-B. MFA fatigue spams push prompts; Kerberoasting requests RC4 service tickets; token theft reuses a stolen session token; pass-the-ticket reuses a Kerberos ticket without the password.

Q7. [cysa-701] An EDR alert flags suspicious SMB traffic from a workstation at CertCrack.com. Before isolating the host, the analyst wants to confirm exactly which services and versions the host is running and which ports are actually listening. Which action gives the analyst this information?

A. Run a vulnerability scan template check for one named CVE against the host
B. Run an Nmap service and version detection scan against the host
C. Review the host's entries in Zeek's conn.log for the last hour
D. Sweep the subnet with Angry IP Scanner to list live hosts

> Correct answer: B Explanation: Correct: Nmap's service and version detection probes each open port and fingerprints the service and version behind it — exactly the enumeration the analyst needs before deciding on isolation. A template check such as Nuclei's verifies one specific known issue; it does not enumerate all listening services and versions. Zeek's conn.log records connections the host made or received, but it does not fingerprint service versions on listening ports. Angry IP Scanner gives a fast live-host list with a few common ports; it performs no service or version fingerprinting.

Q8. [cysa-704] CertCrack.com is investigating a suspected DNS-tunnelling channel. The responder needs two weeks of historical DNS queries and connection records for one workstation, but full packet captures are kept for only 48 hours. Which source gives the responder what they need?

A. A new Wireshark capture started today
B. Nikto's web server findings for the workstation
C. The YARA match report from the antivirus console
D. Zeek's dns.log and conn.log records for the period

> Correct answer: D Explanation: Correct: Zeek is a network security monitor that writes structured logs — dns.log for queries, conn.log for connections — designed to be retained and queried long after raw packets are discarded, which is exactly the historical reconstruction this investigation needs. A Wireshark capture started today cannot recover traffic from two weeks ago. Nikto reports web server files and misconfigurations, not workstation DNS history. YARA matches describe files that match a pattern; they contain no record of the host's DNS lookups.

Q9. [cysa-712] CertCrack.com is migrating from one SIEM platform to another and does not want to rewrite its library of log detections a second time. The detection engineer writes each detection once, in a vendor-neutral format, and converts it to each platform's query language. Which format is the engineer using?

A. EVE JSON, Suricata's event output schema
B. YARA, converted per platform by the SIEM vendor
C. Sigma, the generic signature format for SIEM detection rules
D. NSE, Nmap's scripting format

> Correct answer: C Explanation: Correct: Sigma is an open, vendor-neutral format for log detection rules, with converters that translate one rule into the query languages of many SIEM platforms — write once, run anywhere. EVE JSON is the format of Suricata's output events, not a portable way to author detections. YARA patterns match files and memory, not log events, and no per-SIEM conversion ecosystem exists for it. NSE is Nmap's scripting engine for network probes; it has nothing to do with SIEM detection content.

Q10. [cysa-715] After a detection-engineering sprint, CertCrack.com's SOC wants to verify that its EDR now alerts on one specific MITRE ATT&CK technique — OS credential dumping — before the technique is added to the coverage report. The test must be small, controlled, and repeatable. Which approach is BEST?

A. Run a Caldera operation across the test range
B. Run a full Nessus credentialed scan of the test host
C. Run the matching Atomic Red Team test for that technique
D. Sweep the test range with Masscan to confirm the host is reachable

> Correct answer: C Explanation: Correct: Atomic Red Team is a library of individual tests, each mapped to a single ATT&CK technique, designed for validating one detection at a time in a controlled way — precisely this task. A Caldera operation chains many techniques into an automated campaign, which is broader than the single-technique proof required and muddies the result. A Nessus scan checks for vulnerabilities, not whether a detection fires on an emulated technique. Masscan only reports open ports; it says nothing about EDR behaviour.

Domain 2: Vulnerability Management (26%)

Q11. [cysa-250] Which two factors should drive how often a system is vulnerability scanned? (Choose two.)

A. The system's exposure and the criticality of the data it handles.
B. The vendor's marketing schedule for new scanner features.
C. The number of employees in the department that owns the system.
D. Regulatory or contractual requirements that mandate a scan cadence.
E. The physical location of the data center where the system sits.

> Correct answer: A and D Explanation: Correct: exposure and data criticality determine risk, and compliance obligations often set minimum frequencies. The other options are wrong: vendor marketing, headcount, and data center geography do not determine scan cadence.

Q12. [cysa-275] An analyst notices that scan results for a subnet vary widely between runs, with hosts appearing and disappearing. What is the most likely explanation?

A. The subnet contains DHCP-assigned or ephemeral addresses, so the same IP may represent different hosts across runs.
B. The scanner is randomly dropping results.
C. The subnet is under active attack during every scan.
D. The hosts are all powered off between scans.

> Correct answer: A Explanation: Correct: DHCP churn and ephemeral workloads make IP-based tracking inconsistent across runs. The other options are wrong: scanners do not randomly drop results; an attack is not the likely cause of routine variance; and powered-off hosts would not appear at all.

Q13. [cysa-300] An analyst sees a finding for a vulnerability that was already remediated and verified last month. What is the most likely cause of its reappearance?

A. A system was restored from an old image or backup, or a new identical host was deployed from an unpatched template.
B. The scanner is generating random findings.
C. The remediation was never actually performed.
D. The vulnerability mutated into a new variant.

> Correct answer: A Explanation: Correct: restores from old images and unpatched templates are the classic cause of regressed findings. The other options are wrong: scanners do not generate random findings; the remediation was verified; and vulnerabilities do not mutate.

Q14. [cysa-330] A remediation team consistently misses SLAs for one application because its owner never approves maintenance windows. What is the best escalation path?

A. Present the accumulated risk to the application owner
B. Silently extend the SLAs for that application.
C. Close the findings as accepted without owner involvement.
D. Stop scanning the application.

> Correct answer: A Explanation: Correct: risk escalation with a deadline and a formal acceptance option forces an accountable decision. The other options are wrong: silent SLA extension hides the problem; closing without the owner is unauthorized; and stopping scans removes visibility.

Q15. [cysa-360] A vulnerability affects a system that is already scheduled for a major upgrade next month, which will replace the vulnerable component. What is the best interim handling?

A. Do nothing, since the upgrade will fix it.
B. Close the finding now in anticipation of the upgrade.
C. Patch the old component at full cost this week.
D. Apply mitigations and monitoring now

> Correct answer: D Explanation: Correct: interim mitigations cover the gap, and closure must wait for proof the upgrade removed the flaw. The other options are wrong: doing nothing leaves a month of exposure; premature closure is unverified; and full patching may waste effort days before replacement.

Q16. [cysa-725] A critical CVE in a widely used web framework was published this morning. CertCrack.com needs to know, within the hour, whether any of its 500 web servers is exposed to this one specific vulnerability. Which approach is fastest?

A. Schedule full credentialed Nessus scans of all 500 servers and review the results tomorrow
B. Run Nikto against each server and read every hygiene finding for a mention of the CVE
C. Point ScoutSuite at the web tier and review its flagged rules
D. Run Nuclei with the template written for that CVE against all 500 targets

> Correct answer: D Explanation: Correct: Nuclei's template model exists for this exact moment — one YAML template describes how to check for the specific CVE, and the scanner runs that single check quickly across hundreds of targets. Full credentialed scans are the deepest option but the slowest; they cannot deliver a one-CVE answer across 500 hosts within the hour. Nikto reports server hygiene issues such as dangerous files and outdated software; it is not a targeted CVE-verification tool. ScoutSuite audits cloud account configuration, not the vulnerability state of running web servers.

Q17. [cysa-733] CertCrack.com is preparing for a cloud security review of its production AWS environment. The security lead wants a broad posture audit of the live accounts — which storage buckets are public, which IAM policies are over-privileged, and which services have logging disabled — presented as a visual report the team can work through service by service. Which tool is the BEST fit?

A. Prowler
B. ScoutSuite
C. Trivy
D. Checkov

> Correct answer: B Explanation: Correct: ScoutSuite audits live cloud accounts through the provider APIs and produces an HTML dashboard of flagged rules per service — exactly the broad, visual posture review described. Prowler also audits live accounts, but its strength is benchmark-mapped PASS/FAIL compliance evidence (for example against CIS), not the service-by-service posture dashboard asked for here. Trivy scans artefacts such as container images and repositories for CVEs and secrets; it does not audit a live AWS account's configuration. Checkov scans infrastructure-as-code before deployment, and this environment is already running.

Q18. [cysa-737] CertCrack.com's CI pipeline is about to publish a new container image for the payments API. Policy states the build must fail if the image contains any CRITICAL vulnerability or an embedded secret, and the release record must include a software bill of materials. Which tool enforces this in the pipeline?

A. Nessus
B. OpenVAS
C. ScoutSuite
D. Trivy

> Correct answer: D Explanation: Correct: Trivy scans container images for known vulnerabilities by severity, exposed secrets, and misconfigurations, and can generate an SBOM — all three pipeline requirements in one tool. Nessus and OpenVAS are host and network vulnerability scanners; neither is the image-scanning gate in a build pipeline, and neither produces the image SBOM. ScoutSuite audits live cloud accounts through provider APIs and never inspects the contents of a container image.

Q19. [cysa-739] A developer at CertCrack.com opens a pull request whose Terraform would create a storage bucket with public access enabled. The security team wants this caught automatically before merge, with the failing policy named in the pull request. Which tool provides this gate?

A. ScoutSuite, scheduled to audit the account after the next deployment
B. Checkov, scanning the Terraform code in CI and failing the pull request on the public-access policy
C. Prowler, run manually by the cloud team each Friday
D. Trivy, scanning the bucket after it has been created

> Correct answer: B Explanation: Correct: Checkov performs static analysis of infrastructure-as-code — it reads the Terraform in the pull request, evaluates it against policies such as 'no public buckets', and fails the CI check before anything is deployed. ScoutSuite audits accounts that already exist; the public bucket would be live before the scheduled audit noticed it. A weekly manual run cannot gate a specific pull request and happens long after merge. Trivy scans artefacts such as images and repositories; by the time a created bucket could be reviewed, the exposure the gate exists to prevent has already happened.

Domain 3: Incident Response and Management (24%)

Q20. [cysa-430] MATCHING PBQ. Match each observed activity to the MITRE ATT&CK tactic it best represents. Activities: 1. An attacker queries Active Directory for all domain administrators. 2. Malware writes a registry Run key so it starts at logon. 3. A script dumps LSASS process memory. 4. Stolen documents are uploaded to an external file-sharing site. Which mapping is correct?

A. 1-Discovery, 2-Persistence, 3-Credential Access, 4-Exfiltration
B. 1-Reconnaissance, 2-Defense Evasion, 3-Discovery, 4-Command and Control
C. 1-Collection, 2-Privilege Escalation, 3-Credential Access, 4-Impact
D. 1-Discovery, 2-Persistence, 3-Defense Evasion, 4-Collection

> Correct answer: A Explanation: Correct: Querying AD for administrators is Discovery (learning about the environment); a Run key is Persistence (surviving restart); LSASS dumping is Credential Access; uploading stolen documents externally is Exfiltration. The second mapping is wrong because the AD query happens inside the network (Discovery, not external Reconnaissance) and a Run key is Persistence, not Defense Evasion. The third is wrong because the AD query is Discovery rather than Collection, and uploading data out is Exfiltration, not Impact. The fourth is wrong because LSASS dumping targets credentials, not evasion, and external upload is Exfiltration, not Collection.

Q21. [cysa-460] A threat intelligence report describes an adversary that consistently uses spear-phishing for initial access, then credential dumping, then lateral movement via SMB. An analyst maps this to the adversary's typical playbook. Which framework best supports describing this repeated sequence of tactics?

A. MITRE ATT&CK, because it catalogs tactics and techniques in the order adversaries use them
B. The Diamond Model, because it sequences attacks into phases from adversary action through victim impact
C. The Cyber Kill Chain, because it lists adversary tools
D. CVSS, because it scores vulnerability severity

> Correct answer: A Explanation: Correct: ATT&CK catalogs adversary tactics and techniques, making it the right framework for describing a repeated tactic sequence. The Diamond Model is wrong because it models adversary-capability-infrastructure-victim relationships, not tactic sequences. The Kill Chain is wrong as the best fit because it describes high-level phases, not the technique-level playbook. CVSS is wrong because it scores vulnerabilities, not adversary behavior.

Q22. [cysa-490] A responder is documenting an incident timeline. The responder records each event with a timestamp, source, and description, and notes where timestamps came from different time zones. Why is time-zone normalization important?

A. Normalization slows documentation, so it should be skipped
B. Events from systems in different zones appear out of order unless converted to a common reference, which corrupts the sequence
C. Time zones do not affect timelines, because all systems use the same clock and synchronize to the same time source, so offsets never occur in practice
D. Normalization is only needed for legal cases, not internal reviews

> Correct answer: B Explanation: Correct: Timeline reconstruction requires a common time reference; unconverted zone offsets misorder events and mislead analysis. Same clock is wrong because systems commonly run in local zones. Legal-only is wrong because accurate sequencing matters for all analysis. Skipping is wrong because an incorrect timeline is worse than a slower correct one.

Q23. [cysa-520] YES/NO PBQ. For each statement about evidence handling, select Yes if it follows proper procedure or No if it does not. Statements: 1. Original evidence should be analyzed directly rather than working from a forensic image. 2. Every transfer of evidence must be recorded with date, time, sender, and receiver. 3. A hash of the evidence should be recorded at acquisition and verified before analysis. 4. Evidence may be stored on any available shared drive as long as it is password-protected. Which answer sequence is correct?

A. 1-No, 2-Yes, 3-Yes, 4-No
B. 1-Yes, 2-Yes, 3-No, 4-Yes
C. 1-No, 2-No, 3-Yes, 4-Yes
D. 1-Yes, 2-No, 3-No, 4-No

> Correct answer: A Explanation: Correct: Statement 1 is No because analysts work from forensic images to preserve originals; statement 2 is Yes because chain of custody requires every transfer to be documented; statement 3 is Yes because hashing at acquisition and re-verification proves integrity; statement 4 is No because evidence requires access-controlled, integrity-monitored storage, not any password-protected share. The second sequence is wrong because it approves analyzing originals, rejects hashing, and approves casual storage. The third is wrong because it rejects transfer documentation and approves casual storage. The fourth is wrong because it rejects transfer documentation and hashing.

Q24. [cysa-550] A responder is handling an incident involving a compromised email account. The responder checks the account's login history, mailbox rules, and sent items. Which investigation principle does this sequence follow?

A. Comprehensive account review: checking access, persistence mechanisms, and actions taken by the attacker
B. Minimal review, because only login history matters
C. Random review, because order does not matter
D. Recovery review, because the account is being restored to service and its settings are being reset to defaults

> Correct answer: A Explanation: Correct: A compromised account review must cover how the attacker entered (login history), how they persisted (mailbox rules), and what they did (sent items). Minimal review is wrong because it misses persistence and actions. Random review is wrong because the sequence is systematic, not random. Recovery review is wrong because this is investigation, not restoration.

Q25. [cysa-742] An incident at CertCrack.com began with exploitation of a known CVE on an Internet-facing server. Before finalising the root-cause report, the IR lead wants technical confirmation that the CVE is genuinely exploitable on that server's exact software version — not merely listed by a scanner. Which action provides that confirmation?

A. Run the corresponding Metasploit exploit module against a clone of the server in the lab and observe whether it yields a session
B. Re-run the vulnerability scan with a second scanner and compare the two reports

C. Check the CVE's CVSS score and record it as proof of exploitation
D. Ask the server's administrator whether the patch was applied

> Correct answer: A Explanation: Correct: Metasploit exists to prove exploitability — running the CVE's module against a lab clone of the server shows definitively whether that version can be exploited, turning a scanner's inference into confirmed root-cause evidence. Two scanners agreeing is still inference from banners and versions; neither attempts the exploit, so neither confirms it. A CVSS score measures potential severity in general, not whether this specific build is exploitable. An administrator's recollection is not technical evidence, and patch records are exactly what the investigation should verify independently.

Domain 4: Reporting and Communication (16%)

Q26. [cysa-600] An analyst includes screenshots in an incident report as evidence of attacker activity. What makes those screenshots most useful as documentation?

A. Capturing them in high resolution regardless of content
B. Including visible timestamps, hostnames, and a caption explaining what each image shows
C. Cropping out all system information to keep the images clean and presentable for executives
D. Saving them only on the analyst's personal workstation

> Correct answer: B Explanation: Correct: timestamps, hostnames, and explanatory captions tie each screenshot to a specific system and moment, which is what makes it usable evidence. High resolution alone does not help if the image cannot be placed in context. Cropping out system information removes the very details that establish relevance. Storing evidence only on a personal workstation breaks chain of custody and availability.

Q27. [cysa-630] A SOC report cites an MTTD of 4 hours. What does that metric measure?

A. The average time from when an incident begins until it is detected
B. The average time to close a ticket
C. The average time to restore service after an incident
D. The average time between a user report and analyst assignment, from ticket creation

> Correct answer: A Explanation: Correct: mean time to detect measures the interval from when an incident begins to when it is detected, which is the definition of MTTD. Time to close a ticket measures workflow speed, not detection. Time to restore service is recovery, captured by MTTR. Time between report and analyst assignment is triage delay, not detection time.

Q28. [cysa-660] A retailer discovers cardholder data was exposed. Who must be notified first under PCI DSS obligations?

A. Notify the acquiring bank and card brands promptly, as required by PCI DSS and contractual obligations
B. Notify only internal IT staff
C. Wait until the next annual PCI audit
D. Notify customers directly and no one else

> Correct answer: A Explanation: Correct: PCI DSS and card-brand contracts require prompt notification to the acquiring bank and card brands, because they manage fraud risk across the payment ecosystem. Notifying only internal IT is insufficient. Waiting for the annual audit is far too late. Customers are notified through the card brands and issuers per the process, not by the merchant alone.

Q29. [cysa-690] Which two signs show a lessons-learned program is working? Choose two.

A. The repeat-incident rate for the same weakness declines
B. Incident reports get longer
C. Action items from reviews are closed on time, verified as done
D. Every review produces the same recommendations as the last

Correct answer: A and C Explanation: Correct: a falling repeat-incident rate shows the same weaknesses are not recurring, and on-time closure of action items shows lessons became changes. Longer reports measure verbosity, not learning. Identical recommendations every time indicate nothing was actually learned or fixed.

Q30. [cysa-747] An external auditor asks CertCrack.com for evidence of how its Azure subscriptions score against recognised cloud security benchmarks, control by control, in a format the audit team can import into its own tooling. Which tool should the security team run?

A. Prowler
B. ScoutSuite
C. Nikto
D. Nuclei

Correct answer: A Explanation: Correct: Prowler's checks are mapped to benchmarks and compliance frameworks such as CIS, and it reports PASS/FAIL per control with exportable output (CSV, JSON, compliance reports) — the auditor-ready evidence the scenario requires. ScoutSuite gives a broad posture overview in an HTML report, but benchmark-by-benchmark evidence is Prowler's speciality. Nikto scans web servers for dangerous files and misconfigurations; it knows nothing about cloud subscriptions. Nuclei verifies template-described issues on running targets; it does not assess cloud benchmark compliance.