PREPKLOUD FIELD GUIDE SERIES

CompTIA CySA+ CS0-004 Field Guide

An analyst-first guide to security operations, vulnerability management, incident response, and reporting.

First edition · 2026 · prepkloud.com

About this guide

CompTIA Cybersecurity Analyst+ (CySA+) validates the skills of a security analyst: turning telemetry into findings, prioritizing vulnerabilities by real risk, handling incidents, and communicating clearly to technical and business audiences. Version 4, exam **CS0-004**, is the current version.

Exam facts, as published by CompTIA when this guide was written: launched 23 June 2026; up to 85 questions, including multiple-choice and performance-based questions; 165 minutes; passing score 750 on a 100–900 scale; about four years of hands-on experience as a SOC or vulnerability analyst recommended. Always confirm current details on the official CompTIA CySA+ page and the official CS0-004 exam objectives before you schedule.

CS0-004 domain	Weight	Chapters
1. Security Operations	34%	2–6
2. Vulnerability Management	26%	7–9
3. Incident Response and Management	24%	10–12
4. Reporting and Communication	16%	13

or copied exam items. Practise only in environments you own or are explicitly authorized to test.

Contents

1. How CS0-004 works and how to prepare

2. Visibility architecture and telemetry quality

3. Log analysis, SIEM, and EDR

4. Network, email, and identity indicators

5. Threat intelligence and threat hunting

6. Automation, SOAR, and AI in the SOC

7. Vulnerability assessment methods

8. Analysing and prioritizing findings

9. Mitigation, controls, and exceptions

10. Incident response frameworks and preparation

11. Detection, analysis, and evidence

12. Containment, eradication, recovery, and lessons learned

13. Reporting and communication

14. Original practice questions and explanations

15. Six-week plan, cheat sheet, and glossary

CHAPTER 1 · EXAM METHOD

How CS0-004 works and how to prepare

CySA+ is an intermediate, hands-on exam. Many questions present evidence, such as a log excerpt, a scan result, packet details, or an alert timeline, and ask what it means or what to do next. Performance-based questions may ask you to classify findings, order response steps, or pick the right tool for a task. You need judgment, not only definitions.

The analyst's five questions

Analysts often go wrong by mixing up five separate questions. Keep them apart:

1. **Evidence:** what was actually observed?
2. **Coverage:** what could have been observed, and were the sources healthy?
3. **Confidence:** how strongly does the evidence support the conclusion?
4. **Authority:** who is allowed to act, and under what procedure?
5. **Outcome:** did the action reduce risk without unacceptable business impact?

> **Exam mindset**
>
> An alert is evidence, not a verdict. A query with no results means something only if the data source was healthy. A CVSS score is an input to the business decision, not the decision itself. And a "success" response from a containment API does not prove recovery until you check the state again.

Question strategy

- Read the last sentence first. Is it asking for the *first*, *best*, or *most likely* answer?
- Consider where you are in the incident lifecycle. You do not eradicate before you have scoped and contained.
- Prefer answers that preserve evidence and follow authority and procedure.
- Flag long performance-based questions and return to them so they do not consume your time early.

> **Recall**
>
> - Can I name the four domains and their weights?
> - Can I explain the difference between evidence and coverage?

Visibility architecture and telemetry quality

Security operations begins with knowing where your visibility comes from. Map every source: identity provider, endpoints, DNS, proxy, firewall, network flows, email, SaaS applications, cloud control plane, containers, and business applications. For each, know how events are authenticated and transported, how long they take to arrive, how long they are kept, and how they are parsed.

Source	Answers the question
Identity provider	Who signed in, from where, with which factors, and was it risky?
EDR	Which processes ran, what they touched, and which network connections they made.
DNS and proxy	Which domains and URLs hosts tried to reach.
Flow logs / NetFlow	Who talked to whom, how much, and when, without payloads.
Cloud control-plane logs	Who changed cloud configuration and permissions.
Email gateway	Sender, authentication results, attachments, links, and delivery actions.

Telemetry health

Monitor the monitoring. Watch for missing agent heartbeats, ingestion lag, dropped events, parser failures, broken field mappings, and clock skew. Keep both the original event time and the ingestion time. Normalize fields to a common schema for correlation, but retain the raw event for evidence.

> **The silent parser**
>
> A detection for suspicious PowerShell has not fired in two weeks. The analyst checks coverage first and finds that an agent update changed the log format, so the parser has been dropping the command-line field. The fix covers three things: repair the parser, backfill and re-run the detection for the gap period, and add a health alert for that field.

Architecture concepts

Expect to reason about segmentation, zero trust, secure remote access, cloud and hybrid designs, and the security value and limits of each. Know the difference between on-premises, cloud, and SaaS visibility: in SaaS you may only have the logs the provider exposes.

Recall

- Why keep both event time and ingestion time?
- What should you check before trusting a zero-result query?

Log analysis, SIEM, and EDR

A SIEM centralizes, normalizes, correlates, and retains logs so analysts can search and detect across sources. EDR records detailed endpoint behaviour and can respond on the host, for example by isolating it, killing a process, or quarantining a file. XDR extends detection and response across endpoint, identity, email, and cloud signals.

Reading logs well

- **Windows:** know the common security event IDs, such as successful and failed logons (4624, 4625), special-privilege logons (4672), new processes (4688), new services (7045), and audit-log clearing (1102). Recognize logon types such as interactive, network, and remote interactive.
- **Linux:** authentication logs, sudo usage, cron changes, new SSH keys, and unexpected listening services.
- **Web servers:** HTTP methods, status codes, user agents, and patterns such as path traversal (`../`), SQL injection strings, or bursts of 401/403 responses.

Detection quality

Every detection trades false positives against false negatives. Tune using context: baselines, allow-lists that are reviewed, and correlation of independent signals. A single failed logon means

little. Many failed logons across many accounts from one source, followed by one success, suggests password spraying.

> **Behaviour over signatures**
>
> Signatures catch known-bad artefacts. Behavioural analytics catch unusual activity: a finance user running reconnaissance commands, a service account signing in interactively, or a host suddenly uploading gigabytes to a new domain. The strongest detections combine both.

> **Living off the land**
>
> EDR shows `winword.exe` spawning `powershell.exe` with an encoded command, which then connects to a newly registered domain. None of these binaries is malicious on its own. The parent-child relationship and the destination make it a high-confidence finding.

Recall

- What distinguishes EDR from SIEM?
- Which log pattern suggests password spraying rather than brute force?

Network, email, and identity indicators

Network indicators

- **Beaconing:** regular, repeated connections to the same destination at a fixed interval, sometimes with jitter.
- **DNS anomalies:** very long or high-entropy subdomains (possible DNS tunnelling), newly registered domains, or algorithmically generated domain names.
- **Data exfiltration:** unusual outbound volume, uploads to personal cloud storage, or traffic at odd hours.
- **Lateral movement:** workstation-to-workstation SMB, RDP, or remote management traffic that does not match normal administration.
- **Scanning:** one source touching many ports or hosts in a short period.

Packet captures show full payloads when they are available. Flow data scales better but shows only metadata. Tools such as Wireshark and tcpdump appear in scenarios. Know what a TCP three-way handshake looks like and what repeated SYNs without replies suggest.

Email indicators

Check sender authentication results: SPF (is the sending server authorized for the domain?), DKIM (is the message signed and unaltered?), and DMARC (does the domain's policy align, and what should happen on failure?). Look for display-name spoofing, look-alike domains, mismatched reply-to addresses, urgent payment requests, and links whose visible text differs from the real destination.

Identity indicators

Impossible travel, MFA fatigue (many push requests), new MFA methods registered right after a risky sign-in, consent grants to unfamiliar applications, and service accounts used interactively are all strong identity signals.

> **MFA fatigue**
>
> A user reports dozens of MFA prompts overnight and approved one to make them stop. The analyst treats the account as compromised: revoke sessions, reset credentials, review sign-ins and any mailbox rules or MFA methods added, and move to number-matching or phishing-resistant MFA.

Recall

- What do SPF, DKIM, and DMARC each verify?
- What pattern in DNS logs suggests tunnelling?

Threat intelligence and threat hunting

Threat intelligence is information about adversaries that has been evaluated and put in context so you can act on it. It comes in levels: **strategic** (trends and risk for leaders), **operational** (campaigns and actor intent), and **tactical** (indicators and techniques for defenders).

Indicators and their value

Hash values and IP addresses are easy for attackers to change. Domains are slightly harder. Tools, and above all tactics, techniques, and procedures (TTPs), are the hardest to change and the most valuable to detect. This is often described as the pyramid of pain.

Frameworks

- **MITRE ATT&CK:** a knowledge base of adversary tactics and techniques used to map detections and find coverage gaps.
- **Cyber Kill Chain:** the stages of an intrusion from reconnaissance to actions on objectives.
- **Diamond Model:** relates adversary, capability, infrastructure, and victim for each event.

- **STIX and TAXII:** a standard format for describing threat intelligence, and a protocol for sharing it.

Evaluate each source for reliability, timeliness, relevance, and confidence. An indicator with no context or age information can create noise.

Threat hunting

Hunting is proactive: start with a hypothesis ("an actor targeting our sector uses scheduled tasks for persistence"), identify the data needed, search, and record the result either way. A hunt that finds nothing is useful only if coverage was confirmed. Successful hunts should become automated detections.

Recall

- Why are TTPs more valuable than file hashes?
- What makes a hunt with no findings trustworthy?

Automation, SOAR, and AI in the SOC

SOAR (security orchestration, automation, and response) platforms run playbooks that enrich alerts, gather context, and take approved actions. Good automation removes repetitive work and speeds up response. Poor automation acts on weak evidence at machine speed.

What to automate

- **Enrichment:** reputation look-ups, asset owner, user details, and related alerts. Low risk, high value.
- **Triage:** de-duplication, grouping, and severity scoring.
- **Response:** blocking a known-bad indicator, isolating a host, or disabling an account. These need confidence thresholds, approval rules, and rollback steps.

Scripting (for example PowerShell, Python, and Bash), regular expressions, and APIs (REST with JSON) appear in scenarios. Read a short script and say what it does. Store API keys in a secrets vault, not in the script.

AI in security operations

AI assistants can summarize incidents, draft queries, and explain unfamiliar artefacts. Govern them like any other tool: decide which data they may process, keep a human in the loop for impactful

actions, verify outputs against the evidence, and protect AI-connected workflows from prompt injection and data leakage.

Process improvement

Measure the SOC with meaningful metrics, such as mean time to detect and to respond, detection coverage against ATT&CK, false-positive rates, and backlog age. Use them to guide improvement, not to reward closing tickets fast.

Recall

- Which SOAR actions are safe to automate fully, and which need approval?
- How should AI output be treated in an investigation?

Vulnerability assessment methods

Vulnerability management is a continuous cycle: discover assets, assess them, prioritize findings, remediate or mitigate, verify, and report. You cannot assess what you do not know about, so an accurate asset inventory comes first.

Method	Strength	Limitation
Unauthenticated network scan	Shows the attacker's external view	Misses local software and configuration issues
Authenticated (credentialed) scan	Accurate, detailed findings on installed software and settings	Needs managed credentials with the right privileges
Agent-based scanning	Works for remote and intermittently connected devices	Requires agent deployment and maintenance
Passive discovery	No impact on fragile systems (useful for OT/ICS)	Limited detail
Web application scanning (DAST)	Finds runtime web flaws	Can affect application data; coverage depends on crawling
Cloud posture and container image scanning	Finds misconfigurations and vulnerable packages before or after deployment	Needs integration with the pipeline and cloud APIs

Planning a scan

Define scope and authorization in writing, schedule around business-critical windows, rate-limit scans on fragile systems, and coordinate with operations. Special considerations apply to operational technology, medical devices, and legacy systems, where active scanning can cause outages.

Recall

- Why do credentialed scans produce fewer false positives?
- When is passive discovery the better choice?

Analysing and prioritizing findings

Validate first

Not every finding is real. A false positive reports a vulnerability that is not present, for example because a vendor backported a fix without changing the version string. A false negative misses a real vulnerability, often because a scan was unauthenticated or blocked. Validate important findings before escalating them.

Understanding CVSS

The Common Vulnerability Scoring System rates technical severity. Know the CVSS v3.1 base metrics: attack vector (network, adjacent, local, physical), attack complexity, privileges required, user interaction, scope, and impact on confidentiality, integrity, and availability. CVSS v4.0 keeps the same idea but adds an attack requirements metric and replaces scope with separate impact metrics for the vulnerable system and subsequent systems. A network-reachable flaw that needs no privileges and no user interaction is generally more urgent than a local one that needs administrator rights.

Contextual prioritization

CVSS measures severity, not your risk. Combine it with:

- **Exploitation evidence:** is it being exploited in the wild (for example, listed in CISA's Known Exploited Vulnerabilities catalogue)? What is the exploit prediction score (EPSS)?
- **Exposure:** is the asset internet-facing?
- **Asset criticality:** does it hold sensitive data or support critical services?
- **Compensating controls:** is the vulnerable feature disabled or already filtered?

> **Two findings, one patch window**
>
> Finding A has CVSS 9.8 on an isolated lab server with no sensitive data. Finding B has CVSS 7.5 on an internet-facing VPN appliance and is known to be exploited in the wild. Patch B first: exposure and active exploitation outweigh the higher base score.

Recall

- Can I list the CVSS base metrics?
- Which three context factors can change a finding's priority?

Mitigation, controls, and exceptions

Remediation removes the vulnerability, for example by patching, upgrading, or reconfiguring. Mitigation reduces its likelihood or impact when remediation is not immediately possible. Examples include disabling the vulnerable feature, restricting network access, applying a web application firewall rule, or adding monitoring.

Common vulnerability classes

- **Injection** (SQL, command, LDAP): validate and parameterize input.
- **Cross-site scripting:** encode output and use a content security policy.
- **Broken access control and insecure direct object references:** enforce authorization on the server for every request.
- **Server-side request forgery:** restrict outbound requests and protect metadata services.
- **Insecure deserialization, buffer overflows, and race conditions:** use safe libraries, memory-safe patterns, and patching.
- **Misconfiguration:** default credentials, open storage, unnecessary services, and weak TLS settings.

Control types

Controls can be preventive, detective, corrective, deterrent, or compensating. They can also be technical, managerial, operational, or physical. A compensating control is an alternative that meets the intent of a requirement when the primary control is not feasible.

Exceptions and risk acceptance

Sometimes a system cannot be patched, for example because of vendor support or uptime requirements. Document an exception with the business owner's risk acceptance, the compensating controls, an expiry date, and a review schedule. An exception without an expiry quietly becomes permanent.

Verification

Close a finding only after a rescan or other evidence confirms the fix. A change ticket marked "done" is not proof.

Recall

- What is the difference between remediation and mitigation?
- What must every risk exception include?

Incident response frameworks and preparation

Frameworks such as NIST's incident handling guidance describe a lifecycle: **preparation**; **detection and analysis**; **containment, eradication, and recovery**; and **post-incident activity**. Know the order and what belongs in each phase.

Preparation

- An approved incident response plan with roles, severity levels, escalation paths, and authority to act.
- Playbooks for common scenarios: phishing, ransomware, account compromise, data exposure, and insider activity.
- Contact lists for legal, communications, HR, executives, regulators, law enforcement, and key vendors.
- Tools and access ready in advance: forensic workstations, evidence storage, and responder accounts.
- Tabletop exercises and simulations to test the plan.

Severity and classification

Classify incidents by functional impact, information impact (was data exposed or altered?), and recoverability. Severity drives who is notified and how fast. Know that legal and regulatory

notification requirements may apply to data breaches, and that the legal team usually decides on them.

> **Event vs incident**
>
> An event is any observable occurrence. An incident is an event that violates, or threatens to violate, security policy. Not every alert is an incident. Analysis decides.

Recall

- Can I list the four phases of the NIST lifecycle in order?
- What should be decided before an incident rather than during it?

Detection, analysis, and evidence

Analysis confirms whether an incident has happened, how far it reaches, and how serious it is. Build a timeline, identify patient zero, map activity to ATT&CK techniques, and find every affected account, host, and data store before you act broadly.

Scoping questions

- What was the initial access vector?
- Which identities and hosts show related activity?
- Is there persistence, such as scheduled tasks, services, new accounts, mailbox rules, or cloud access keys?
- Was data accessed, staged, or exfiltrated?
- Is the attacker still active?

Evidence handling

Collect evidence in **order of volatility**: CPU and memory contents first, then network state and running processes, then disk, then logs and backups. Create forensic images, verify them with cryptographic hashes, and work on copies. Keep a **chain of custody** that records who collected each item, when, how, and every transfer. When litigation is possible, a **legal hold** preserves relevant data.

Ransomware note on one server

Before rebuilding, the analyst checks whether other hosts show the same encryption process, whether domain admin credentials were used, and whether data left the network before encryption. Rebuilding one server while the attacker still holds domain credentials would lead to reinfection.

Recall

- What is the order of volatility?
- Why hash forensic images?

Containment, eradication, recovery, and lessons learned

Containment

Containment stops the damage from spreading while you preserve evidence. Options include isolating hosts with EDR, segmenting networks, disabling or resetting accounts, revoking sessions and tokens, and blocking indicators. Choose in proportion to the incident, consider the business impact, and act under documented authority. Containing too early can alert the attacker and trigger destructive actions. Containing too late lets damage spread. Scoping guides the timing.

Eradication

Remove the attacker's foothold: malware, persistence mechanisms, backdoor accounts, malicious rules, and compromised credentials. Fix the root cause, whether that is an unpatched vulnerability, a misconfiguration, or a weak authentication method.

Recovery

Restore from known-good backups or rebuild from trusted images, patch before reconnecting, and return systems to service in stages with increased monitoring. Validate that security controls

are working and that business processes function. Watch for re-infection.

Post-incident activity

Hold a blameless lessons-learned review. Record the timeline, what worked, what did not, the root cause, and specific improvements with owners and deadlines, covering detections, playbooks, controls, and training. Retain evidence according to policy.

> **Exam trap**
>
> Reimaging a machine is eradication and recovery. It is not the first step. Unless the question says evidence has already been captured, preserve evidence and scope first.

Recall

- Why can early containment sometimes be harmful?
- What makes a lessons-learned review effective?

Reporting and communication

An analyst's work only reduces risk when the right people understand it and act. This domain tests whether you can match the message to the audience.

Audience	Needs
Executives and board	Business impact, risk trend, decisions required, cost and time. Avoid jargon.
System owners and IT	Specific affected assets, technical detail, remediation steps, deadlines.
Legal, compliance, and privacy	Facts about data involved, timelines, regulatory implications, evidence status.
Communications / PR	Confirmed facts that can be shared, what is still unknown, approved messaging.

Vulnerability reporting

Report trends, not only totals: open critical findings over time, mean time to remediate by severity, SLA compliance, recurring issues, and exceptions that are close to expiry. Separate the metrics you act on from vanity metrics.

Incident reporting

A good incident report includes an executive summary, a timeline, scope and impact, root cause, actions taken, the current status,

remaining risks, and recommendations with owners. Communicate confirmed facts, clearly labelled unknowns, options, and next steps. Never speculate as if it were fact.

Metrics and KPIs

Useful KPIs include mean time to detect (MTTD), mean time to respond (MTTR), percentage of assets scanned, patch SLA compliance, and detection coverage. Explain what changed and why, not only the number.

Escalation and stakeholder management

Follow the communication plan: who must be told, by whom, how fast, and through which channel. Use out-of-band channels if the normal ones might be compromised.

Recall

- What does an executive need that an engineer does not?
- Which elements belong in every incident report?

Practice questions and explanations

Answer each question before reading the explanation. These are original scenarios, not recalled CompTIA questions.

Security operations (1–10)

1. A detection has produced no alerts for two weeks after an agent update. What should the analyst check first?

Telemetry coverage and parsing health for that data source.
A query with no results is meaningful only if the data was collected and parsed correctly.

2. Logs show hundreds of accounts each failing to sign in once or twice from one IP, then one success. What is the most likely attack?

Password spraying.
Brute force targets one account with many passwords; spraying tries a few passwords across many accounts.

3. EDR shows a word processor spawning PowerShell with an encoded command. Why is this high-confidence?

The parent-child process relationship is unusual and commonly associated with malicious documents.
Both binaries are legitimate; the behaviour is the indicator.

4. A host connects to the same external IP every 60 seconds with small, similar payloads. What does this suggest?

Command-and-control beaconing.
Regular timing and uniform size are typical beacon characteristics.

5. DNS logs show many queries with very long, random-looking subdomains of one domain. What is the likely activity?

DNS tunnelling or data exfiltration over DNS.
High-entropy, long labels encode data inside queries.

6. An email passes SPF but fails DMARC alignment. What does this indicate?

The sending server was authorized for some domain, but not one aligned with the visible From domain.
DMARC checks alignment with the domain the user sees.

7. Which indicator type is hardest for an attacker to change?

Tactics, techniques, and procedures (TTPs).
Hashes and IP addresses are trivial to change.

8. Which standard is used to describe threat intelligence in a structured format?

STIX.
TAXII is the transport protocol for sharing it.

9. A threat hunt finds nothing. What makes that result useful?

Documented confirmation that the relevant data sources were complete and healthy during the hunt window.
Without coverage, "nothing found" proves nothing.

10. Which SOAR action is safest to fully automate?

Alert enrichment, such as reputation, asset, and user look-ups.
Disruptive actions like disabling accounts need confidence thresholds and approval.

Vulnerability management (11–18)

11. Why do credentialed scans usually produce more accurate results?

They can inspect installed software, patches, and configuration directly.
Unauthenticated scans infer versions from external responses.

12. Which discovery method is safest for fragile industrial control systems?

Passive discovery.
Active scanning can disrupt some OT devices.

13. A scanner flags an old version string, but the vendor has backported the fix. What is this finding?

A false positive.
Validate with vendor advisories or package details before escalating.

14. Which CVSS base metric describes whether an attacker must already have an account on the system?

Privileges required.
Attack vector describes where the attacker must be; user interaction describes whether a victim must act.

15. A CVSS 7.5 flaw on an internet-facing appliance is known to be exploited in the wild. A CVSS 9.8 flaw is on an isolated lab host. Which should be fixed first?

The exploited, internet-facing 7.5 finding.
Exposure and active exploitation outweigh a higher base score.

16. A critical system cannot be patched for 60 days. What is the correct handling?

A documented exception with owner risk acceptance, compensating controls, an expiry date, and review.
Undocumented delays leave risk unmanaged.

17. Which control best mitigates SQL injection in application code?

Parameterized queries (prepared statements) with input validation.
A WAF is a useful compensating layer, not the root fix.

18. When should a vulnerability finding be closed?

After a rescan or other evidence confirms the fix.
A completed change ticket alone is not proof.

Incident response (19–26)

19. What is the first phase of the NIST incident response lifecycle?

Preparation.
It is followed by detection and analysis; containment, eradication, and recovery; and post-incident activity.

20. What is the difference between an event and an incident?

An incident is an event that violates or threatens to violate security policy.
Analysis decides whether an event is an incident.

21. Which evidence should be collected first?

The most volatile, such as memory contents and running processes.
Disk and logs persist longer.

22. Why is a hash calculated for a forensic image?

To prove the image has not changed since collection.
It supports integrity and chain of custody.

23. A user approved one of many unexpected MFA prompts. What should happen first?

Treat the account as compromised: revoke sessions and reset credentials, then investigate.
Also review added MFA methods, mailbox rules, and recent activity.

24. Before reimaging an infected server, what should the analyst ensure?

Evidence is preserved and the scope, including other affected hosts and credentials, is understood.
Otherwise the attacker may return using remaining access.

25. Which activity belongs to eradication rather than containment?

Removing persistence mechanisms and fixing the root cause.
Containment limits spread; eradication removes the foothold.

26. What is the main output of a lessons-learned review?

Specific improvements with owners and deadlines.
A review that assigns blame without actions does not reduce future risk.

Reporting and communication (27–30)

27. What should an executive incident summary emphasize?

Business impact, current risk, decisions needed, and next steps, in plain language.
Technical detail belongs in the appendix or the technical report.

28. Which metric best shows whether vulnerability remediation is improving?

Mean time to remediate by severity, tracked over time against SLA.
Total finding counts alone can hide trends.

29. During an incident, a manager asks the analyst to guess the root cause for a press statement. What should the analyst provide?

Confirmed facts, clearly labelled unknowns, and when more information is expected.
Speculation presented as fact creates legal and reputational risk.

30. Email may be compromised during an incident. How should responders communicate?

Use an out-of-band channel defined in the communication plan.

Attackers may be reading the normal channels.

Six-week plan, cheat sheet, and glossary

Week	Focus	Output
1	Telemetry, log sources, SIEM, and EDR	A source map with health checks for your lab or workplace.
2	Network, email, identity indicators; threat intelligence and hunting	Three hunting hypotheses with the data each needs.
3	Automation, scripting, AI governance; start vulnerability management	Read and explain five short scripts; one SOAR playbook outline.
4	Scanning, CVSS, contextual prioritization, mitigation	Prioritize a ten-finding list and justify the order.
5	Incident response lifecycle and evidence handling	A phishing and a ransomware playbook; first timed practice set.
6	Reporting, metrics, mixed timed practice	An executive and a technical report for one scenario; book when practice is consistently strong.

Decision cheat sheet

- **Zero results:** check coverage before concluding.
- **Spraying vs brute force:** many accounts, few tries / one account, many tries.

- **Beaconing:** regular interval, uniform size, same destination.
- **SPF / DKIM / DMARC:** authorized sender / signed and unaltered / aligned policy.
- **Pyramid of pain:** TTPs > tools > artefacts > domains > IPs > hashes.
- **CVSS vs risk:** severity plus exploitation, exposure, and asset criticality.
- **Remediate vs mitigate:** remove the flaw / reduce its likelihood or impact.
- **IR order:** prepare, detect and analyse, contain, eradicate, recover, learn.
- **Evidence:** order of volatility, hash, chain of custody.
- **Reporting:** audience first; facts, unknowns, options, owners.

Glossary

ATT&CK

MITRE's knowledge base of adversary tactics and techniques.

Chain of custody

Documented record of who handled evidence and when.

CVSS

Common Vulnerability Scoring System for technical severity.

EDR

Endpoint detection and response: endpoint telemetry and host actions.

EPSS
Exploit Prediction Scoring System: estimated likelihood of exploitation.

IoC
Indicator of compromise, an artefact suggesting malicious activity.

MTTD / MTTR
Mean time to detect / mean time to respond or remediate.

SIEM
Platform that centralizes, correlates, and retains security logs.

SOAR
Security orchestration, automation, and response playbooks.

STIX / TAXII
Threat intelligence format / sharing protocol.

TTPs
Tactics, techniques, and procedures used by an adversary.

XDR
Extended detection and response across endpoint, identity, email, and cloud.

Before test day, compare this guide with CompTIA's official CS0-004 objectives, practise reading real log formats in a lab you control, and rework every question you missed. The best answer is supported by evidence, proportionate, and within your authority.