CERTLABZ CERTIFICATION BOARD

Certified Adversary Emulation Specialist (CAES)

Ten Sample Questions

These ten questions are written at the difficulty of the live exam. Six carry a figure, a data table, a configuration or a code listing; four are plain scenario questions, as on the live paper. Every one is a scenario a working adversary emulation specialist would meet, and none of them appears in the live question bank.

The live exam is multistage adaptive: how you score on one stage selects the difficulty of the next, so no two candidates sit the same paper. It is a performance-based exam: each form can carry up to 5 to 12 performance-based questions alongside the multiple-choice questions, which means you build, sequence or complete something rather than only choosing from a list. Answers and full explanations follow each question here; the live exam does not disclose item-level answers.

Question 1

Domain 1: Adversary Emulation Foundations and Threat Intelligence

DATA

Emma Wilson drafts an emulation plan and maps four observed behaviours to ATT&CK tactics.

Observed behaviour	Mapped tactic
A malicious attachment lands in a targeted inbox	Initial Access
A scheduled task is created to survive reboot	Persistence
NTLM hashes are dumped from memory	Impact
Staged files are sent to an external server	Exfiltration

Table 1.1 Behaviours and the tactic assigned to each.

Which mapping is wrong, and what is the correct tactic?

A. The attachment mapping is wrong; delivering a malicious attachment to an inbox is an Execution-tactic behaviour.

B. The scheduled-task mapping is wrong; a task that survives reboot is a Privilege Escalation behaviour, not Persistence.

C. The hash-dumping mapping is wrong; dumping credentials from memory is Credential Access, not Impact.

D. The exfiltration mapping is wrong; sending staged files to an external server is a Command and Control behaviour.

Correct answer: C
Dumping NTLM hashes from memory is OS Credential Dumping under the Credential Access tactic; Impact covers destructive or disruptive effects such as encryption or wiping. The attachment is Initial Access, the scheduled task is Persistence and sending staged files out is Exfiltration, so those three are correct.

Question 2

Domain 2: Planning, Scoping and Rules of Engagement

CONFIG

James Carter reviews an extract of the rules of engagement before an engagement begins.

```
scope_in    : 10.20.0.0/16, app.example.com
scope_out   : payroll-*, any partner-owned system
windows     : Mon-Fri 20:00-02:00 only
prohibited  : denial-of-service, destructive actions
stop_if     : production instability OR real intrusion found
contact     : soc-lead (deconfliction), ciso (escalation)
```

Listing 2.1 Extract from the rules of engagement.

During the window the team can reach a partner-owned server that trusts the client domain. What is correct?

A. Test it, because the trust relationship makes it effectively part of the client environment and therefore in scope.

B. Test it lightly and note it, because a brief check of a reachable system is acceptable under most rules of engagement.

C. Leave it alone, because partner-owned systems are explicitly out of scope and reachability is not authorisation.

D. Test it after informing the SOC lead, because the deconfliction contact can authorise expanding the scope on the day.

Correct answer: C
The rules of engagement place partner-owned systems out of scope, and only the system owner can authorise testing. Reachability, a trust relationship, or a verbal nod from a deconfliction contact do not grant authorisation, so the server must be left alone.

Question 3

Domain 4: Command and Control, and Infrastructure

DATA

Olivia Davis compares four candidate C2 configurations for an engagement.

Configuration	Behaviour
1	HTTPS beacon, fixed 60-second interval, no jitter, one domain
2	HTTPS beacon via a redirector, long interval with jitter, malleable profile
3	Plaintext HTTP beacon straight to the team server
4	Beacon straight to the client production cloud tenant

Table 3.1 Candidate C2 configurations.

Which configuration best emulates a resilient, stealthy actor while staying safe?

A. Configuration 1, because a fixed short interval keeps the operator in constant control of the compromised host.

B. Configuration 3, because plaintext traffic straight to the team server is the simplest and easiest path to debug.

C. Configuration 4, because routing through the client production tenant makes the traffic blend in perfectly with theirs.

D. Configuration 2, because a redirector, jittered long interval and malleable profile mirror how a real actor evades detection.

Correct answer: D
A redirector protecting the team server, a long jittered interval and a malleable profile that shapes traffic are how mature actors stay stealthy and resilient, and they create realistic detection opportunities. A fixed short interval and plaintext HTTP are trivially detected, and abusing the client production tenant as C2 is unsafe and inappropriate.

Question 4

Domain 5: Post-Exploitation, Persistence and Lateral Movement

CONFIG

Liam Anderson lists post-exploitation actions from a foothold. They are shown out of order.

```
A. Pass-the-hash to reach a file server (lateral movement)
B. Enumerate accounts, hosts and shares (discovery)
C. Exfiltrate benign marker files (exfiltration)
D. Dump credentials from the host (credential access)
E. Create a scheduled task for durability (persistence)
```

Listing 4.1 Actions to sequence.

Which order most faithfully emulates a typical actor after gaining a foothold?

A. A, then B, then D, then E, then C, moving to the file server before learning anything about the environment first.

B. B, then D, then A, then E, then C, discovering the environment, taking credentials, then moving, persisting and exfiltrating.

C. C, then B, then D, then A, then E, exfiltrating marker files first to prove the path before any other action is taken.

D. E, then A, then B, then D, then C, establishing persistence first so the foothold cannot be lost during the operation.

Correct answer: B
A realistic chain runs discovery first to understand the environment, then credential access to obtain material, then lateral movement to a higher-value system, then persistence, and finally collection and exfiltration. Moving or exfiltrating before discovery, or persisting before knowing where you are, does not match how actors actually operate.

Question 5

Domain 6: Detection, Purple Teaming and Reporting

DATA

Ava Thompson records the outcome of four emulated techniques during a purple-team session.

Technique	Outcome
Kerberoasting	Ticket request logged, no alert fired
Credential dumping	EDR alert fired correctly
DNS tunnelling	No log or telemetry found at all
Scheduled-task persistence	Event caused by the emulation team

Table 5.1 Techniques and their outcomes.

How should the Kerberoasting outcome be reported?

A. As a true positive, because the ticket request was logged and a log entry is equivalent to a working detection alert.

B. As a detection gap, noting the ticket-request telemetry exists and recommending a rule that turns it into an alert.

C. As a visibility gap, because the absence of an alert proves that no telemetry for the technique was collected at all.

D. As out of scope, because Kerberoasting is a credential-access technique and the session only measured persistence.

Correct answer: B
The ticket request was logged but no alert fired, so the telemetry exists and the fix is a detection rule: that is a detection gap, not a true positive. A visibility gap is where no telemetry exists at all, which describes the DNS tunnelling result, and nothing here is out of scope.

Question 6

Domain 3: Reconnaissance, Initial Access and Execution

CONFIG

The plan calls for spearphishing, but the client cannot risk real employees clicking during business hours. The team drafts options.

```
opt1: phish all staff live during business hours
opt2: run the phishing against a consenting test group
opt3: cancel the initial-access phase entirely
opt4: phish a competitor to prove the technique works
```

Listing 6.1 Options under discussion.

Which option lets the emulation stay faithful and safe?

A. Option 1, because phishing all staff live is the only way to measure the real click-through rate across the whole company.

B. Option 3, because cancelling initial access removes the risk while the rest of the engagement continues unaffected by it.

C. Option 4, because phishing a competitor exercises the same technique without putting the client's own employees at risk.

D. Option 2, because a consenting test group or seeded mailboxes exercises the technique and its detection within the rules.

Correct answer: D
A consenting test group or seeded mailboxes exercises the phishing technique and lets the team measure detection without disrupting the wider workforce, and the adjustment is documented. Phishing all staff ignores the constraint, cancelling loses the detection data, and phishing a competitor is unauthorised access to a third party.

Question 7

Domain 1: Adversary Emulation Foundations and Threat Intelligence

A client asks why the team insists on emulating a specific named threat actor rather than running a generic set of attacks. What is the best answer?

A. A generic attack set is always superior, because covering more techniques at random gives the broadest possible test.

B. Emulating a specific actor answers whether the organisation can detect the real threats that target its sector and assets.

C. A named actor is chosen only because it makes the final report sound more impressive to the executive audience.

D. The choice makes no practical difference, because detection either works against all techniques or it works against none.

Correct answer: B
Threat-informed emulation reproduces the TTPs of an actor that realistically targets the organisation, so the result answers a concrete question: could we detect this threat as it actually operates? A random set is not tied to real risk, the choice is not cosmetic, and detection coverage varies technique by technique.

Question 8

Domain 5: Post-Exploitation, Persistence and Lateral Movement

During an engagement the team reaches a database of real customer records and needs to demonstrate exfiltration. What is the correct approach?

A. Exfiltrate the entire database, because only a complete copy proves beyond doubt that the data could be stolen.

B. Exfiltrate a single real customer record, because one record is a small enough sample to avoid causing a data breach.

C. Exfiltrate benign canary or synthetic records that prove the path without moving any real customer data out.

D. Encrypt the database in place, because demonstrating impact is a more convincing proof than exfiltration ever is.

Correct answer: C
Proving the exfiltration path with benign canary or synthetic records demonstrates the risk without creating a real breach. Moving real customer data, even a single record, is itself a breach, and encrypting production is a destructive act that a faithful, safe emulation replaces with a benign marker.

Question 9

Domain 2: Planning, Scoping and Rules of Engagement

Midway through an engagement the team finds clear evidence of a separate, genuine intrusion already in progress. What should they do?

A. Continue the planned emulation and describe the genuine intrusion in the final report delivered at the end of the engagement.

B. Quietly take over the genuine intrusion to study the real attacker's behaviour and enrich the emulation findings.

C. Delete the evidence of the genuine intrusion so it does not get confused with the emulation team's own activity.

D. Stop the affected activity and follow the agreed escalation path to notify the client immediately, per the rules of engagement.

Correct answer: D

Discovery of a real, active intrusion is a defined stop-and-escalate condition. The team notifies the client immediately through the agreed contact so incident response can begin. Continuing regardless, interfering with the real attacker, or destroying evidence are all unacceptable.

Question 10

Domain 6: Detection, Purple Teaming and Reporting

The client detected 3 of 20 emulated techniques and wants to publish that the environment is "85% attacker-proof". How should the specialist respond?

A. Agree and publish the figure, because a single headline percentage is the clearest way to communicate security posture.

B. Explain that detecting some of one actor's techniques is not an overall security score, and reframe the result as coverage.

C. Refuse to share any results at all, because coverage figures can always be misread by a non-technical audience.

D. Raise the number to make it more reassuring, since the exact figure matters less than the confidence it gives leadership.

Correct answer: B

Emulation measures detection coverage against one threat model, not an overall security percentage, and 3 of 20 detected is in fact a low coverage result. The honest response is to reframe the finding as coverage against that actor with prioritised gaps. Publishing or inflating a headline number would mislead, and withholding results helps no one.

Published by the CertLabz Certification Board, a division of CertLabz LLC. Companion to objectives document CAES-001. https://certlabz.com/certifications/certified-adversary-emulation-specialist.html