CERTLABZ CERTIFICATION BOARD

Certified Threat Hunting Specialist (CTHS)

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 threat hunter 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 4: Analysis Techniques — QUERY

Emma Wilson writes a first hunt for encoded PowerShell across the fleet. Her query is below.

```
DeviceProcessEvents
| where FileName =~ "powershell.exe"
| where ProcessCommandLine has_any ("-enc", "-EncodedCommand", "-e ")
| summarize count() by DeviceName, InitiatingProcessFileName
| sort by count_ asc
```

Listing 1.1 The hunt query, stacking least-frequent first.

Which statement best assesses this query as a hunt?

A. It is too broad to be useful, because almost all PowerShell in an enterprise is launched with an encoded command.

B. It is a sound first hunt: encoded commands are rare and adversary-linked, and the parent process helps separate the two.

C. It is invalid, because an encoded PowerShell command cannot be matched once it has been base64 obfuscated by the caller.

D. It should be run only on the domain controllers, since encoded PowerShell is never launched on ordinary workstations.

Correct answer: B
Encoded commands are common for adversaries and comparatively rare for legitimate administration, so stacking them and reviewing the least frequent first is productive. The initiating process is the key discriminator: a management agent is benign, winword.exe or an unusual parent is not. Refine by excluding known deployment tooling rather than widening or narrowing to one host class.

Question 2

Domain 2: Adversary Tradecraft and ATT&CK

DATA

James Carter reviews a burst of Windows security events from one domain account within a few minutes.

Time	Event	Detail
09:14	4769	Service ticket, SPN svc-sql, encryption RC4 (0x17)
09:14	4769	Service ticket, SPN svc-backup, RC4 (0x17)
09:15	4769	Service ticket, SPN svc-web, RC4 (0x17)
09:15	4769	Service ticket, SPN svc-report, RC4 (0x17)
09:16	4769	Service ticket, SPN svc-erp, RC4 (0x17)

Table 2.1 Kerberos service-ticket requests from user jdoe.

Which technique does this pattern most likely indicate?

A. Pass the hash, because the account is reusing one stolen NTLM hash to reach several different backend services.

B. A golden ticket, because the account is presenting a forged ticket-granting ticket signed with the krbtgt hash.

C. Password spraying, because one account is testing a single common password against many service accounts at once.

D. Kerberoasting, because one account is pulling RC4 service tickets for many SPNs to crack their passwords offline.

Correct answer: D
Bulk 4769 requests for many service principal names, forced to weak RC4 encryption, from a single user in a short window is the Kerberoasting signature: the attacker harvests service tickets to crack the service-account passwords offline. Legitimate clients request tickets as they need them, not in a sweep. Pass the hash and golden tickets show different footprints, and spraying appears in logon failures, not ticket requests.

Question 3

Domain 2: Adversary Tradecraft and ATT&CK

DATA

Olivia Davis runs a beaconing analytic and reviews one candidate: a host contacting a single external domain at a steady interval for three weeks.

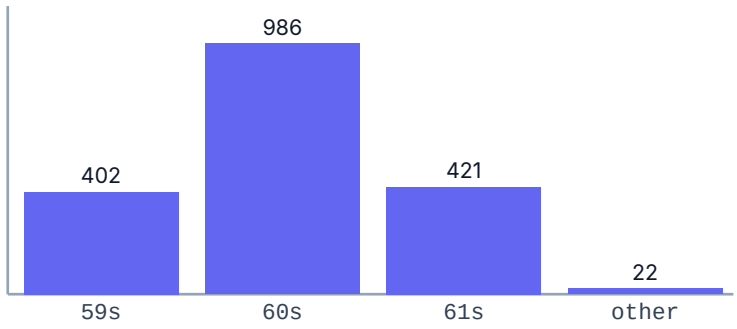

Figure 3.1 Intervals between the host's connections to the domain, over 21 days.

What is the correct next step before drawing any conclusion?

A. Dismiss it as normal keep-alive traffic, because a 60-second interval is exactly what a browser session uses.

B. Block the destination domain for the whole organisation, because the tight interval alone proves it is command and control.

C. Check the process, the full destination and whether other hosts share the pattern, since agents beacon too.

D. Reimage the host immediately, because three weeks of regular connections is conclusive evidence of an implant.

Correct answer: C
The near-constant interval with small jitter is the statistical fingerprint of beaconing, but legitimate agents beacon as well. The process making the connection, the full URL or SNI, and whether the pattern appears fleet-wide are what separate an EDR agent from an implant. Reputation and interval alone are weak discriminators, so verify before blocking or reimaging.

Question 4

Domain 5: Hunting in Endpoint, Network, Cloud and Identity

DATA

Liam Anderson finds the process lineage below on twelve servers within the same hour, each preceded by a network logon.

```
services.exe  (pid 684)
  % %  cmd.exe  /c powershell -nop -w hidden -enc <base64>
       % %  powershell.exe

preceding event: 4624 logon type 3 from 10.20.4.11
preceding event: 7045 new service "SysUpdate" created
```

Listing 4.1 Lineage and preceding events on one of the twelve servers.

What is the most likely explanation?

A. Routine patching, because services.exe launching a shell is how Windows Update applies fixes on server systems.

B. Remote service execution for lateral movement, given the service creation, the network logon and the hidden shell.

C. A scheduled backup job, because the new service named SysUpdate runs overnight maintenance across the servers.

D A time synchronisation fault, because twelve servers acting together within an hour points to a clock or NTP issue.

Correct answer: B

services.exe spawning a hidden, encoded PowerShell shell, paired with a new service (7045) and a type-3 network logon, is the classic PsExec-style remote service execution used for lateral movement. The same source IP across twelve servers in an hour is the spread. Windows Update does not create shells this way, and the pattern is nothing to do with backups or clocks.

Question 5

Domain 3: Telemetry and Data Sources

CONFIG

Ava Thompson plans a lateral-movement hunt and checks the estate's logging first. The coverage summary is below.

```
EDR sensors:        workstations 98%,  servers 0%
Sysmon:             not deployed
Security auditing:  4624 on; 4688 on (no command line); 1102 on
SIEM retention:     30 days
Log forwarding:     workstations only
```

Listing 5.1 Telemetry coverage for the estate.

Which improvement would most increase this hunt's capability?

A Extend endpoint telemetry to servers, because that is where lateral movement lands and there is currently no coverage.

B Add full packet capture on every network segment, because lateral movement is only visible in raw network traffic.

C Reduce SIEM retention to seven days, because a shorter window makes the lateral-movement queries run much faster.

D Enable command-line logging on workstations, because the servers already report their process command lines in full.

Correct answer: A

Servers are exactly where lateral movement lands, and they have no EDR, no Sysmon and no log forwarding, so the hunt is blind on the systems that matter most. Extending endpoint telemetry to servers is the biggest gain. Packet capture everywhere is disproportionate, shorter retention hurts long-dwell hunts, and command-line logging on workstations does not help because the servers are the gap.

Question 6

Domain 4: Analysis Techniques

DATA

A stack of scheduled task names across 5,000 hosts returns the tail below. The team must decide what to look at first.

Scheduled task name	Hosts	Binary path
GoogleUpdateTaskMachineCore	4,880	C:\Program Files\Google\...
OneDriveStandaloneUpdater	4,120	C:\Users\...\OneDrive\...
MicrosoftEdgeUpdateTaskMachine	4,905	C:\Program Files\Microsoft\...
GoogleUpdateTaskMachine	1	C:\ProgramData\gu.exe

Table 6.1 Least-frequent scheduled tasks after stacking.

Which task should be investigated first, and why?

A. The Edge update task, because it is present on the largest number of hosts and so has the widest possible blast radius.

B. The OneDrive updater, because it runs from a per-user profile path rather than from the Program Files directory tree.

C. The one-host Google task, because a real updater name is running an unsigned binary from an unexpected ProgramData path.

D. None of them yet, because a name appearing on only one host out of five thousand is statistically too rare to matter.

Correct answer: C
Rarity plus name mimicry with a wrong path is the lead: a task named like GoogleUpdate but running gu.exe from ProgramData, on a single host, is a classic masquerading persistence pattern. Stacking exists to surface exactly this long-tail item. The widespread signed updaters are expected, and rarity is the reason to look, not to dismiss.

Question 7 Domain 1: Hunting Foundations and Methodology

A team lead proposes the hunt hypothesis "look through the logs for anything suspicious this week". Why is this a poor hypothesis, and what would make it a good one?

A. It is too narrow, because one week is not long enough; widening it to a full year would make the hunt worthwhile.

B. It names too many data sources at once, so restricting it to a single source such as DNS would make it testable.

C. It is unfalsifiable and unbounded; a good hypothesis names a behaviour, a data source, a scope and a timeframe.

D. It is acceptable as written, because experienced hunters work best from open-ended prompts rather than fixed scopes.

Correct answer: C
A usable hypothesis can end in a clear result: found, not found, or data gap. "Anything suspicious" cannot, because it names no behaviour, source, scope or timeframe. A good version reads like "an adversary is using scheduled tasks for persistence on finance workstations, testable with event 4698 over the last 30 days".

Question 8

Domain 6: Detection Engineering, Reporting and Programme

A hunt reliably found malicious activity, and the team wants that value to persist. What should happen to the successful hunt query?

A. It should be kept private to the hunter who wrote it, so that an adversary cannot learn which behaviours are watched.

B. It should be converted into a documented, tuned detection with a response playbook so the SOC catches it continuously.

C. It should be deleted once the intrusion is handled, because a query that has found something has served its purpose.

D. It should be rerun by hand every week indefinitely, because scheduled automation cannot reproduce a hunter's judgement.

Correct answer: B

Hunt-to-detection conversion is the core output of a hunt programme: turn the query into a rule, tune its false positives, map it to ATT&CK, and attach a response playbook so the SOC catches the behaviour without a hunter present. Keeping it private or rerunning it manually wastes the value, and deleting it discards it entirely.

Question 9

Domain 1: Hunting Foundations and Methodology

A hunt for a specific persistence technique finds nothing after querying complete, well-covered telemetry for the right timeframe. What is the correct conclusion and action?

A. Declare the environment permanently free of that adversary, and close the hunt without any further documentation.

B. Record the negative with its scope and queries, and propose a detection so the technique is covered from now on.

C. Repeat the identical hunt every day indefinitely, since a technique found absent today may well appear tomorrow.

D. Discard the hunt as a failure, because a hunt that does not find an intrusion has produced no useful result at all.

Correct answer: B

A negative result over complete data is a real result, but only as good as the data it searched. Record the scope, timeframe and queries so it is reproducible, and turn the query into a detection so the technique is covered going forward. It does not prove the environment is clean forever, and a documented negative plus a new detection is a valuable outcome, not a failure.

Question 10

Domain 5: Hunting in Endpoint, Network, Cloud and Identity

In a cloud tenant, a hunter sees a user sign in from two countries forty minutes apart, then a new mailbox rule forwarding all mail externally, then a consent grant to an unknown OAuth application with mail.read. What is the assessment and the response?

A. Normal remote working, because staff travel and connect through different regional endpoints and cloud proxies routinely.

B A mail server defect, because forwarding rules and application consents are created automatically during service outages.

C Account takeover with persistence and durable access; revoke sessions and tokens, remove the rule and app, and hunt tenant-wide.

D Nothing actionable yet, because no alert has fired and the individual events are each permitted by current tenant policy.

Correct answer: C

The chain is textbook: impossible travel indicates takeover, the forwarding rule is persistence, and the illicit OAuth consent is durable access that survives a password reset. Identity hunts join sign-in, mailbox and application events into one story. The response is to revoke sessions and tokens, remove the rule and the app, then sweep the tenant for the same rule and consent pattern.

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