CERTLABZ CERTIFICATION BOARD

Certified Vulnerability Management Professional (CVMP)

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 vulnerability analyst 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, Scoring and Prioritisation

DATA

Emma Wilson can remediate one finding this window. The candidates and their signals are below.

Finding	CVSS	On KEV?	EPSS	Asset
A	9.8	No	0.02	Internal isolated test box, no data
B	7.5	Yes	0.71	Internet-facing customer database
C	8.1	No	0.09	Internal file server, some sensitive data
D	5.0	No	0.01	Internal print server

Table 1.1 Candidate findings with severity, exploitation and asset context.

Which finding should be remediated first, and on what reasoning?

A. Finding A, because a CVSS base score of 9.8 is the highest number present and therefore always the most urgent.

B. Finding C, because it sits on a file server and internal assets should always be prioritised ahead of external ones.

C. Finding B, because it is on the KEV list, has a high EPSS and sits on an exposed, sensitive asset despite its lower score.

D. Finding D, because print servers are widely deployed and a flaw there gives the broadest possible reach for an attacker.

Correct answer: C
Risk combines severity, exploitation and context. B is confirmed exploited (KEV), likely to be exploited (EPSS 0.71) and sits on an exposed, sensitive asset, so its real risk far exceeds a theoretical 9.8 on an isolated test box. The base score alone (A) and blanket "internal first" or "print servers" reasoning all miss the actual risk.

Question 2

Domain 3: Scanning and Assessment

DATA

James Carter compares two scans of the same Linux server taken minutes apart.

Scan type	Result for OpenSSL	Basis
Credentialed (authenticated)	Patched, not vulnerable	Read the installed package and backported fix
Unauthenticated network	Vulnerable, outdated version	Inferred from the service version banner

Table 2.1 Conflicting results for the same host.

How should the conflict be resolved?

A. Trust the network scan, because a version banner is a direct reading and credentialed scans are easily misconfigured.

B. Trust the credentialed scan: distributions backport fixes without changing the banner, so the banner result is a false positive.

C. Treat them as two distinct vulnerabilities, one at the package level and one separately at the network service level.

D. Discard both scans as unreliable and schedule a manual penetration test before drawing any conclusion at all.

Correct answer: B
Linux distributions routinely backport security fixes without bumping the version string, so banner-based detection false-positives while the credentialed scan reads the real, patched package. The credentialed result is more trustworthy; verify, then dismiss the banner finding rather than double-counting or discarding good data.

Question 3

Domain 1: Vulnerability Management Foundations and Governance

DATA

Olivia Davis documents her team's process. The steps, as currently practised, are listed below.

```
1. Run a quarterly authenticated scan of the documented ranges
2. Export the findings to a spreadsheet
3. Email the spreadsheet to IT and consider the job done
(no ownership, no SLAs, no verification, no inventory reconciliation)
```

Listing 3.1 The team's current process.

What is the most important gap to fix first?

A. Increase the scan frequency from quarterly to weekly, because scanning more often is the single biggest improvement here.

B. Buy a second scanner from another vendor, so that findings can be cross-checked between two independent products.

C. Assign accountable owners and severity-based SLAs and verify fixes, because emailing findings with no ownership fixes nothing.

D. Add more columns to the spreadsheet, so that the findings sent to IT carry more technical detail for each affected host.

Correct answer: C
The process finds vulnerabilities but nothing gets fixed: no owners, no deadlines, no verification. The biggest gap is governance, ownership, SLAs and verification, which turn findings into remediation. Scanning more often or adding tools and columns just produces a larger list nobody is accountable for.

Question 4

Domain 5: Remediation, Mitigation and Verification

DATA

Liam Anderson labels the actions taken on four findings this month.

Action taken	Finding
Applied the vendor patch that removes the flaw	A web server RCE
Added a WAF rule blocking the exploit, patch pending	A second web server RCE
Segmented and access-restricted an unpatchable appliance	A legacy device
Rescanned to confirm the flaw is gone	The first web server

Table 4.1 Actions and the findings they were applied to.

Which classification of these actions is correct?

A. All four are remediation, because each one reduces the risk associated with the finding it was applied to in some way.

B. The patch and rescan are remediation; the WAF rule and the segmentation are both mitigations, not fixes.

C. The patch is remediation; the WAF rule is remediation; the segmentation and the rescan are both mitigations.

D. The patch and the segmentation are remediation; the WAF rule and the rescan are both classed as verification steps.

Correct answer: B
Remediation removes the flaw (the patch), and the rescan is the verification that confirms it. The WAF rule and the segmentation reduce risk without removing the underlying vulnerability, so both are mitigations. Treating a mitigation as a fix, or grouping the rescan with the controls, misstates what has actually been resolved.

Question 5

Domain 6: Reporting, Metrics and Programme

DATA

A dashboard shows three months of programme metrics.

Metric	Month 1	Month 2	Month 3
Total open findings	22,000	19,000	16,000
Critical MTTR (days)	11	18	27
Scan coverage	95%	95%	95%
Critical findings on KEV, open	40	55	70

Table 5.1 Programme metrics over three months.

What is the most accurate reading of this dashboard?

A. The programme is improving, because the total number of open findings has fallen steadily every month across the quarter.

B. The programme is stable, because scan coverage has held at 95 percent and coverage is the metric that matters most here.

C. The programme is degrading on risk: critical MTTR and open exploited criticals are rising even as easy low-risk findings clear.

D. The dashboard is inconclusive, because three months is far too short a period to read any trend from these four metrics.

Correct answer: C
Falling totals look good but hide the real story: the time to fix critical findings has more than doubled and the number of open, actively exploited criticals is climbing. The team is clearing easy, low-risk findings while the dangerous ones pile up. Risk is rising, which the raw total conceals.

Question 6

Domain 2: Asset Discovery and Inventory

DATA

Ava Thompson reconciles asset data from four sources for one environment.

Source	Hosts seen	Note
Documented scan scope	1,200	What the scanner is configured to cover
Passive network monitoring	1,500	Sees traffic from live hosts
Cloud provider API	340	Autoscaling instances, minutes-lived
CMDB	1,150	Last bulk-updated 8 months ago

Table 6.1 Asset counts by discovery source.

What should the reconciliation conclude?

A. The scanner scope is authoritative, so the extra hosts seen by passive monitoring can safely be disregarded as noise.

B. There are unmanaged and ephemeral assets outside scope; expand coverage, use cloud APIs and reconcile the stale CMDB.

C The CMDB is the single source of truth, so the count of 1,150 should be used and the other three sources discarded.

D The counts broadly agree, so the inventory is accurate and no further discovery or reconciliation work is required now.

Correct answer: B
Passive monitoring sees 300 hosts the scan scope misses, the cloud API reveals short-lived instances a periodic scan would never catch, and the CMDB is eight months stale. The conclusion is unmanaged and ephemeral assets outside coverage: expand scope, use continuous cloud discovery and reconcile the CMDB, rather than trusting any single source.

Question 7 Domain 4: Analysis, Scoring and Prioritisation

A team lead insists on remediating strictly in descending CVSS base-score order, arguing it is the fairest and most objective method. Why is this approach flawed?

A It is not flawed at all; the CVSS base score is a complete measure of risk and sorting by it is the correct method.

B It is flawed only because CVSS scores are frequently miscalculated, so the ordering they produce is usually inaccurate.

C It ignores exploitation and context, so theoretical high scores jump ahead of lower-scored findings under active attack.

D It is flawed because newer vulnerabilities should always be remediated before older ones regardless of their severity.

Correct answer: C
The CVSS base score is intrinsic severity, not risk. Sorting by it alone puts a theoretical 9.8 on an isolated box ahead of a 7.5 that is on the KEV list, exposed and on a sensitive asset. Real prioritisation folds in exploitation (KEV, EPSS), exposure and business criticality, not just the number.

Question 8 Domain 5: Remediation, Mitigation and Verification

A finding is closed in the tracker marked "patched", but no rescan was run to confirm it. Why is verification by rescan an essential step before closing a finding?

A It is not essential; once the responsible team marks a patch as applied, the finding can be closed with full confidence.

B A patch may fail to apply, need a reboot or miss a second instance, so a rescan is what confirms the flaw is truly gone.

C Rescanning is only required for compliance audits, so it can be skipped whenever the programme is not being audited.

D Verification matters only for critical findings, so lower-severity findings can always be closed without any rescan.

Correct answer: B
Closed does not mean fixed. Patches fail to apply, wait on reboots, or leave a second vulnerable instance untouched, and a claimed fix is only confirmed by a verification rescan. Skipping it, or reserving it for audits or criticals only, lets unremediated risk be recorded as resolved.

Question 9

Domain 5: Remediation, Mitigation and Verification

A critical, actively exploited vulnerability affects a production system whose next maintenance window is two weeks away. What is the correct response?

A. Wait for the scheduled maintenance window in two weeks, since following the normal change process is always the priority.

B. Accept the risk permanently and close the finding, because the maintenance window makes timely patching impractical here.

C. Escalate for an emergency change given active exploitation, and apply interim mitigations until the patch is deployed.

D. Remove the system from the scan scope so the finding no longer appears, then revisit it at the next quarterly review.

Correct answer: C
Active exploitation of a critical flaw justifies breaking the normal cadence: escalate for an emergency change, and in the meantime reduce exposure with isolation, access restriction or virtual patching. Waiting two weeks, accepting the risk outright, or hiding the finding by removing it from scope all leave an exploited weakness in production.

Question 10

Domain 1: Vulnerability Management Foundations and Governance

The board asks for a single number to represent the organisation's vulnerability risk. What is the soundest answer to give, and what caveat must accompany it?

A. The raw count of open vulnerabilities, because it is simple, objective and easy for the board to understand at a glance.

B. The number of scans completed this quarter, because activity is the most reliable proxy for how much risk is being managed.

C. A risk-weighted exposure score trended over time, reported alongside coverage so it cannot be improved by scanning less.

D. The single highest CVSS score currently open, because the worst individual finding best represents the overall risk level.

Correct answer: C
A risk-weighted, trended exposure score is the best single indicator, because it reflects severity, exploitability and asset importance together. The essential caveat is coverage: without it, a team could improve the number simply by scanning fewer assets. Raw counts, activity metrics and a single worst score all mislead.

Published by the CertLabz Certification Board, a division of CertLabz LLC. Companion to objectives document CVMP-001. https://certlabz.com/certifications/certified-vulnerability-management-professional.html