Exam - Database Security and Forensics

(June 3rd, 2026 - Duration: 90 minutes)

Last name:	First name:	Group:

Instructions: exactly one correct answer per question. **Circle the letter corresponding to the correct answer.** Choose the most precise and most complete option. Read all options before answering. No documents are allowed. No negative marking.

Question 1. A bank deploys very strict multi-factor authentication (MFA) that adds 8 seconds to every login. Following complaints, some operators share a single already-open account to save time. Which CIA chain of events most accurately describes what happens?

A	A gain in Availability degraded Confidentiality: MFA speeds up access but exposes shared credentials
B	Only Integrity is affected: shared accounts do not modify data, so Confidentiality and Availability remain intact
C	A gain in Confidentiality degraded perceived Availability, leading to a human workaround that recompromises Confidentiality and traceability
D	No pillar is affected: MFA stays active, session sharing remains internal, and overall security is preserved

Question 2. A system enforces CHECK constraints, ACID transactions, and audit triggers on every write. Users notice that bulk inserts have become very slow. From a CIA standpoint, how should this situation be correctly characterized?

A	An Integrity violation: the write slowdown ends up silently corrupting the inserted data
B	A Confidentiality violation: the audit triggers expose sensitive data in cleartext to unauthorized third parties
C	A design error: integrity and performance are two completely independent dimensions
D	A deliberate Integrity ↔ Performance trade-off: integrity checks add overhead that reduces effective availability

Question 3. An administrator accidentally deletes 50,000 rows, then immediately runs *ROLLBACK* before any COMMIT. The data reappears intact. Which ACID property made this recovery possible, and which pillar was preserved?

A	Isolation → Confidentiality: concurrent transactions are partitioned, which protects data secrecy
B	Durability → Availability: once written, data persists and therefore remains accessible to users
C	Atomicity → Integrity: the uncommitted transaction is rolled back as a whole, restoring an exact and consistent state
D	Consistency → Confidentiality: constraints ensure that only valid data remains secret

Question 4. We have 4 disks each with an individual annual reliability of $R_d = 0.90$, arranged as two parallel groups, each group consisting of two disks in series. Which value is closest to the overall reliability of the system?

A	0.6561: the four disks are in series, so the reliability equals 0.90^4
B	0.9639: each series group equals $0.90^2 = 0.81$ and the two parallel groups give $1-(1-0.81)^2$
C	0.9999: such an arrangement tolerates the failure of any disk, so the reliability tends toward 1
D	0.9000: parallelism brings nothing here, the system equals the reliability of a single disk

Question 5. Three disks with reliability $R_d = 0.90$ are arranged in RAID 5 (tolerates the failure of a single disk). Which expression and approximate value give the reliability of the volume?

A	$0.90^3 + 3\times0.90^2\times0.10 = 0.972$: the volume survives zero or exactly one disk failure
B	$0.90^3 = 0.729$: to function, RAID 5 requires all three disks to be operational
C	$1-(1-0.90)^3 = 0.999$: a single one of the three disks remaining functional is enough for the data
D	$0.90^2 = 0.81$: the parity disk stores nothing, so only two disks actually count

Question 6. A provider raises its SLA from 99.9% to 99.99% availability. Two statements are proposed: (1) the maximum annual downtime drops from about 8.8 hours to about 53 minutes; (2) the cost typically grows linearly with this gain. Which one(s) is/are correct?

A	Both statements are correct: the downtime calculation is right and the cost grows proportionally to the gain
B	Only statement (2) is correct: the stated durations are wrong but the cost does grow linearly
C	Only statement (1) is correct: the cost of availability is non-linear, one extra “nine” potentially costing ~10× more
D	Neither is correct: the durations are wrong and the cost stays constant regardless of the target level

Question 7. A company sets two continuity objectives: lose no more than 15 minutes of data in case of an incident (RPO, Recovery Point Objective = 15 min) and restore service in under 5 minutes (RTO, Recovery Time Objective = 5 min). Which type of backup site meets both requirements?

A	A COLD site, because its 48–72 h activation leaves ample time to restore the backups
B	A WARM site, because its few hours of restoration time are enough to meet a 5-minute RTO
C	Any site will do: RPO and RTO are equivalent and interchangeable objectives
D	A HOT site, because its near-real-time replication and failover within minutes satisfy a 5-minute RTO

Question 8. Statistics show that about 55% of breaches are internal while about 80% of budgets target external threats. A manager concludes: “We should therefore shift 80% of the budget toward fighting internal threats.” Which criticism is the most rigorous?

A	Frequency does not translate linearly into budget: the trade-off must weigh impact, cost, and risk, not frequency alone
B	The conclusion is correct: budget allocation must exactly mirror the observed proportion of internal and external incidents
C	The conclusion is wrong because internal threats are in fact negligible compared to targeted external attacks
D	On the contrary, the external budget should be increased, because external attackers deploy means that are far more expensive to counter

Question 9. Match each real-world case to the most precise insider-threat subcategory: (1) Capital One 2019, (2) Tesla 2019 (theft of the Autopilot source code, departure to competitor XPeng).

A	(1) negligent insider who misconfigured the firewall, later exploited by an external hacker; (2) malicious insider who stole source code
B	(1) malicious insider who stole customer data; (2) negligent insider who made a cloud misconfiguration error
C	(1) and (2) are purely external attacks, with no internal factor or prior legitimate access
D	(1) former employee retaining their access; (2) administrator whose account was hacked by a third party

Question 10. Why is a malicious insider threat structurally harder to counter than an equivalent external attacker, and which combination of countermeasures should be preferred?

A	Because it uses more advanced attack tools; the remedy is to deploy a more powerful perimeter firewall
B	Because the insider already has access, knowledge, and trust; the remedy combines least privilege, separation of duties, and auditing
C	Because it is always state-sponsored; the remedy relies on TLS encryption of network communications
D	Because it necessarily constitutes a persistent APT; the remedy is to disconnect the company network

Question 11. Consider the server query: SELECT * FROM users WHERE user='$u' AND pass='$p'. An attacker enters in the username field ' OR '1'='1' -- (the -- is followed by a space) and leaves the password empty. What is the exact effect?

A	The query triggers an SQL syntax error and is immediately rejected by the database engine
B	user='' is false but OR '1'='1' is always true and -- comments out the password check: the query returns all rows and logs in the first account
C	Only users whose name literally equals the value 1 are selected and returned by the query
D	The double dash encrypts the submitted query, which systematically triggers a WAF (web application firewall) alert

Question 12. An application authenticates via SELECT * FROM users WHERE user='$u' AND pass='$p'. To block injections, a developer forbids the keyword OR (regardless of case) in inputs. Why does this “blacklist” remain insufficient?

A	The attacker bypasses the filter without the word OR: by entering admin' -- as the name, they comment out the password check and log in as admin
B	It is actually sufficient: forbidding the word OR neutralizes all possible SQL injections
C	The mistake was targeting the word OR: in reality only whitespace characters should have been filtered
D	It is enough, in addition, to strengthen the password robustness of the database administrator account

Question 13. A page is vulnerable to SQL injection, but it never displays any data or error message. The attacker injects: ... AND IF(SUBSTRING(password,1,1)='a', SLEEP(5), 0). The page takes 5 seconds to respond. What does the attacker conclude?

A	That the server is overloaded: the delay is due to traffic and reveals nothing about the stored password
B	That the injection failed: a response delay means the database rejected the malicious query
C	That the 1st letter of the password is “a”; by repeating the test letter by letter, the attacker reconstructs the entire password
D	That the password is exactly 5 characters long, one per second of delay measured by the attacker

Question 14. WannaCry (2017) is described as hybrid. Which combination and causal chain best explain its worldwide virulence within 4 days?

A	Virus + Trojan horse: the virus replicates from file to file and the Trojan then opens a persistent backdoor
B	Worm + ransomware: the worm spread by itself via EternalBlue (SMBv1) and the ransomware encrypted files, hence the speed and impact
C	Botnet + ransomware: a central C&C server registered in advance remotely coordinated the encryption of all machines
D	Rootkit + worm: the rootkit hid the code at the lowest level, making the threat completely undetectable by any antivirus

Question 15. Which operational difference, rather than purely a matter of scale, separates the ILOVEYOU worm (2000) from WannaCry (2017)?

A	ILOVEYOU encrypted the victim's files whereas WannaCry merely spread without causing damage
B	WannaCry targeted only Linux servers whereas ILOVEYOU exclusively targeted Windows workstations
C	The only difference is scale: WannaCry simply affected a far larger number of machines
D	ILOVEYOU required opening an attachment; WannaCry exploited a network vulnerability (SMBv1) and spread without interaction

Question 16. Why does a Trojan horse often evade a purely signature-based antivirus, whereas a known classic worm is more easily blocked?

A	Because it presents itself as legitimate software: the wrapper looks clean and hides the payload, whereas a known worm has a signature
B	Because the Trojan spreads much faster than the distribution of antivirus signature-database updates
C	Because the Trojan resides only in RAM and never writes a single file to the hard disk
D	Because the Trojan systematically encrypts the entire disk from the very first second of execution

Question 17. A startup collects emails to send a newsletter, then resells them a month later to an advertiser. Beyond consent, which GDPR principle is violated, and within what time frame should a breach of this data be notified?

A	The right to data portability is violated, and any breach must be notified to the authority within 30 days
B	Purpose limitation is violated: a new use requires new consent, and a breach is notified within 72 h
C	The right to erasure is violated since the emails are retained, and any breach must be reported within 24 h
D	No principle is violated: once an email is obtained lawfully, reselling it to a third party is entirely unrestricted

Question 18. In a defense-in-depth setup, an attacker nonetheless succeeds with an SQL injection that reaches the DBMS. Which “database”-layer mechanism still limits the damage, and why does encryption at rest remain useful even after access?

A	The network firewall blocks the query after the fact, and data encryption simply prevents connection to the DBMS
B	The antivirus installed on the server removes the detected injection, and encryption incidentally speeds up query processing
C	No further protection is possible: once the network layer is breached, the attacker freely accesses everything
D	RBAC (Role-Based Access Control) restricts the account to read-only and limits the impact; encryption renders stolen data unreadable and auditing logs everything

Question 19. Correctly order the continuous security cycle and identify the most frequent mistake.

A	Detection → Prevention → Reaction → Recovery; the most frequent mistake is devoting too many resources to auditing
B	Assessment → Design → Deployment → Management then back to assessment; the mistake is stopping after the first phases
C	Design → Assessment → Management → Deployment; the common mistake is wanting to start the process with an audit phase
D	No cycle is necessary: security is a product you install once and that requires no further maintenance afterward

Question 20. A defense-in-depth setup forces an SQL injection to pass through 3 independent barriers: the WAF (Web Application Firewall) stops 70% of attempts, application-level validation stops 80% of what gets through, and parameterized queries stop 90% of the rest. What is the approximate probability that an attempt passes all three layers, and which principle does this illustrate?

A	≈ 24%: this high rate shows that in practice a single layer of protection is almost always enough to stop the attack
B	≈ 60%: the pass-through probabilities of the different layers add up to give the total
C	≈ 0.6%: the product of the pass-through rates (0.30 × 0.20 × 0.10) shows that independent layers reduce risk multiplicatively
D	0%: stacking three independent layers guarantees, in theory, absolute security against injection

Good Luck - Mr. L. FOUGHALI