Annotated Answer Key — Database Security and Forensics

(June 3rd, 2026 — Instructor document: justifications, without the questions)

Quick answer grid

Q1 : C	**Q11 : B**
Q2 : D	**Q12 : A**
Q3 : C	**Q13 : C**
Q4 : B	**Q14 : B**
Q5 : A	**Q15 : D**
Q6 : C	**Q16 : A**
Q7 : D	**Q17 : B**
Q8 : A	**Q18 : D**
Q9 : A	**Q19 : B**
Q10 : B	**Q20 : C**

Detailed justifications

Question 1 — Correct answer: C

Hardening security (MFA) targets Confidentiality. The perceived slowness degrades perceived Availability, which pushes operators toward a human workaround (session sharing) that recompromises both Confidentiality and traceability. This is a direct illustration of the Confidentiality ↔ Availability tension in the CIA Triangle (slide 25). A reverses the roles (MFA does not gain Availability); B denies the harm to Confidentiality; D denies any impact even though account sharing breaks accountability.

Question 2 — Correct answer: D

CHECK constraints, ACID transactions, and audit triggers strengthen Integrity but add overhead that slows down writes and therefore reduces effective Availability. This is the deliberate Integrity ↔ Performance trade-off (slide 25). A is wrong: a slowdown does not corrupt data; B is off-topic (auditing discloses nothing to third parties); C is wrong because integrity and performance are precisely not independent.

Question 3 — Correct answer: C

A ROLLBACK before COMMIT illustrates the Atomicity (“all or nothing”) of the ACID properties: the uncommitted transaction is rolled back as a whole, which restores an exact and consistent state — hence Integrity (slides 12-14). B confuses Durability (which concerns what is already committed); A and D wrongly tie the recovery to Confidentiality.

Question 4 — Correct answer: B

Series-parallel arrangement. Each group = 2 disks in series: $0.90 \times 0.90 = 0.81$. Two groups in parallel: $1 - (1 - 0.81)^2 = 1 - (0.19)^2 = 1 - 0.0361 = 0.9639$. A treats the 4 disks as pure series ($0.90^4 = 0.6561$); C overestimates; D ignores the contribution of parallelism. Application of the series (product) and parallel $1-(1-R)^n$ formulas (slides 21-22).

Question 5 — Correct answer: A

RAID 5 on 3 disks tolerates exactly one failure. Reliability = P(0 failures) + P(1 failure) = $0.90^3 + 3 \times 0.90^2 \times 0.10 = 0.729 + 0.243 = 0.972$. B wrongly assumes all 3 disks are required; C applies a pure-parallel formula (RAID 1 with 3 disks); D ignores the actual role of parity (slide 19, RAID 5 “tolerates 1 failed disk”).

Question 6 — Correct answer: C

Only statement (1) is correct. 99.9% ≈ 8.76 h/year and 99.99% ≈ 52.6 min/year: the stated durations are accurate (slide 15). However, the cost of availability is NON-linear: moving from 99.9% to 99.99% can cost ~10× more (slide 25), so statement (2) is false. A and B wrongly accept linearity; D wrongly rejects the durations.

Question 7 — Correct answer: D

RTO = 5 min and RPO = 15 min are strict requirements. Only a HOT site, with near-real-time replication (low RPO) and failover within minutes (low RTO), satisfies them (slide 24). A (COLD, 48-72 h) and B (WARM, a few hours) violate the RTO; C is wrong because RTO and RPO measure two different things (recovery time vs data loss).

Question 8 — Correct answer: A

The rigorous criticism: the frequency of a type of incident does not convert linearly into budget. A correct trade-off weighs impact, cost of the countermeasure, and residual risk, not merely the observed frequency (slides 12-13, the 55%/80% “paradox”). B takes the proportion as an allocation rule; C wrongly downplays the internal threat; D makes the opposite, equally simplistic argument.

Question 9 — Correct answer: A

Subcategorization consistent with the course cases. Capital One 2019 (slide 31): NEGLIGENT insider (cloud firewall misconfiguration), the flaw then exploited by an external attacker. Tesla 2019 (slide 32): MALICIOUS insider (downloading and exfiltrating the Autopilot source code, ahead of a move to the Chinese competitor XPeng). B swaps the two profiles; C denies any internal factor; D invents scenarios (former employee / hacked account) inconsistent with the studied cases.

Question 10 — Correct answer: B

The insider is structurally more dangerous because they combine three advantages: legitimate access already granted, internal knowledge of where the data is located, and the system's implicit trust (slide 30). The remedy combines least privilege, separation of duties, and auditing (slides 11, 13, 33). A reduces the problem to tools and the perimeter; C and D rely on false generalizations (necessarily state-sponsored / necessarily an APT).

Question 11 — Correct answer: B

The input ' OR '1'='1' -- makes the WHERE clause always true (OR '1'='1') and the -- comments out the password check. The query then returns all rows and authenticates the first account (slide 35, the ' OR 1=1 -- mechanism). A is wrong: the query is syntactically valid; C misreads the literal; D is fabricated (the -- encrypts nothing).

Question 12 — Correct answer: A

A blacklist of just the OR keyword can be bypassed without OR: entering admin' -- as the name comments out the password check and logs in as admin. This is why blacklist filtering is insufficient; the real remedy is the parameterized query that separates code and data (slide 13). B is wrong (the filter does not cover everything); C and D propose ineffective fixes.

Question 13 — Correct answer: C

Time-based blind SQL injection (time-based Blind SQLi). The 5 s delay triggered by SLEEP(5) confirms that the condition SUBSTRING(password,1,1)='a' is true: the 1st letter is “a”. By repeating letter by letter, the attacker reconstructs the entire password. A attributes the delay to chance; B believes it failed; D misinterprets the duration.

Question 14 — Correct answer: B

WannaCry (2017) is the Worm + Ransomware hybrid: the worm spread by itself via the EternalBlue (SMBv1) flaw, without interaction, and the ransomware encrypted files — hence the global speed (≈230,000 machines, 150 countries in 4 days) and impact (slide 29). A, C, and D combine the wrong families (virus/Trojan, botnet/C&C, rootkit).

Question 15 — Correct answer: D

An operational difference, not one of scale: ILOVEYOU (2000) required a human action (opening the attachment), whereas WannaCry exploited a network vulnerability (SMBv1) and spread without interaction (slides 29, 41). A reverses the behaviors; B invents OS targets; C wrongly reduces the difference to scale alone.

Question 16 — Correct answer: A

A Trojan horse presents itself as legitimate software: the wrapper looks clean and hides the payload, which defeats a signature-based antivirus, whereas a known worm has a detectable signature (slides 39, 42). B, C, and D describe behaviors that are not the Trojan's distinctive mechanism (speed, memory residence, encryption).

Question 17 — Correct answer: B

Beyond consent, it is purpose limitation that is violated: emails collected for a newsletter cannot be resold without new consent for that new use. And a data breach must be notified within 72 h (slides 4, GDPR). A invents a 30-day deadline and the wrong right; C gets both the deadline (24 h) and the principle wrong; D is false (reselling is not unrestricted).

Question 18 — Correct answer: D

Even after a successful injection, the “database” layer limits the damage: RBAC restricts the application account to read-only (least privilege), auditing logs all actions, and encryption at rest renders stolen data unreadable without the key (slides 48-49, defense in depth). A, B, and C describe mechanisms that are ineffective or nonexistent at this stage.

Question 19 — Correct answer: B

Continuous security cycle: Assessment → Design → Deployment → Management, then back to Assessment. The most frequent mistake is stopping after the first phases, which leads to obsolescence (slide 46). A, C, and D give a wrong order or deny the need for a continuous cycle.

Question 20 — Correct answer: C

Independent layers: the probability of passing all three is the PRODUCT of the pass-through rates. The WAF lets through 30%, validation 20%, parameterized queries 10%: $0.30 \times 0.20 \times 0.10 = 0.006 \approx 0.6\%$. This illustrates the MULTIPLICATIVE reduction of risk through defense in depth (slides 49, 20). A, B, and D use the wrong model (a single layer would suffice / addition / absolute security).