CERTLABZ CERTIFICATION BOARD

Certified Post-Quantum Cryptography Specialist (CPQS)

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 security engineer 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 5: Protocols and Implementation

CONFIG

Emma Wilson enables hybrid post-quantum key exchange on a fleet of TLS 1.3 load balancers. The configuration and a week of negotiated-group telemetry are below.

```
# tls.conf (excerpt)
min_protocol      = TLSv1.3
groups            = X25519MLKEM768:X25519:secp256r1
fallback_on_error = true
```

Listing 1.1 Load balancer TLS configuration after the change.

Negotiated group	Sessions	Handshake failures	Median handshake
X25519MLKEM768	612,400	0.02%	31 ms
X25519	401,900	0.03%	28 ms
secp256r1	18,300	0.04%	33 ms
(none, failed)	9,870	100%	-

Table 1.1 Seven days of telemetry. Failed handshakes have no negotiated group.

The failed handshakes cluster on clients behind one vendor's outbound proxy and succeed when the client retries without the hybrid group. What is the most likely cause and the correct fix?

A. The server is rejecting the larger ClientHello; raise its record size limit and the failures on every client will stop.

B. The proxy cannot pass a ClientHello that spans more than one segment; fix the proxy and keep hybrid first with retry as a stopgap.

C. ML-KEM decapsulation fails on the server for those clients; move X25519 first in the group list so they negotiate it.

D. The proxy is intercepting traffic as a man in the middle and the failures are correct; no configuration change is needed.

Correct answer: B
A ClientHello offering X25519MLKEM768 carries a 1,184-byte encapsulation key, so it no longer fits in one segment. Middleboxes that assumed a small first flight broke in exactly this way during the 2024 browser rollouts. The server cannot be the cause: 612,400 hybrid sessions succeeded through it. Reordering groups would give up post-quantum protection for everyone to spare one proxy.

Question 2

Domain 2: Post-Quantum Algorithms and Standards

DATA

James Carter must choose a signature scheme for three artefacts using the measured characteristics below.

Scheme	Public key	Signature	Sign / s	Verify / s	Assumption
ML-DSA-65	1,952 B	3,309 B	9,800	31,000	Module lattice
SLH-DSA-SHA2-128s	32 B	7,856 B	2	1,900	Hash only
FN-DSA-512	897 B	666 B	5,600	38,000	NTRU lattice, float sampler
LMS (H20, W8)	60 B	1,620 B	12	4,200	Hash only, stateful

Table 2.1 Measured on one core. Signatures per second are rounded.

The artefacts: (1) TLS server certificates for 40,000 new connections per second; (2) a bootloader verification key burned into read-only memory for 15 years; (3) DNSSEC records that must fit a 1,232-byte UDP response with two signatures. Which assignment is correct?

A. (1) SLH-DSA-SHA2-128s, (2) ML-DSA-65, (3) LMS: the most conservative scheme for TLS, the fastest for the bootloader.

B. (1) FN-DSA-512, (2) LMS, (3) SLH-DSA-SHA2-128s: the smallest for TLS, the stateful scheme for the bootloader.

C. (1) LMS, (2) FN-DSA-512, (3) ML-DSA-65: the stateful scheme for TLS, the compact scheme for the bootloader.

D. (1) ML-DSA-65, (2) SLH-DSA-SHA2-128s, (3) FN-DSA-512: fast verification for TLS, hash-only for the bootloader.

Correct answer: D
TLS needs fast verification and tolerable size at high volume, which is ML-DSA. A verifier frozen in silicon for 15 years wants the most conservative assumption, hash only, and signs so rarely that two signatures per second is irrelevant. Only FN-DSA's 666-byte signatures leave room for two in a DNS response. LMS is stateful and belongs in firmware pipelines with hardware state management, not in DNS.

Question 3

Domain 1: Cryptographic Foundations and the Quantum Threat

TIMELINE

Olivia Davis presents Mosca's inequality to a hospital board. Imaging records must stay confidential for 25 years, the migration is scoped at 6 years, and the risk committee's horizon for a cryptographically relevant quantum computer is 10 to 15 years.

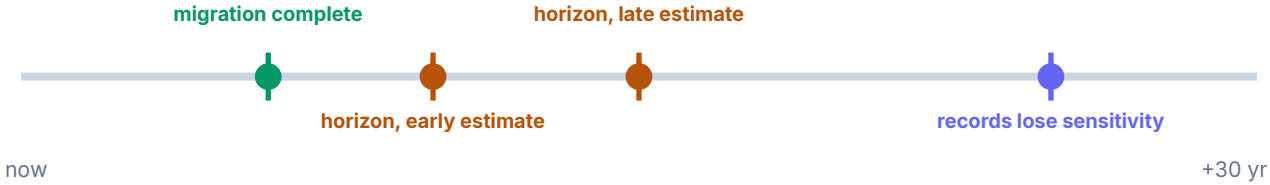

Figure 3.1 The timeline, in years from today.

Which statement should she make?

A. The programme is safe because the migration completes at year 6, before even the earliest horizon estimate at year 10.

B. Only the late estimate matters for planning, and since 6 years is less than 15 there is no exposure to report to the board.

C. Records sent today under classical key exchange are already exposed, so the paths carrying imaging must move to hybrid first.

D. The records become exposed only if the migration slips past year 10, so the board should fund a schedule contingency.

Correct answer: C
Mosca's inequality compares shelf life plus migration time (31 years) with the horizon (10 to 15 years). Finishing the migration before the horizon does not help data recorded before it finished, because that data stays sensitive for 25 years. The only effective response is to protect the recordable channels for this data class immediately, ahead of the general estate.

Question 4

Domain 3: Cryptographic Inventory and Risk Assessment

DATA

A discovery run produced the inventory extract below. Liam Anderson must rank the four assets for migration.

Asset	Cryptography	Data shelf life	Recordable by adversary?	Change effort
A. Site-to-site VPN	IKEv2, DH group 14, AES-256	20 years	Yes, leased line	Low: firmware supports RFC 9370
B. Internal PKI root	RSA-4096, 10-year cert	n/a (signature)	n/a	High: 30% of relying parties cannot verify ML-DSA
C. Public website	TLS 1.3, X25519, ECDSA	None	Yes	Low
D. Device firmware key	ECDSA P-256 in OTP memory	n/a (signature)	n/a	Impossible after shipment; ships in 4 months

Table 4.1 Inventory extract. "Recordable" means an adversary could plausibly capture the traffic.

Which ranking is correct, from first to last?

A. D, A, B, C: the unchangeable firmware key first, then long-lived VPN data, then the verifier-gated PKI, then the public site.

B. B, A, D, C: the oldest and largest key first, then long-lived VPN data, then the firmware key, then the public site.

C. A, C, B, D: the internet-facing assets first in order of exposure, then the internal PKI, then the device firmware key.

D. C, A, B, D: the public website first because it is the most exposed, then the VPN, then the PKI, then the firmware key.

Correct answer: A
Two assets have an unmovable deadline. The firmware key is a trust anchor that cannot be updated once burned, so it must carry a post-quantum scheme before the ship date. The VPN is the harvest-now-decrypt-later case: long shelf life over a recordable link, cheap to fix. The PKI is gated by the relying parties, so it is scheduled. The public website has no confidentiality to lose.

Question 5

Domain 4: Migration Planning and Crypto Agility

CODE

Ava Thompson reviews a pull request that adds "post-quantum support" to a key-wrapping library used by 60 services.

```
def wrap(dek: bytes, pub) -> bytes:
    if isinstance(pub, MLKEMPublicKey):
        ct, ss = mlkem768.encaps(pub)
        return ct + aes_kw(kdf(ss), dek)
    return rsa_oaep(pub, dek)              # legacy path

def unwrap(blob: bytes, priv) -> bytes:
    if len(blob) > 1088:                   # "must be ML-KEM"
        ct, wrapped = blob[:1088], blob[1088:]
        return aes_unkw(kdf(mlkem768.decaps(priv, ct)), wrapped)
    return rsa_oaep_decrypt(priv, blob)
```

Listing 5.1 The proposed change.

Which review comment identifies the most serious problem?

A. ML-KEM must not be combined with AES key wrap; the data-encryption key should be encrypted directly under the KEM.

B. The blob carries no algorithm identifier and infers the scheme from length; tag every blob with a versioned algorithm id.

C. RSA-OAEP must be removed in the same pull request, because a library with any classical path is not post-quantum.

D. The KDF step is unnecessary and should be deleted, because the ML-KEM shared secret is already 32 uniform bytes.

Correct answer: B
Inferring the scheme from length breaks the moment a second parameter set, a hybrid or a different KEM is added, and it makes rollback ambiguous. Tagged data lets old and new blobs coexist and supports staged rollout. KEM plus symmetric key wrap is the correct construction, keeping the RSA path for existing blobs is necessary during migration, and a KDF remains good practice for domain separation.

Question 6

Domain 6: Governance, Compliance and Assurance

DATA

A multinational must write one cryptographic standard covering a US defence subsidiary, EU operations and a commercial cloud product.

Source	Applies to	Parameter requirement	Key dates
NIST IR 8547	Everything	Any FIPS 203/204/205 set	Deprecate 112-bit RSA/ECC after 2030; disallow after 2035
CNSA 2.0	Defence subsidiary	ML-KEM-1024, ML-DSA-87, LMS/XMSS for firmware	Prefer by 2025; exclusive 2030 to 2033 by category
EU roadmap 2025	EU operations	Not specified	Start by end 2026; high-risk by 2030; all by 2035
Cloud customers	Cloud product	Hybrid TLS "as soon as possible"	Contractual, varies

Table 6.1 Requirements gathered for the standard.

Which policy structure is correct?

A. Apply CNSA 2.0 everywhere as the strictest regime, so one parameter set and one date apply to every system and audit.

B. Apply the EU roadmap everywhere as the latest-dated regime, so the programme carries the lowest cost and least disruption.

C. One standard with a global floor of the NIST and EU milestones plus a stricter CNSA 2.0 profile for the defence subsidiary.

D. A separate standard for each source, with each business unit choosing the regime that best fits its own risk appetite.

Correct answer: C
The regimes converge on 2030 and 2035; they differ in parameter sets and scope. A floor plus a profile keeps one policy, one inventory and one audit trail while meeting the strictest requirement where it applies. Forcing ML-KEM-1024 on the commercial estate adds cost with no compliance benefit, and choosing the loosest regime ignores harvest-now-decrypt-later exposure the dates do not cover.

Question 7

Domain 2: Post-Quantum Algorithms and Standards

A code reviewer finds that a hand-written ML-KEM decapsulation raises an exception when the re-encrypted ciphertext does not match the received one, and compares the two with an ordinary byte comparison. What is wrong?

A. Nothing; FIPS 203 requires decapsulation to reject an invalid ciphertext explicitly so the caller can retry.

B. The explicit failure and the variable-time comparison together give an attacker a chosen-ciphertext oracle.

C. Re-encryption is wasteful and should be removed; the recovered message can be returned directly as the secret.

D. The failure is fine but should be logged with the ciphertext so that repeated invalid inputs can be investigated.

Correct answer: B
FIPS 203 decapsulation never fails visibly: on a mismatch it returns a pseudorandom key derived from a secret value stored with the decapsulation key, selected in constant time. A visible failure or a timing difference lets an attacker probe malformed ciphertexts and recover the key over many queries. Re-encryption is required; it is the Fujisaki-Okamoto check.

Question 8

Domain 5: Protocols and Implementation

An operations team switches a public web property to ML-DSA-65 certificates and finds the bytes sent per handshake rise from 3.5 KB to about 27 KB, against a budget of 12 KB. Which combination of standardised or in-progress measures gets under budget without giving up post-quantum authentication?

A. Sign the leaf with ECDSA and the intermediate with ML-DSA, since only the intermediate must be post-quantum.

B. Replace ML-DSA with SLH-DSA-128s throughout the chain, because its public keys are only 32 bytes each.

C Use the smallest permitted leaf parameter set, omit the root, and adopt chain-reduction such as abridged certificates.

D Send only the leaf certificate and rely on clients fetching intermediates over HTTP from the authority information access URL.

Correct answer: C
Every signed element in the chain grows, including SCTs and OCSP. The sound response combines a smaller leaf set, not sending what the client already holds, and the IETF chain-reduction proposals written for exactly this problem. An ECDSA leaf leaves authentication classical, SLH-DSA has tiny keys but 7.8 KB signatures on every element, and AIA fetching adds a round trip and a plaintext dependency.

Question 9 Domain 4: Migration Planning and Crypto Agility

A programme office reports that 54 percent of live sessions now negotiate hybrid key exchange, up from 31 percent, while the count of unclassified assets has risen from 95 to 140 and no asset has yet had its classical-only option disabled. Which observation should lead the report?

A The programme is ahead of plan, because hybrid sessions rose by more than twenty points in a single quarter.

B The programme cannot yet claim to know its exposure, because unknowns are growing and no asset is complete.

C The programme should move its effort to lower-priority assets, because the high-priority work is largely done.

D The programme should pause hybrid rollout until the classical-only options have been disabled on every asset.

Correct answer: B
Progress numbers are only meaningful against a known estate. A growing unknown count means discovery is finding assets faster than the team classifies them, and some will be high priority. The hybrid share measures capability, not completion: an asset is migrated only when the vulnerable option can no longer be negotiated.

Question 10 Domain 5: Protocols and Implementation

A platform team runs OpenSSH 10 bastions with mlkem768x25519-sha256 first in the key-exchange list and Ed25519 host keys. A partner's SFTP client, OpenSSH 8.9, connects successfully. Which statement about that connection is correct?

A It is post-quantum protected, because the server lists a post-quantum algorithm first and the client accepted the connection.

B It is not post-quantum protected, because the 8.9 client offers only classical key exchange and negotiation picks what both support.

C It is post-quantum protected for confidentiality but not for integrity, because Ed25519 host keys are used for the transport MAC.

D It is not post-quantum protected, because the Ed25519 host key downgrades the negotiated key exchange to a classical group.

Correct answer: B

SSH negotiates the first key-exchange algorithm both sides support, so the server's preference cannot help a client that never offers a post-quantum option. Host keys authenticate the server and do not affect the key exchange chosen; they are still classical on every connection and are the next migration step, though not urgent in the harvest-now-decrypt-later sense.

Published by the CertLabz Certification Board, a division of CertLabz LLC. Companion to objectives document CPQS-001. https://certlabz.com/certifications/certified-post-quantum-cryptography-specialist.html