CERTLABZ CERTIFICATION BOARD

Certified Zero Trust Specialist (CZTS)

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 2: Identity and Access — DATA

Emma Wilson investigates an account takeover. The identity provider's sign-in log for one afternoon is below. Push MFA with number matching is enforced for everyone.

Time	Result	Source	Device	MFA	Session
13:02	Success	Home ISP, known city	Managed laptop, compliant	Push approved	S-1
13:41	Success	Hosting provider, other country	Unknown browser	Push approved	S-2
13:42	Token refresh	Hosting provider, other country	Unknown browser	n/a	S-2
14:10	Mail rule created: forward all to external address	Hosting provider	Unknown browser	n/a	S-2

Table 1.1 Sign-in and activity log. Session S-2 was later confirmed as the attacker.

The user approved a push "because it appeared right after I logged in". Which combination of controls would have stopped this?

A. A longer password policy and a conditional access rule that blocks every sign-in from outside the home country.

B. Number matching on the push prompt, which the organisation has already enforced for all users and all applications.

C Phishing-resistant authentication, plus a device policy that denies an unknown browser for a confidential resource.

D Disabling token refresh entirely, so that every session expires within minutes and must be re-authenticated.

Correct answer: C
The attacker used an adversary-in-the-middle page; the user completed a real MFA prompt for the attacker's session, and number matching does not help when the user types the number the attacker's page shows. FIDO2 fails on the look-alike origin. Independently, device policy denies an unknown browser for mail, and continuous evaluation would have revoked S-2 on the hosting-provider signal. Geo-blocking is a location control and fails against a proxy in the right country.

Question 2

Domain 4: Networks and Segmentation — CONFIG

James Carter reviews the cloud security groups protecting a three-tier application.

```
# security-groups.tf (excerpt)
sg_web:  ingress tcp/443  from 0.0.0.0/0
         ingress tcp/22   from 203.0.113.0/24        # office
sg_app:  ingress tcp/8080 from 10.20.0.0/16          # whole VPC
sg_db:   ingress tcp/5432 from 10.20.0.0/16          # whole VPC
sg_ops:  ingress tcp/22   from 203.0.113.0/24
all:     egress  any      to 0.0.0.0/0
```

Listing 2.1 Current rules. The VPC is 10.20.0.0/16 and also hosts unrelated workloads.

Which rewrite is consistent with zero trust?

A Reference the calling tier's group on each rule, replace port-22 rules with brokered access, and restrict egress to named destinations.

B Widen 10.20.0.0/16 to 10.20.0.0/8 so autoscaled instances are always covered, and keep the office range for administrative SSH.

C Insert a network firewall appliance between the three tiers and leave the security groups exactly as they are today.

D Move the database to a private subnet with no route to the internet and remove the sg_db security group altogether.

Correct answer: A
Every "from the whole VPC" rule is an implicit trust zone: anything in the VPC, including unrelated workloads, can reach the application and database tiers. Referencing the calling tier's security group is identity-based policy that survives autoscaling. Office-range SSH is location-based trust and is replaced by brokered access. Unrestricted egress is the exfiltration path a compromised worker uses.

Question 3

Domain 6: Visibility, Automation and Migration — DATA

Olivia Davis rates a business unit against the CISA Zero Trust Maturity Model 2.0 before writing its roadmap.

Pillar	Finding
Identity	SSO for 70% of apps; MFA by push; standing admin rights; quarterly manual reviews
Devices	Complete inventory with EDR; posture checked at sign-in only; BYOD unmanaged
Networks	Two firewall zones; flat inside each; VPN for remote; no east-west encryption
Applications and workloads	Internal apps trust the office network; 40 apps with local passwords; APIs behind a gateway with static keys
Data	Classification policy exists; three of nine systems tagged; encryption at rest with provider-managed keys

Table 3.1 Findings. Cross-cutting: logs centralised from IdP and EDR only; no automated response.

Which rating and first move are correct?

A. Every pillar rates Advanced because SSO, EDR and firewalls exist; the first move is to buy a ZTNA product for remote users.

B. Every pillar rates Traditional; the first move is a year of identity work, followed by a year of device work.

C. Networks rate Optimal because two zones exist; the first move is to add more firewall zones between departments.

D. Identity and Devices rate Initial, Networks and Applications Traditional; the first move is one protect surface, all pillars.

Correct answer: D

ZTMM 2.0 rates each pillar separately. Push MFA and standing privilege are Initial at best; two flat zones with a VPN are Traditional; apps that trust the network are Traditional; partial tagging is Initial. Applying all pillars to one protect surface proves the pattern and lifts every rating, which beats a year-per-pillar waterfall that leaves each resource half-protected.

Question 4

Domain 3: Devices and Endpoints

DATA

A policy engine receives the signals below for four devices requesting the same confidential application.

Device	Identity	Posture source	Signals	Freshness
L-01	TPM certificate	MDM + EDR	Compliant; EDR healthy	5 min
L-02	TPM certificate	MDM + EDR	MDM compliant; EDR reports agent tampered	MDM 6 h, EDR 3 min
P-03	None (personal phone)	App container only	Container policy applied; OS current	1 min
S-04	Software certificate on disk	None	No posture available	n/a

Table 4.1 Device signals at request time.

Which decision set is correct?

A. Allow all four devices, because each of them presented some form of device identity to the policy engine.

B. Allow L-01; deny L-02 and alert; allow P-03 only through the container with no download; deny S-04.

C. Allow L-01 and L-02 because both hold TPM certificates; deny P-03 and S-04 because neither is managed.

D. Deny all four devices until each user contacts the help desk and has the device manually re-enrolled.

Correct answer: B
Trust scales with assurance. A hardware-bound identity plus fresh healthy posture earns access. A fresh tamper signal must override a six-hour-old compliance statement; conflicting signals resolve conservatively. A personal phone gets container-scoped access at reduced trust. A copyable software certificate with no posture is no better than a password from an unknown machine.

Question 5

Domain 5: Applications, Workloads and Data

CODE

Ava Thompson reviews the authorisation middleware of an internal orders service behind a service mesh with mTLS.

```
def authorize(request):
    peer_ip = request.transport.peer_ip
    if peer_ip.startswith("10.30."):          # mesh CIDR
        return ALLOW
    token = request.headers.get("X-User")     # set by the gateway
    if token:
        return ALLOW
    return DENY
```

Listing 5.1 The middleware as deployed.

Which rewrite makes the service zero trust ready?

A. Narrow the CIDR check to the exact subnet of the gateway, so that only requests arriving through the gateway are allowed.

B. Keep the address check but add a rate limit per source, so that a compromised workload cannot flood the service.

C. Authorise on the caller's workload identity from the mTLS certificate and a verified user token, then check the action.

D. Remove every check, because the service mesh already encrypts all traffic between workloads with mutual TLS.

Correct answer: C
The service trusts a network location and an unsigned header, both of which any workload in the mesh can present. mTLS gives it a verified caller identity for free; a signed user token gives it the subject; policy then decides the action. Encryption is not authorisation, and a narrower CIDR is still a location.

Question 6

Domain 6: Visibility, Automation and Migration

DATA

After enabling automated response, the SOC reviews one week of triggers and outcomes.

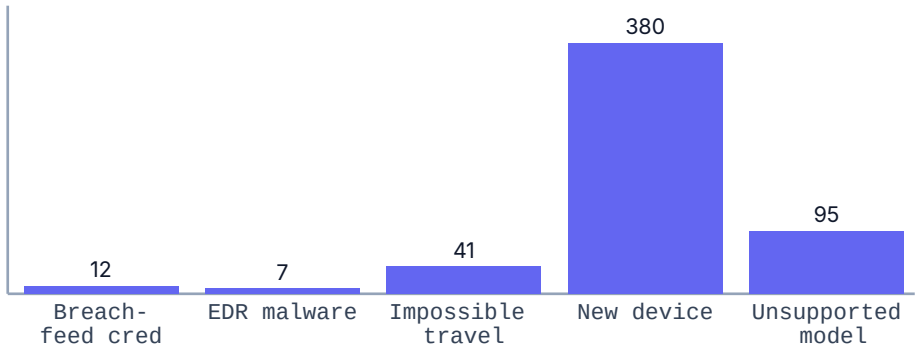

Figure 6.1 Signals received in one week, by type.

Signal	Automated action	Confirmed true positive
Credential in breach feed	Revoke sessions, force FIDO2 re-enrolment	12 of 12
EDR active malware	Quarantine device in inventory	7 of 7
Impossible travel	Revoke sessions	3 of 41
New device for user	Revoke sessions	1 of 380
Laptop model unsupported by MDM	Deny access	0 of 95 (data gap)

Table 6.1 Response mapping and outcomes.

Which change should the SOC make?

A. Keep hard responses for compromise signals; move the two behavioural signals to step-up; fix the MDM gap instead of denying.

B. Disable all automated response, because the large majority of the triggers in the week turned out to be false positives.

C. Apply session revocation to every signal type for consistency, so that analysts have one predictable response to learn.

D. Increase log retention from 30 to 90 days so that the false positives can be analysed in more depth next quarter.

Correct answer: A

Response strength should match signal confidence. Compromise signals earned their hard actions; behavioural signals with low precision should change the assurance required, not end the session; a denial caused by missing data is a data-quality problem and the fix is the data, not the policy.

Question 7

Domain 1: Zero Trust Principles and Architecture

A vendor proposes a "zero trust platform": a cloud broker with app connectors, policy written by IP range and user group, posture defined as "agent installed", and applications that keep their own login pages but trust the broker's source address. Which assessment is correct?

A. The proposal is a complete zero trust architecture, because a broker with connectors is what SP 800-207 describes.

B. The broker is a valid enforcement point, but the design keeps implicit trust in IP ranges, agent presence and source addresses.

C. The proposal must be rejected outright, because SP 800-207 does not permit cloud-hosted policy decision points.

D. The proposal is acceptable with one change: add a second broker in another region so that the control plane is resilient.

Correct answer: B
A product can implement a PEP without the organisation reaching zero trust. SP 800-207 requires dynamic policy on identity, asset state and context; applications that trust a source IP are the same implicit trust the broker was bought to remove. The fix is architectural: identity- and posture-based policy, signed identity headers verified by apps, and logs exported to central analytics.

Question 8 Domain 4: Networks and Segmentation

Four weeks of flow discovery for a payments application shows all high-volume flows covered by the draft allow policy, plus three unmatched flows: a monitoring host to the database twice a day, a batch job to the database at 03:00 on Sundays, and an unknown host reaching the database fourteen times. What is the right action before enforcing default deny?

A. Enforce immediately, because three unmatched flows out of millions is well within any reasonable error margin.

B. Add a rule allowing any source to reach the database port, so that none of the three flows can break in production.

C. Resolve each flow with its owner, add scoped rules for the legitimate two, and investigate the unknown host as a finding.

D. Abandon segmentation of this application, because the flow pattern is too complex to capture safely in policy.

Correct answer: C
Monitor mode exists to surface exactly these flows. Low-volume, scheduled flows are the ones that break a week after enforcement; an unknown host reaching the database is a finding, not an exception. Blanket rules would turn the exercise into the flat network with extra steps.

Question 9 Domain 2: Identity and Access

A platform team discovers that the payments API accepts a static key stored in a ConfigMap, that the orders service and a reporting service both mount it, and that the reporting mount was added "temporarily" two years ago. The intended policy is that only orders may call payments. Which migration is correct?

A. Rotate the static key and keep it in the ConfigMap, then remind the reporting team that the mount was meant to be temporary.

B. Move the key from the ConfigMap to a Secret object and consider the migration complete once the deployment is updated.

C Allow the reporting service in the policy as well, because it already holds the key and removing it might break reports.

D Issue attested workload identities to orders and payments, enforce the caller policy with mTLS, and remove the shared key.

Correct answer: D
A shared static key is a standing credential any workload with the mount can use, and "temporarily" is how implicit trust persists. Workload identity binds the caller to an attested identity; the policy names callers and operations. Storage hygiene for a shared secret does not change who can use it, and reporting gets its own scoped identity only if it has a legitimate need.

Question 10 Domain 5: Applications, Workloads and Data

A data platform team proposes per-user identity pass-through, tag-driven masking and just-in-time grants for a lake that mixes public and personal data. An objection is raised that per-user identity will break the dashboards that use the shared analytics role. What is the correct response?

A Keep the shared role for the dashboards and apply the per-user identity only to ad-hoc queries by analysts.

B Give each dashboard service its own workload identity with minimal grants, and retire the shared role on a date.

C Abandon the per-user identity design, because the dashboards matter more to the business than the audit trail.

D Configure the dashboards to run under the personal credentials of the analyst who most recently edited them.

Correct answer: B
Services are subjects too. A non-human identity with least-privilege grants keeps the dashboards working and keeps every query attributable. Leaving the shared role in place for "just the dashboards" is the classic standing exception that outlives the programme, and running services under a person's credentials breaks attribution in the other direction.

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