CERTLABZ CERTIFICATION BOARD

Certified Edge Computing Professional (CECP)

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 edge computing 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 1: Edge Computing Foundations — DATA

Emma Wilson decides where each workload should run on the device-edge-cloud continuum.

Workload	Key requirement
Machine safety shutdown	Respond within milliseconds
Model training	Heavy compute over large data
Cross-site reporting	Aggregate all sites
On-site video analytics	Real-time, high bandwidth locally

Table 1.1 Workloads and requirements.

Which placement is correct?

A. Safety shutdown in the cloud, training at the edge, reporting on one device, video analytics in the cloud, to centralise compute.

B. Safety shutdown at the edge, training in the cloud, cross-site reporting in the cloud, video analytics at the edge, matching each need.

C. All four in the cloud, because centralising every workload in one region is always the simplest and most reliable design choice.

D. All four on the device, because running everything locally removes any dependence on networks or the cloud in every situation.

Correct answer: B
A millisecond safety shutdown must run locally at the edge, heavy model training suits the cloud, cross-site reporting needs cloud aggregation, and real-time on-site video analytics belongs at the edge. Putting safety control or live analytics in a distant cloud adds fatal latency, and forcing training or cross-site aggregation onto one device is impractical.

Question 2 — Domain 3: Connectivity, 5G and MEC — DATA

James Carter chooses connectivity and design for four sites.

Site	Situation
Remote mine	No fibre, unreliable cellular
City AR service	Sub-20ms for many users
Large factory	Mobile assets, dense sensors
Any control loop	Must not stop if the link drops

Table 2.1 Sites and their situations.

Which set of choices fits best?

A. Mine on fibre, AR in one distant region, factory on public Wi-Fi, control loop over the WAN, standardising on the cheapest links.

B. Mine on satellite with local autonomy, AR on MEC over 5G, factory on private 5G, control loop running locally so a link loss cannot stop it.

C. Mine streaming all data live, AR processed overnight in batch, factory with no network, control loop calling the cloud each cycle for decisions.

D. All four sites on public cellular only, because a single connectivity type keeps the design simple regardless of each site's constraints.

Correct answer: B
A fibre-less mine needs satellite plus local autonomy, city-wide sub-20ms AR needs MEC over 5G near users, a factory with mobile assets suits private 5G, and any control loop must run locally so a link loss does not halt it. The other options assume connectivity that does not exist, add fatal latency, or make control depend on a network that can fail.

Question 3 — Domain 4: Edge Workloads and Orchestration — CONFIG

Olivia Davis reviews how software is rolled out to 500 edge sites.

```
rollout:
  strategy: all-sites-at-once
  health_checks: none
  rollback: manual, requires on-site visit
  # a bad build last month bricked 50 sites
```

Listing 3.1 The current rollout process.

Which change most reduces the risk of a bad update?

A. Roll out to all sites even faster, because shortening the window is what limits the impact of a bad build across the fleet.

B. Use a canary rollout to a few sites with automatic health checks and rollback before deploying to the rest of the fleet.

C. Send an engineer to each of the 500 sites during every rollout, so any failure can be corrected by hand on site immediately.

D. Stop updating the fleet entirely, because leaving the current software in place is the only way to avoid a bad update again.

Correct answer: B
A canary rollout to a few sites with automated health checks and automatic rollback catches a bad build before it reaches the whole fleet, which is exactly what would have contained last month's incident. Updating faster still hits every site, staffing 500 sites does not scale, and never updating leaves security and bugs unaddressed.

Question 4

Domain 5: Data, Analytics and AI at the Edge

DATA

Liam Anderson faces four edge data and AI problems.

Problem	Constraint
Vision model too big for device	Limited memory and compute
Patient data must stay on site	Privacy and regulation
Link drops and data is lost	Intermittent connectivity
Huge video bill to the cloud	Bandwidth cost

Table 4.1 Edge data and AI problems.

Which set of approaches fits these problems?

A. Send video to the cloud, upload raw patient data, accept the dropped data, and raise the video resolution to improve the models.

B. Quantise or prune the model, use federated learning, add store-and-forward buffering, and run inference locally sending only events.

C. Buy a data centre for the site, email the patient data, ignore the dropped data, and add more cameras to compensate for the loss.

D. Abandon edge AI, centralise all raw data, remove the sensors causing traffic, and stop improving the models to control the costs.

Correct answer: B
Optimising the model with quantisation or pruning fits it on the device, federated learning improves a shared model without moving raw patient data, store-and-forward buffering preserves data through link drops, and local inference sending only events cuts the video bill. The other options either violate privacy, lose data, or give up the benefits of the edge.

Question 5 — Domain 6: Security, Operations and Management — CONFIG

Ava Thompson audits the security of an edge device fleet.

```
device credentials : one shared password on all devices
boot               : no secure boot
data at rest       : unencrypted
updates            : unsigned, pulled over plain HTTP
```

Listing 5.1 Current edge device security.

An attacker with physical access to one device could compromise the whole fleet. Which fix set addresses the findings?

A. Keep the shared password but change it monthly, skip secure boot, and leave updates unsigned, since rotation alone secures the fleet.

B. Give each device a unique identity, enable secure boot, encrypt data at rest, and require signed updates over an authenticated channel.

C. Encrypt data at rest only, and leave the shared password, unsigned updates and missing secure boot unchanged as acceptable risks.

D. Disconnect the devices from the network entirely, because an air gap removes the need for identity, secure boot or signed updates.

Correct answer: B

Per-device identity stops one physical compromise spreading fleet-wide, secure boot blocks tampered firmware, encryption at rest protects a stolen device, and signed updates over an authenticated channel prevent malicious software being pushed. Rotating a shared password, encrypting alone, or assuming an air gap all leave critical gaps open.

Question 6 — Domain 2: Edge Architecture and Infrastructure — DATA

A company must deploy identical software to 500 unstaffed edge sites reliably.

Approach	Method
A	Send an engineer to configure each site by hand
B	Zero-touch provisioning with central fleet management
C	Email setup instructions to store staff
D	Configure each device uniquely with no standard image

Table 6.1 Deployment approaches.

Which approach fits, and why?

A. Approach A, because a trained engineer visiting each of the 500 sites is the most reliable way to guarantee a correct configuration.

B. Approach B, because zero-touch provisioning with central, declarative fleet management is the only approach that scales to 500 unstaffed sites.

C Approach C, because store staff following emailed instructions is the fastest and cheapest way to configure a large number of sites.

D Approach D, because configuring every device uniquely gives each site exactly what it needs without the constraints of a standard image.

Correct answer: B
Zero-touch provisioning with centralised, declarative fleet management is the only approach that scales to 500 unstaffed sites: devices auto-configure and join the managed fleet with no on-site work. Visiting every site or relying on store staff does not scale, and configuring each device uniquely makes updates and support unmanageable.

Question 7 Domain 1: Edge Computing Foundations

A factory needs a machine to stop within milliseconds of detecting a fault. Why is a cloud-only design unsuitable, and what is the right approach?

A Cloud-only is fine, because modern cloud regions respond fast enough for millisecond machine safety control in any factory setting.

B Round-trip latency and connectivity risk make cloud-only too slow and unreliable, so the safety control must run locally at the edge.

C The problem is only cost, because cloud compute is expensive, and moving the same cloud-only design to a cheaper region resolves it.

D Cloud-only is unsuitable because the cloud has no compute for control tasks, so a second cloud region should run the safety logic instead.

Correct answer: B
A safety shutdown needing millisecond response cannot depend on a network round trip to a distant cloud, and a dropped link would leave the machine uncontrolled. The control logic must run locally at the edge so it is fast and keeps working even when disconnected. The issue is latency and reliability, not cost or a lack of cloud compute.

Question 8 Domain 3: Connectivity, 5G and MEC

A design assumes the 5G link to the cloud is always available before a control decision is made. Why is this risky?

A It is not risky, because 5G networks are engineered never to drop and always deliver decisions with no meaningful latency at all.

B Any network can drop, so control logic should run locally and continue safely even when the link to the cloud is unavailable.

C The only risk is cost, because 5G data is expensive, and the control design itself is sound and certain to work whenever it is called.

D The risk is that 5G is too fast, so the control loop should be slowed down deliberately to match the speed of the cloud connection.

Correct answer: B
No network, including 5G, is guaranteed always available, so making a control decision depend on a live link risks halting operations when the link drops. Control logic should run locally and continue safely through a disconnection, syncing with the cloud when connectivity returns. The concern is reliability, not cost or excessive speed.

Question 9

Domain 5: Data, Analytics and AI at the Edge

A healthcare edge device must not send raw patient data off-site but should still help improve a shared model over time. What approach fits?

A. Upload all raw patient data to the cloud for training, because a central copy of the data always produces the most accurate model.

B. Use federated learning, sharing only model updates rather than raw data, so patient data stays on the device while the model improves.

C. Stop trying to improve the model, because the only way to train a model is to centralise the raw data it will be trained on.

D. Email a small sample of the raw patient records to the data science team, since a small enough sample is not a privacy concern.

Correct answer: B
Federated learning trains a shared model by exchanging model updates rather than raw data, so the patient data never leaves the device and privacy and regulation are respected while the model still improves. Uploading or emailing raw records, even a sample, breaches the requirement, and abandoning improvement is unnecessary.

Question 10

Domain 6: Security, Operations and Management

Why is the security of an edge fleet generally harder than that of a single central data centre?

A. It is easier, because edge devices are small and therefore inherently more secure than large central data centre servers are.

B. Many distributed devices sit in physically accessible, uncontrolled locations, which enlarges the attack surface and adds physical risk.

C. There is no difference, because a device in a shop or on a pole is exactly as physically protected as a server in a locked data centre.

D. It is only about cost, because edge security is simply a matter of buying more devices and has nothing to do with physical exposure.

Correct answer: B
An edge fleet spreads many devices across physically accessible, uncontrolled locations, so the attack surface is far larger than one hardened site and physical tampering or theft is a real risk. That is why per-device identity, secure boot, encryption and signed updates matter so much at the edge. Edge devices are not inherently more secure, and the difference is not merely cost.

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