CERTLABZ CERTIFICATION BOARD

Certified Multi-Cloud Professional (CMCP)

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 multi-cloud 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: Multi-Cloud Foundations and Strategy — DATA

Emma Wilson reviews four proposed reasons for going multi-cloud at her company.

Proposed reason	Claim
Data residency	Keep regulated data in a specific country
Resilience	Survive a single provider outage
Cost	Multi-cloud will automatically be cheaper
Best-of-breed	Use each provider's strongest service

Table 1.1 Reasons proposed for adopting multi-cloud.

Which claim is not a sound reason on its own?

A. Data residency, because keeping regulated data in a required country is never a valid reason to use more than one provider.

B. Resilience, because a single provider with several regions is always exactly as resilient as using two providers.

C. Cost, because multi-cloud adds complexity and does not automatically reduce spend, and often increases it.

D. Best-of-breed, because using each provider's strongest service is never worth the added operational effort involved.

Correct answer: C
Multi-cloud rarely reduces cost by itself; the extra providers add skills, security, monitoring and egress costs that must be managed, so "it will be cheaper" is not a sound standalone driver. Data residency, genuine cross-provider resilience and best-of-breed service selection are all legitimate reasons when tied to a real need.

Question 2

Domain 2: Identity, Access and Federation Across Clouds

CONFIG

James Carter reviews how an application on Cloud A authenticates to an API on Cloud B.

```
# Cloud A workload config
CLOUD_B_ACCESS_KEY = "AKIA...long-lived static key..."
CLOUD_B_SECRET     = "wJal...committed to the repo..."
# key created 2 years ago, shared across 4 services, never rotated
```

Listing 2.1 The current cross-cloud credential setup.

What is the best fix for this cross-cloud access?

A. Rotate the static key every night and keep it in the repository, since frequent rotation makes a committed key safe to use.

B. Use workload identity federation so the Cloud A workload obtains short-lived credentials for Cloud B with no static key.

C. Move the static key into an environment variable, because a key in an environment variable cannot be leaked from a repo.

D. Share one broader admin key across all four services, so there is a single credential to rotate instead of several keys.

Correct answer: B
Workload identity federation lets the Cloud A workload exchange its own identity for short-lived Cloud B credentials, removing the static key entirely, which is the leading cause of cloud key leaks. Nightly rotation of a committed key, moving it to an environment variable, or sharing one broad admin key all keep a long-lived secret and make the problem worse.

Question 3

Domain 3: Networking and Connectivity

DATA

Olivia Davis plans to peer two clouds' networks. Their address ranges are below.

Network	CIDR range
Cloud A production	10.10.0.0/16
Cloud A non-prod	10.20.0.0/16
Cloud B production	10.10.0.0/16
Cloud B non-prod	10.30.0.0/16

Table 3.1 Planned network address ranges.

Which problem will block connectivity, and what is the fix?

A. The non-prod ranges clash, so Cloud A non-prod must be renumbered to a range that does not overlap with Cloud B non-prod.

B. The production ranges overlap (both 10.10.0.0/16), so one must be renumbered before the networks can be peered or routed.

C. All four ranges are too small, so every network must be widened to a /8 before any peering between the clouds can work.

D. The ranges are fine as planned, so the peering will work immediately once a VPN or interconnect link is established.

Correct answer: B
Cloud A production and Cloud B production both use 10.10.0.0/16, and overlapping CIDR ranges make routing between the networks impossible. One of them must be renumbered to a non-overlapping range. The non-prod ranges do not clash, the prefix sizes are fine, and a link alone cannot fix overlapping addresses.

Question 4

Domain 4: Compute, Storage and Data Portability

DATA

Liam Anderson rates four design choices for how portable they are between clouds.

Design choice	Portability
Containers on Kubernetes	High
Open data formats (Parquet)	High
Provider-specific serverless functions	Low
Managed proprietary database with no export	Low

Table 4.1 Design choices and their portability.

A workload must move between two clouds with minimal rework. Which combination should it favour?

A. Provider-specific serverless functions and a managed proprietary database, because both are the most productive to build with.

B. Containers on Kubernetes with state in open formats, because both are portable and minimise per-provider rework on the move.

C. A managed proprietary database with no export path, because keeping all data in one managed service simplifies the migration.

D. Provider-specific serverless functions with open data formats, because the functions will run unchanged on either provider.

Correct answer: B
Containers on Kubernetes plus state in open formats are the portable combination, so the workload moves with minimal rework. Provider-specific serverless and a proprietary database with no export are exactly the low-portability choices to avoid when movement is the goal, and serverless functions do not run unchanged across providers.

Question 5 — Domain 5: Security, Governance and Cost Management — CONFIG

Ava Thompson sees the multi-cloud bill climbing and finds this state.

```
resources tagged with owner : 12%
accounts mapped to teams    : none
idle dev environments       : running 24/7
public storage buckets      : 3 found
```

Listing 5.1 Findings from a cost and governance review.

Which action should come first?

A. Buy large committed-use discounts immediately, because a long-term commitment is the fastest way to bring the bill down.

B. Turn off all monitoring and logging, because reducing the telemetry the providers charge for will cut the bill quickly.

C. Enforce a tagging and account structure so cost maps to owners and workloads, since you cannot optimise what you cannot attribute.

D. Move everything onto a single provider this week, because consolidating to one cloud always removes the cost problem entirely.

Correct answer: C

Attribution comes first: with only 12% of resources tagged and no account-to-team mapping, no one can see who owns the spend, so tagging and account structure are the foundation. Committing before you understand usage risks waste, disabling telemetry harms security, and a rushed single-cloud move is a huge project, not a first cost fix. The public buckets are a security issue to remediate in parallel.

Question 6 — Domain 6: Operations, Observability and Resilience — DATA

A team compares two replication choices for a multi-region service.

Choice	Behaviour
Synchronous replication	Writes confirmed in both regions before success
Asynchronous replication	Writes confirmed locally, copied to the other region after

Table 6.1 Replication choices.

The service needs the lowest possible data loss on a regional failure and can accept higher write latency. Which fits, and why?

A. Asynchronous replication, because copying writes after they are confirmed locally guarantees that no data is ever lost in a failure.

B. Synchronous replication, because confirming writes in both regions before success minimises data loss at the cost of higher latency.

C. Asynchronous replication, because it both minimises data loss and lowers write latency at the same time with no trade-off at all.

D Synchronous replication, because it lowers write latency and is therefore always the correct choice for any multi-region service.

Correct answer: B
Synchronous replication confirms each write in both regions before returning success, so a regional failure loses little or no data, which is the requirement, and the cost is higher write latency, which the service can accept. Asynchronous replication lowers latency but leaves a window of unreplicated writes (a higher RPO), so it does not meet the lowest-data-loss goal.

Question 7 Domain 1: Multi-Cloud Foundations and Strategy

A team proposes replicating every workload on all three major providers so they are "fully multi-cloud". What is the main problem with this plan?

A Nothing is wrong, because running every workload on every provider is the definition of a mature multi-cloud strategy.

B It multiplies operational cost and complexity with no business reason, when multi-cloud should be driven by specific needs.

C It is impossible, because a workload can never technically run on more than one cloud provider at the same time at all.

D The only issue is naming, because the strategy is sound and simply needs a clearer internal label to gain approval.

Correct answer: B
Blanket duplication across every provider is a classic anti-pattern: it multiplies cost, skills, security surface and monitoring for no business reason. Multi-cloud should be driven by specific goals such as residency, resilience or best-of-breed, applied where they matter, not applied to everything by default.

Question 8 Domain 3: Networking and Connectivity

An architecture repeatedly transfers large volumes of data from one cloud to another. Which cost factor most needs attention, and how is it best managed?

A DNS query fees, which dominate multi-cloud bills, so the fix is to reduce the number of DNS lookups the application makes.

B Compute licensing, which grows with transfer volume, so the fix is to buy more committed compute on the sending provider.

C Egress charges, which can dominate multi-cloud bills, so the fix is to place data and compute to minimise cross-provider transfer.

D Storage class fees, which rise with each transfer, so the fix is to move all data to the cheapest available archive tier.

Correct answer: C
Data transfer out (egress) is a leading hidden cost of multi-cloud and grows directly with cross-provider traffic, so the right response is to place data and the compute that uses it to minimise transfer between providers. DNS fees, compute licensing and storage classes are not what large repeated cross-cloud transfers primarily cost.

Question 9

Domain 2: Identity, Access and Federation Across Clouds

A departing contractor turns out to have had access to three clouds through separate local accounts, and one was missed at offboarding. What does this reveal, and what is the fix?

A. It reveals that local accounts per cloud are best practice, and the fix is simply to remind managers to offboard more carefully.

B. It reveals fragmented identity, and the fix is central federation so joiner-mover-leaver is reliable across every provider at once.

C. It reveals that MFA was missing, and the fix is to add MFA, which by itself would have removed the access when they left.

D. It reveals nothing important, because a single missed contractor account across three clouds is not a meaningful security risk.

Correct answer: B

Separate local accounts per cloud make deprovisioning error-prone, which is exactly how orphaned access happens. Central federation through one identity provider means a single offboarding action removes access everywhere, making joiner-mover-leaver reliable. MFA is valuable but does not remove access on departure, and lingering access across clouds is a real risk.

Question 10

Domain 6: Operations, Observability and Resilience

A company believes it can fail over to a second provider if its main cloud has a major outage, but has never tested it. Why is this a problem?

A. It is not a problem, because a documented failover plan is sufficient and testing it would only risk causing an outage.

B. Untested failover often fails when it is finally needed, so regular drills are required to prove the mechanism actually works.

C. The only problem is cost, because running a second provider is expensive, and the failover itself is certain to work when called upon.

D. The problem is that failover should never cross providers, so the plan should be replaced with a larger single-provider deployment.

Correct answer: B

Resilience that has never been exercised frequently fails at the moment it is needed, because configuration drift, data gaps and process gaps only surface under a real failover. Regular drills prove the mechanism and expose those gaps safely. A plan on paper is not proof, and cross-provider failover is a legitimate design when the business needs it.

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