CERTLABZ CERTIFICATION BOARD

# Certified Disaster Recovery Specialist (CDRS)

## 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 disaster recovery specialist 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: DR Foundations, BIA and Risk

DATA

Emma Wilson reviews the recovery objectives set for four systems.

| System | Set objective |
|---|---|
| Trading platform | RTO 5 min, RPO near zero |
| Internal wiki | RTO 5 min, RPO near zero |
| Email | RTO 4 hours, RPO 1 hour |
| Archive store | RTO 3 days, RPO 24 hours |

Table 1.1 Systems and their recovery objectives.

Which objective is most likely wrong, and why?

A. The trading platform, because a critical revenue system should never be given such a tight recovery objective in any circumstances.

B. The internal wiki, because a five-minute RTO and near-zero RPO are over-engineered and costly for a low-value internal tool.

C. The email objective, because four hours is far too fast and email should always be given a multi-day recovery objective instead.

D. The archive store, because any stored data must always be recovered within minutes regardless of how rarely it is actually accessed.

**Correct answer: B**
Objectives must be proportionate to business impact. A five-minute RTO with near-zero RPO on an internal wiki is over-engineered and expensive for its value. The trading platform genuinely warrants tight objectives, and email and the archive have reasonable, proportionate objectives.

## Question 2 — Domain 2: Backup Strategy and Data Protection — DATA

James Carter reviews the backup setup against a stated one-hour RPO.

| Aspect | Current state |
|---|---|
| Backup frequency | Once nightly |
| Copies | One, on the same storage array |
| Immutability | None; backups are writable |
| Restore tests | Never performed |

Table 2.1 Current backup setup.

Which finding most directly breaks the one-hour RPO?

A. The lack of restore testing, because untested backups always fail and that failure is what causes the recovery point to be missed.

B. The nightly backup frequency, because a once-a-day backup can lose up to a day of data, far exceeding the one-hour RPO required.

C. The single copy on the same array, because keeping one copy is the specific reason the recovery point objective cannot be met here.

D. The writable backups, because immutability is what determines the recovery point objective and its absence is what breaks the one hour.

**Correct answer: B**
The RPO is set by how much recent data you can afford to lose, which backup frequency directly controls. A nightly backup can lose almost a full day, far more than one hour, so it breaks the RPO. The single copy, missing immutability and lack of testing are serious resilience problems too, but they are not what defines the recovery point.

## Question 3 — Domain 3: DR Sites, Replication and Architecture — DATA

Olivia Davis compares DR options for a system needing an RTO of minutes and near-zero RPO.

| Option | Description |
|---|---|
| A | Cold site with nightly tape restore |
| B | Backups only, no second site |
| C | Active-active across separated sites with synchronous replication |
| D | Single site with extra servers |

Table 3.1 DR options.

Which option meets the objectives, and why?

A. Option A, because a cold site with nightly tape restores can be brought online within minutes to meet a near-zero data loss target.

B. Option B, because keeping good backups alone is always sufficient to achieve an RTO of minutes and a near-zero recovery point objective.

C. Option C, because active-active across separated sites with synchronous replication delivers both minutes-level RTO and near-zero RPO.

D. Option D, because adding extra servers in one site protects against a full site loss and therefore meets the stated objectives easily.

**Correct answer: C**
Minutes-level RTO with near-zero RPO requires a running standby and continuous replication, which active-active across geographically separated sites with synchronous replication provides. A cold site and tape restores are far too slow, backups alone cannot hit those objectives, and extra servers in one site do not survive a site loss.

## Question 4

Domain 4: DR Planning, Runbooks and Roles

CONFIG

Liam Anderson maps dependencies before setting recovery order.

```
App A  depends on  Database B
App A  depends on  Auth service C
Database B  depends on  (nothing)
Auth service C  depends on  (nothing)
```

Listing 4.1 System dependencies.

Which recovery order will bring App A up working?

A. Recover App A first, then Database B and Auth service C afterwards, since the front-end application is the highest business priority.

B. Recover Database B and Auth service C first, then App A, so that App A's dependencies are available when it starts.

C. Recover the systems in alphabetical order, because a consistent naming-based sequence avoids any need to consider dependencies.

D. Recover only App A, because once the application server is running the database and authentication service are no longer required.

**Correct answer: B**
App A depends on Database B and Auth service C, so those must be recovered first for App A to function when it starts. Recovering App A first leaves it non-functional because its dependencies are missing, alphabetical order ignores dependencies, and App A genuinely needs both the database and authentication to work.

## Question 5

Domain 5: Testing, Exercising and Recovery Execution

CONFIG

Ava Thompson reviews how a company tests its DR plan.

```
testing:
  tabletop: yearly discussion only
  full failover: never performed
  restore integrity checks: none
  actual RTO measured: no
```

Listing 5.1 Current DR testing.

What is the biggest weakness in this testing regime?

A. The tabletop is held only yearly, and simply running the same discussion more often each year would fully validate the recovery plan.

B. Failover has never actually been performed and recovery time and data integrity are never verified, so real recovery is unproven.

C. The tabletop involves too many people, and reducing it to a single planner reviewing the document would make the testing more reliable.

D. Nothing is seriously wrong, because a documented plan reviewed once a year is sufficient evidence that disaster recovery will work.

**Correct answer: B**

A yearly discussion never proves the technology works. Without an actual failover, measured recovery time and verified data integrity, the organisation only has an assumption that recovery will succeed. Running more discussions or shrinking the group does not test the real failover, and a document alone is not proof.

## Question 6

Domain 6: Resilience, Cyber Recovery and Improvement

DATA

After a ransomware attack, a team plans its recovery. The facts are below.

| Fact | Detail |
|---|---|
| Latest backups | May already contain the ransomware |
| Reachable backups | Some were encrypted by the attacker |
| Immutable copy | Exists from before the intrusion |
| Production | Still potentially compromised |

Table 6.1 Situation after the attack.

What is the correct recovery approach?

A. Restore the most recent backup straight into production quickly, because the fastest possible restore is what matters most after ransomware.

B. Identify a clean recovery point from the immutable copy and restore into an isolated environment, verifying it is clean before going live.

C. Pay the ransom and decrypt in place, because restoring from any backup after ransomware is unreliable and rebuilding is never worthwhile.

D. Restore each reachable backup in turn until one works, accepting that some will still be encrypted, and return whichever restores to production.

**Correct answer: B**
The latest backups may contain the ransomware and production may still be compromised, so the safe path is to identify a clean recovery point (the pre-intrusion immutable copy), restore into an isolated environment, and verify it is clean before returning to production. Restoring the latest backup or into a compromised environment risks immediate reinfection.

## Question 7 Domain 2: Backup Strategy and Data Protection

Why is an immutable or air-gapped backup copy considered essential protection against ransomware?

A. Because immutable backups are faster to restore than ordinary backups, and speed is the main defence against a ransomware attack.

B. Because modern ransomware often deletes or encrypts reachable, writable backups, and an immutable or offline copy cannot be altered by the attacker.

C. Because ransomware can never reach any backup, so making one copy immutable is simply a compliance formality with no real protective value.

D. Because immutability reduces the storage that backups consume, which is the primary reason to adopt it in a disaster recovery strategy.

**Correct answer: B**
Modern ransomware actively hunts and encrypts or deletes writable, reachable backups, so an ordinary online backup can be lost along with production. An immutable or air-gapped copy cannot be altered by the attacker, making it the reliable recovery path. Immutability is about integrity, not speed, storage or mere compliance.

## Question 8 Domain 3: DR Sites, Replication and Architecture

A company places its DR site in the same flood plain as its primary data centre to keep latency low. Why is this a problem?

A. It is not a problem, because keeping the two sites close together lowers replication latency and that benefit outweighs any other concern.

B. A regional event such as a flood could take out both sites at once, so a DR site needs enough geographic separation to survive the same disaster.

C. The only issue is cost, because two sites near each other are more expensive to run than a single larger site would be on its own.

D. The problem is that DR sites must always be adjacent to the primary, so the fix is to move them into the very same building for resilience.

**Correct answer: B**
A DR site exists to survive a disaster that takes out the primary, so sharing the same flood plain means one regional event can destroy both. The site needs geographic separation greater than the likely disaster radius. Low latency does not outweigh that, and placing the sites closer or in the same building makes it worse.

## Question 9

Domain 4: DR Planning, Runbooks and Roles

A DR plan names only the lead engineer as the recovery contact, and during an incident that engineer is unreachable. What flaw does this reveal?

A. None, because naming a single expert contact keeps the plan simple and one well-chosen person is sufficient for any recovery effort.

B. A key-person dependency: with no deputies or escalation chain, recovery stalls when that one person is unavailable during the incident.

C. The plan has too many contacts, and the fix is to remove even the lead engineer so responders improvise without a predefined contact list.

D. The problem is purely the contact's job title, and simply renaming the role in the document resolves the issue without any other change.

**Correct answer: B**

Depending on a single named person is a classic DR gap: if they are unreachable, recovery stops. The plan needs deputies and a clear escalation chain so someone can always act. Having one contact is not sufficient, removing contacts entirely is worse, and the job title is not the issue.

## Question 10

Domain 5: Testing, Exercising and Recovery Execution

After a failover test, the application is up but nobody checks the data, and corruption is discovered in production weeks later. What was missing from the test?

A. Nothing was missing, because confirming that the application started successfully is all that a disaster recovery failover test needs to verify.

B. Data integrity validation: recovery testing must verify that the data is complete and correct, not merely that a restore or failover completed.

C. A faster server, because the corruption was caused by slow hardware during the test and upgrading the hardware would have prevented it.

D. More frequent backups, because the only possible cause of the corruption was that the backups were taken too far apart from each other.

**Correct answer: B**

A failover test that only confirms systems are up does not prove the data is usable. Recovery testing must validate data integrity and correctness, which would have caught the corruption before it reached production. The issue is missing validation, not hardware speed or backup frequency.

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