CERTLABZ CERTIFICATION BOARD

Certified Site Reliability Professional (CSRP)

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 site reliability professional 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: SRE Foundations — DATA

Emma Wilson matches each reliability term to its meaning.

Term	Meaning
Service level indicator	(to choose)
Service level objective	(to choose)
Service level agreement	(to choose)
Error budget	(to choose)

Table 1.1 Reliability terms and their meanings.

Which set of meanings is correct?

A. The indicator sets the external contract, the objective measures the live signal, the agreement is the internal target, and the budget is the uptime promise itself.

B. The indicator is a measured signal, the objective is the internal target for it, the agreement is the external contract, and the budget is the allowed unreliability.

C. The agreement measures the live signal, the budget sets the internal target, the indicator is the contract, and the objective is the spare unreliability you may spend.

D. The budget is the measured signal, the agreement is the internal target, the indicator promises uptime, and the objective is the consequence written into a contract.

Correct answer: B
A service level indicator is a measured signal such as the fraction of successful requests, a service level objective is the internal target for that signal, a service level agreement is the external contract with consequences, and the error budget is how much unreliability the objective still allows. Only option B pairs each term with its real meaning.

Question 2

Domain 2: Monitoring and Observability

CONFIG

James Carter reviews an alerting setup that pages on raw machine signals.

```
alert_on : cpu above 80 percent
alert_on : any single error line
pages_per_night : many, mostly noise
user_impact_checked : no
```

Listing 2.1 The current alerting.

What is the soundest change?

A. Page on even more machine signals, such as memory and disk, because more alerts of every kind always mean the team will notice real problems sooner.

B. Turn off all alerting overnight and rely on user complaints in the morning, since a quiet pager is proof that the service is healthy and needs no watching.

C. Alert on user-facing symptoms tied to the objectives, such as error rate and latency, so the pager fires on real impact rather than on noisy raw signals.

D. Raise every threshold to 99 percent so almost nothing ever pages, on the assumption that fewer alerts of any kind must mean a more reliable running service.

Correct answer: C
Good alerts fire on user-facing symptoms tied to the service objectives, such as rising error rate or latency, so the team is paged for real impact rather than noisy raw signals. Adding more machine alerts increases noise, silencing the pager hides outages, and blanket high thresholds miss real problems.

Question 3

Domain 3: Incident Management

DATA

Olivia Davis matches each incident role to its responsibility.

Role	Responsibility
Incident commander	(to choose)
Communications lead	(to choose)
Operations lead	(to choose)
Scribe	(to choose)

Table 3.1 Incident roles and responsibilities.

Which set of responsibilities is correct?

A. The commander coordinates the response and decisions, the communications lead updates stakeholders, the operations lead applies the fixes, and the scribe records the timeline.

B. The commander records the timeline, the communications lead applies the fixes, the operations lead updates stakeholders, and the scribe coordinates the whole response effort.

C. The commander applies the fixes, the communications lead records the timeline, the operations lead coordinates the response, and the scribe keeps stakeholders informed of status.

D. The commander updates stakeholders, the communications lead coordinates the response, the operations lead records the timeline, and the scribe applies the fixes to the system.

Correct answer: A
The incident commander coordinates the response and makes decisions, the communications lead keeps stakeholders informed, the operations lead applies the technical fixes, and the scribe records the timeline. Only option A gives each role the responsibility it actually holds during an incident.

Question 4

Domain 4: Reliability Engineering

CONFIG

Liam Anderson reviews a design with no redundancy anywhere.

```
web_tier : one server, no spare
database : one node, no replica
retries : none on a failed call
health_checks : none configured
```

Listing 4.1 The current design.

What is the soundest first improvement?

A. Add a faster processor to the single web server, because a quicker machine removes the risk that any one component can bring the whole service down on its own.

B. Delete the database entirely and keep data only in memory, since a system with no stored data at all can never suffer a database outage during normal operation.

C. Leave the design as it is and simply document each single point of failure, on the view that writing the risks down is enough to keep the service reliable.

D. Remove the single points of failure by adding a second web server and a database replica, so the loss of one component no longer takes the whole service down.

Correct answer: D
Each tier here is a single point of failure, so the soundest first step is to add redundancy, such as a second web server and a database replica, so losing one component no longer takes the service down. A faster processor, an in-memory-only store, or documentation alone leaves the single points of failure in place.

Question 5

Domain 5: Automation and Toil Reduction

DATA

Ava Thompson matches each repetitive task to a sound way to reduce its toil.

Repetitive task	Reduction
Restarting a stuck service by hand	(to choose)
Copy-paste deployment steps	(to choose)
Editing config on each host	(to choose)
Renewing a certificate by hand	(to choose)

Table 5.1 Toil tasks and reductions.

Which set of reductions fits these tasks?

A. Do each task twice as often by hand, add a second reviewer to each step, and keep a longer paper log, since more manual effort is the safest way to cut toil over time.

B. A health check that restarts automatically, a scripted deploy pipeline, a config management tool, and automated certificate renewal, one fitting fix per task.

C. Ignore the stuck service, deploy straight to production without steps, edit every host by hand faster, and let the certificate lapse, since skipping work removes the toil.

D. Assign all four tasks to a single new hire permanently, because moving repeated manual work onto one person is what reducing toil across a reliability team truly means.

Correct answer: B
Toil is repetitive manual work that can be automated, so a self-healing health check handles restarts, a scripted pipeline handles deploys, a config management tool handles host settings, and automated renewal handles certificates. Doing the work more often, skipping it, or piling it on one person does not reduce toil.

Question 6

Domain 6: Culture and Operations

DATA

David Brown matches each operating practice to the benefit it brings.

Practice	Benefit
Blameless postmortem	(to choose)
On-call rotation	(to choose)
Written runbooks	(to choose)
Error budget policy	(to choose)

Table 6.1 Practices and their benefits.

Which set of benefits is correct?

A. Postmortems find causes without blame, a rotation shares on-call load fairly, runbooks give repeatable steps, and the budget policy balances new features against reliability.

B. Postmortems assign personal blame, the rotation puts one person on call forever, runbooks are left blank on purpose, and the budget policy simply bans every future release.

C Postmortems skip the causes to save time, the rotation removes on-call entirely, runbooks record only opinions, and the budget policy ignores reliability when shipping features.

D Postmortems punish the last engineer, the rotation doubles everyone hours, runbooks list unrelated tasks, and the budget policy freezes the whole product after a single outage.

Correct answer: A
A blameless postmortem finds causes without punishing people, an on-call rotation shares the load fairly, written runbooks give repeatable steps under pressure, and an error budget policy balances shipping features against protecting reliability. Only option A pairs each practice with its real benefit.

Question 7 Domain 2: Monitoring and Observability

A service feels slow to users, yet every dashboard shows green. What best explains the gap, and what helps?

A Nothing is wrong, because a set of green dashboards is complete proof that no user can be experiencing slowness anywhere in the service at this moment in time.

B The users must be mistaken, so the right response is to close the report unread and trust the existing dashboards rather than looking any deeper into the slow requests.

C The dashboards track only known signals; tracing individual requests reveals where time is spent and surfaces the slow path the aggregate charts are hiding from view.

D The dashboards are the wrong color, so switching every chart from green to red would by itself make the underlying slow requests faster for the affected users.

Correct answer: C
Dashboards show known aggregate signals and can stay green while a subset of requests is slow. Observability tools such as request tracing show where time is spent on individual requests and surface the slow path the averages hide. Dismissing users or recoloring charts does not find the cause.

Question 8 Domain 3: Incident Management

During an outage several engineers each change things at once and nobody knows who is in charge. What practice fixes this?

A Let every engineer keep making independent changes at the same time, because more people acting at once is always the quickest route to resolving any live outage.

B Tell everyone to stop and wait quietly for the problem to resolve itself, since an outage left completely alone will reliably recover without any coordinated action.

C Ask each engineer to post a personal theory in chat and take no action at all, on the view that discussion by itself restores a service faster than any change would.

D Designate one incident commander to coordinate the response, so changes are agreed and tracked and the team acts together instead of colliding during the outage.

Correct answer: D
Uncoordinated simultaneous changes make outages worse and hide what actually helped. Designating a single incident commander to coordinate the response means changes are agreed and tracked and the team acts together. Free-for-all changes, passive waiting, and theory-only chat do not restore the service.

Question 9 Domain 4: Reliability Engineering

When a service fails, every client retries instantly and in a tight loop, which makes the outage worse. What safeguard helps most?

A. Use exponential backoff with jitter and a circuit breaker, so clients wait longer between spaced-out retries and stop hammering a service that is already failing.

B. Have every client retry even faster and with no pause at all, because sending far more requests to a failing service is the surest way to bring it back sooner.

C. Tell each client to retry forever in the same tight loop, on the view that endless immediate retries must eventually force the overloaded service to recover on its own.

D. Remove all error handling from the clients entirely, so that a failed call is silently ignored and no retry, backoff, or circuit breaker is ever needed anywhere at all.

Correct answer: A
Instant synchronized retries pile load onto an already failing service, a pattern known as a retry storm. Exponential backoff with jitter spaces retries out, and a circuit breaker stops calls to a failing service so it can recover. Faster retries or endless loops deepen the outage.

Question 10 Domain 5: Automation and Toil Reduction

An on-call engineer spends hours each week on the same manual cleanup task. What is the soundest way to handle it?

A. Keep doing the manual cleanup by hand every week forever, because a task that someone can already complete by hand never needs to be automated by a reliability team.

B. Hand the weekly cleanup to a brand new engineer each week, on the view that rotating the same manual toil among more people counts as genuinely reducing that toil.

C. Measure how much time the task costs, then automate the repetitive steps so the cleanup runs on its own and the on-call engineer is freed for higher-value work.

D. Simply stop doing the cleanup and hope the leftover data causes no trouble later, since ignoring a recurring maintenance task is the fastest way to remove its toil.

Correct answer: C
Repetitive manual work is toil, and the sound response is to measure its cost and automate the repetitive steps so it runs on its own. That frees the on-call engineer for higher-value work. Doing it by hand forever, rotating it among people, or abandoning it does not remove the toil safely.

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