

Índice

CCAR-P · Examen de práctica

1

CCAR-P · Examen de práctica

63 preguntas · 120 minutos · corte 720/1000

Material no oficial. Preguntas originales escritas contra los objetivos publicados en la guía del examen. No reproducen contenido del examen real. [Ver aviso completo.](#)

Las respuestas están en [respuestas.md](#). Para practicar con corrección automática y cronómetro, usá el [simulador](#).

1. Selecciona 2

D2 · Modelos, prompting e ingeniería de contexto

An application's system prompt is stable but its cache hit rate sits at only 40%. Which TWO changes are most likely to raise it?

- A. Move the session identifier from the top of the prompt to the end
- B. Shorten the static policy block so that it takes up rather fewer tokens
- C. Keep the tool definitions in a fixed order across every request
- D. Raise the maximum output tokens allowed on each of the requests
- E. Send all requests from a single region so they stay grouped

2. Selecciona UNA

D7 · Productividad del desarrollador

Reviews of AI-authored pull requests almost never raise findings, yet defects keep reaching production. What is the most likely cause?

- A. The reviewer shares the context and assumptions that produced the code
- B. The review prompt is too short to describe what quality means here
- C. The model used for the review is smaller than the one used for authoring
- D. The team merges before the automated review has finished running

3. Selecciona UNA

D7 · Productividad del desarrollador

An agent in production starts calling an unexpected tool on a minority of requests. What does the on-call engineer do first?

- A. Read traces of the affected sessions to locate where the choice diverges
- B. Rewrite the tool descriptions, since ambiguity is the most usual explanation
- C. Disable the tool until someone can reproduce the behaviour in staging
- D. Roll the whole service back to the release from the previous quarter

4. Selecciona UNA

D7 · Productividad del desarrollador

A developer joins a team that works with AI assistance daily. What gets them productive fastest?

- **A.** The shared project configuration and standards the team already uses
- **B.** A training session covering prompt techniques and their trade-offs
- **C.** Read-only access for the first weeks, until they know the codebase
- **D.** Their own configuration, so they can work the way they prefer

5. Selecciona 2

D1 · Diseño de solución y arquitectura

A fixed workflow for supplier onboarding now fails on roughly a fifth of cases, where staff step in manually to finish the job. Which TWO signals would justify moving this stage to an autonomous agent?

- **A.** The number of onboarding requests per week has roughly doubled since launch
- **B.** The manual interventions follow no repeatable pattern the team can list
- **C.** The steps differ per supplier and only become clear part-way through
- **D.** Leadership wants the solution to look more advanced to the board
- **E.** The cost per request has risen along with the added volume

6. Selecciona 2

D1 · Diseño de solución y arquitectura

A supervisor coordinates four specialist agents in a regulated claims process. Which TWO responsibilities belong to the supervisor rather than to the specialists?

- **A.** Deciding the order the specialists run in and when to stop the chain
- **B.** Applying the rules that determine whether a claim is fraudulent
- **C.** Recording the ordered outcome of every step for the audit trail
- **D.** Formatting each specialist's own output for its internal use
- **E.** Choosing which model version each specialist should call at runtime

7. Selecciona 2

D1 · Diseño de solución y arquitectura

An architect is splitting a long analysis task into sequenced subtasks. Which TWO practices make the decomposition real rather than cosmetic?

- **A.** Passing structured output between steps so each one can be checked
- **B.** Giving every step only the context that step actually needs
- **C.** Running each of the steps on the largest model in the family
- **D.** Keeping the original prompt intact and appending the steps
- **E.** Merging the steps back into one call once quality has improved

8. Selecciona 2

D1 · Diseño de solución y arquitectura

A sponsor asks the architect to justify the investment in business terms. Which TWO framings belong in that justification?

- **A.** Hours of manual handling removed from the process each month
- **B.** The number of tools currently registered in the agent's catalogue
- **C.** Improvement measured against a contractual response-time SLA
- **D.** The parameter count of the model the team ended up selecting
- **E.** The number of prompt variants the team wrote during the build

9. Selecciona 2

D1 · Diseño de solución y arquitectura

Which TWO conditions indicate that a single augmented model call is enough, with no workflow or agent around it?

- **A.** The task is one transformation whose context can be supplied up front
- **B.** Several specialists must contribute before the output is complete
- **C.** Any reference material needed fits alongside the request itself
- **D.** The next step depends entirely on what the previous step happened to return
- **E.** Each step has to be recorded separately for an external auditor

10. Selecciona UNA

D1 · Diseño de solución y arquitectura

A quoting solution was designed as input, processing, and output, with no further stage. Six months in, nobody can say whether quote accuracy has improved or degraded. What is missing from the architecture?

- **A.** A feedback stage that captures outcomes into an evaluable dataset
- **B.** A caching layer that lowers the cost of repeated quotes
- **C.** A second model that reviews every quote before it goes out
- **D.** A dashboard reporting quote volume and average processing time

11. Clasifica cada fila

D1 · Diseño de solución y arquitectura

For each situation, identify who should decide the next step.

#	Escenario	Tu respuesta
1	The five stages of a monthly regulatory report that has not changed in two years.	
2	Which log query to run next while diagnosing a production outage.	
3	Whether to release a payment that exceeds the agreed approval threshold.	
4	The order in which four specialist agents run inside a regulated pipeline.	
5	How many sources to consult before answering an open research question.	

Opciones: the architect, at design time · the model, at run time · a human reviewer, per case

12. Selecciona UNA

D1 · Diseño de solución y arquitectura

An assistant handles two kinds of request: simple status lookups that arrive constantly, and complex disputes that need deep analysis. Which design serves both?

- **A.** A router that classifies the request and sends it down the fitting path
- **B.** One agent with every tool available, deciding what to do per request
- **C.** A single prompt long enough to cover lookups and disputes together
- **D.** Two separate products, each with its own interface and its own login

13. Selecciona UNA

D1 · Diseño de solución y arquitectura

A pipeline extracts data, transforms it and writes it to a system of record. Where does validation belong?

- **A.** After each step, so a bad result never reaches the following one
- **B.** At the very end, right before the write to the system of record
- **C.** Only at the start, validating the input document that arrived
- **D.** Nowhere: the model is instructed to check its own work as it goes

14. Selecciona 2

D1 · Diseño de solución y arquitectura

A due-diligence review must cover environmental, legal and financial angles. Which TWO facts justify separate agents rather than one?

- **A.** Each angle needs its own sources, tools and specialist context
- **B.** The three reviews do not depend on each other and can run at once
- **C.** The client is paying enough to cover the extra orchestration cost
- **D.** The board asked for the most advanced architecture on the market
- **E.** The volume of reviews per quarter has grown steadily since last year

15. Clasifica cada fila

D1 · Diseño de solución y arquitectura

For each situation, identify the architecture that fits.

#	Escenario	Tu respuesta
1	Turning each inbound invoice into a ledger entry, with the vendor record supplied.	
2	A quarterly audit pack produced through the same six steps every quarter.	
3	Chasing why a shipment was mis-routed, where each answer opens the next query.	
4	A plant expansion review needing environmental, legal and civil assessments.	
5	Rewriting each job advert into the company tone, with the style guide attached.	

Opciones: single augmented call · fixed workflow · autonomous agent · multi-agent system

16. Selecciona 2

D2 · Modelos, prompting e ingeniería de contexto

Which TWO practices make the guardrails written into a system prompt more likely to hold?

- **A.** Stating explicitly what takes precedence when two instructions conflict
- **B.** Placing the guardrails at the start or end, apart from reference material
- **C.** Writing the guardrails in a noticeably more emphatic and commanding tone
- **D.** Repeating each guardrail once inside every section of the prompt
- **E.** Assuming the guardrail also covers the cases it does not mention

17. Selecciona 2

D2 · Modelos, prompting e ingeniería de contexto

A team must cut token spend on a retrieval-heavy assistant without losing answer quality. Which TWO moves fit that goal?

- **A.** Retrieve fewer but better-targeted passages for each question asked
- **B.** Drop the citations from the answer so the output comes out shorter
- **C.** Cache the stable instruction block that every request shares already
- **D.** Switch every workload in the product to the smallest model available
- **E.** Truncate each retrieved passage to its first two sentences

18. Selecciona 2

D2 · Modelos, prompting e ingeniería de contexto

Which TWO tasks are the best candidates for chain-of-thought prompting?

- **A.** Pulling the invoice number and the issue date out of a scanned paper document
- **B.** Deciding whether a contract clause conflicts with three internal policies
- **C.** Translating a paragraph with a supplied glossary of approved terms
- **D.** Estimating a project's risk from several partly conflicting signals
- **E.** Converting a list of names into a fixed JSON structure for import

19. Selecciona UNA

D2 · Modelos, prompting e ingeniería de contexto

A summarisation feature runs on the largest model in the family. Evaluation shows a mid-tier model matches its quality on this task at a third of the cost. What should the architect do?

- **A.** Move the feature to the mid-tier model and keep the evaluation running
- **B.** Stay on the largest model, since summary quality varies a lot with the input
- **C.** Split the traffic evenly between both models to hedge against a drop
- **D.** Move to the smallest model, since the mid-tier one is not the cheapest

20. Clasifica cada fila

D2 · Modelos, prompting e ingeniería de contexto

For each symptom, identify the technique that fixes it.

#	Escenario	Tu respuesta
1	Output format is close but inconsistent from one run to the next.	
2	A critical rule buried in 12,000 tokens of context is followed erratically.	
3	The same 9,000-token policy block is re-sent and re-charged on every call.	
4	Answers about the middle of a large document are unreliable and expensive.	
5	The house table layout is described in prose but never reproduced exactly.	

Opciones: few-shot examples · move the rule to an edge · prompt caching · retrieve only what is relevant

21. Selecciona UNA

D2 · Modelos, prompting e ingeniería de contexto

A classification task has eight well-defined categories and clear written definitions for each. Accuracy is already high without examples. What follows?

- **A.** Keep it zero-shot and spend the tokens on the definitions instead
- **B.** Add few-shot examples anyway, since they always improve accuracy
- **C.** Add one example per category to make the prompt more consistent
- **D.** Move to a larger model to protect the accuracy already achieved

22. Selecciona UNA

D2 · Modelos, prompting e ingeniería de contexto

A support conversation runs long and the earliest messages start falling out of the window. What preserves quality?

- **A.** Summarise the older turns and keep the decisive facts verbatim
- **B.** Let the oldest turns drop, since recent context matters more
- **C.** Move to the model with the largest window and keep everything
- **D.** Start a new conversation whenever the window gets close to full

23. Selecciona 2

D3 · Integración

A support agent has a tool that can delete customer records. Which TWO changes reduce the blast radius without removing the capability from the platform?

- **A.** Scope the tool's credentials so it can only reach the support dataset
- **B.** Require an explicit approval step before any delete call executes
- **C.** Add the word 'careful' to the tool description so that the model hesitates
- **D.** Log every delete after the fact so the team can review them weekly
- **E.** Rename the tool so the model is less likely to select it by mistake

24. Selecciona 2

D3 · Integración

An architect is choosing between MCP and direct API integration for a new internal system. Which TWO facts about the situation point to MCP?

- **A.** Three separate AI applications will need the same access to this system
- **B.** The team owning the system will keep adding and retiring endpoints
- **C.** The integration is a single deterministic call inside a nightly job
- **D.** The system is only ever queried by one service the team owns
- **E.** Only the data engineering team will ever query this system directly

25. Selecciona 2

D3 · Integración

A platform team is defining what every agent session must emit for observability. Which TWO elements matter most for finding failures later?

- **A.** A correlation identifier that ties every step of one session together
- **B.** The retrieved context and tool inputs and outputs for each step
- **C.** The colour-coded severity level of each message shown to the end user today
- **D.** A copy of the full prompt and response stored for all sessions
- **E.** The name of the engineer who last deployed the service version

26. Selecciona 2

D3 · Integración

A RAG assistant returns passages that are topically related but rarely answer the question. Which TWO changes address this directly?

- **A.** Add a re-ranking stage that reorders candidates before they are used
- **B.** Tune chunk boundaries so each passage carries one complete idea
- **C.** Increase the number of passages retrieved from five up to fifty
- **D.** Raise the temperature so the model paraphrases the sources more
- **E.** Replace the vector index with a keyword index built over document titles

27. Selecciona 2

D3 · Integración

An assistant reads documents that belong to individual users. Which TWO conditions make per-user credentials mandatory rather than optional?

- **A.** Each user may only see the subset of documents assigned to them
- **B.** A prompt injection in a document could redirect what the agent asks for
- **C.** All documents in the corpus are already public in the company
- **D.** The assistant is used by a single administrator with full access
- **E.** The corpus is small enough to fit entirely in the context window

28. Selecciona 2

D3 · Integración

Which TWO symptoms suggest that a RAG pipeline's chunking strategy is the problem, rather than the model or the prompt?

- **A.** Answers cite a rule but miss the exception stated in the next paragraph

- **B.** Retrieved passages start mid-sentence and end before the point lands
- **C.** The same question returns different answers on repeated and identical runs
- **D.** Latency rises sharply whenever the corpus grows past a certain size
- **E.** The assistant refuses questions that fall outside the indexed corpus

29. Selecciona UNA

D3 · Integración

A team wants to expose forty internal endpoints to an agent. Tool-selection accuracy in testing drops sharply past roughly a dozen tools. What is the sound design response?

- **A.** Group the endpoints into task-oriented tools and load them on demand
- **B.** Register all forty tools and rely on a more capable model to choose
- **C.** Register all forty and instruct the model to think before selecting one
- **D.** Expose one generic tool that accepts the endpoint name as a parameter

30. Clasifica cada fila

D3 · Integración

For each failure, identify which stage of the retrieval pipeline is at fault.

#	Escenario	Tu respuesta
1	A clause is retrieved without the definition it depends on, which sits two sections away.	
2	Answers still quote a policy that was replaced during last night's refresh.	
3	An exact serial number returns visually similar but incorrect products.	
4	A contractor receives a passage from a document restricted to full employees.	
5	Passages consistently break in the middle of a numbered list of conditions.	

Opciones: chunking · indexing · search strategy · access control

31. Selecciona UNA

D3 · Integración

A retrieval index holds documents from many client organisations. What keeps one client from ever seeing another's content?

- **A.** Filter by tenant at query time using metadata on every chunk
- **B.** Instruct the model to answer only from the requesting client's files
- **C.** Keep one index and rely on relevance to surface the right documents
- **D.** Ask the model to check the client name before quoting a passage

32. Selecciona UNA

D3 · Integración

A knowledge base gets edits all day long, and the assistant must reflect them within minutes. What indexing approach fits?

- **A.** Incremental indexing triggered by each change, with a periodic full rebuild
- **B.** A nightly full rebuild, which is simpler to operate and to reason about
- **C.** Cache answers for an hour so the load on the index stays predictable
- **D.** Retrieve straight from the source database and skip the index entirely

33. Selecciona 2

D3 · Integración

An agent confuses three tools whose descriptions all begin with 'Looks up records'. Which TWO changes fix the selection?

- **A.** State what each tool returns and where its boundary sits
- **B.** Merge or retire the tools whose scope genuinely overlaps
- **C.** Order the catalogue so the most used tool appears first
- **D.** Split each tool further so every case has a dedicated one
- **E.** Have the agent call all three and compare what comes back

34. Clasifica cada fila

D3 · Integración

For each need, identify the connection mechanism that fits.

#	Escenario	Tu respuesta
1	Four internal products need the same standard access to the payroll system.	
2	A scheduled job copies yesterday's transactions into the data warehouse.	
3	A carrier's agent and a retailer's agent negotiate delivery windows.	
4	A new compliance archive should be reachable by any future agent.	
5	A billing service calls the model once to tag each incoming invoice.	

Opciones: MCP server · direct API integration · agent-to-agent protocol

35. Selecciona 2

D4 · Evaluación, testing y optimización

A team is adding production evaluation to a live assistant. Which TWO practices give a trustworthy ongoing signal?

- **A.** Score a sampled slice of real production outputs on a fixed schedule
- **B.** Capture downstream user corrections and feed them into the dataset
- **C.** Re-run the pre-launch suite unchanged every quarter and compare
- **D.** Ask the support team whether the assistant feels better than last month
- **E.** Count how many sessions ended without the user asking again

36. Selecciona 2

D4 · Evaluación, testing y optimización

A latency-sensitive feature must get faster without a measurable quality loss. Which TWO levers fit?

- **A.** Cache the stable prefix so repeated content is not reprocessed
- **B.** Drop chain-of-thought on the steps that are simple extraction
- **C.** Cut the evaluation suite so releases can ship more frequently
- **D.** Reduce the maximum output tokens on every endpoint by half
- **E.** Move all of the reasoning-heavy steps to the smallest available model

37. Selecciona 2

D4 · Evaluación, testing y optimización

Quality has slipped over three weeks with no deploys. Which TWO findings would confirm input drift as the cause?

- **A.** A new customer segment started sending queries in a second language
- **B.** Average query length has doubled since the start of the period
- **C.** The billing total for the month came in above the usual range
- **D.** The p95 latency of the endpoint has stayed flat throughout
- **E.** The number of daily sessions has grown steadily week over week

38. Selecciona 2

D4 · Evaluación, testing y optimización

A team is standing up an LLM-as-judge to score tone at scale. Which TWO steps keep the judge trustworthy over time?

- **A.** Calibrate its scores periodically against a sample rated by experts
- **B.** Write the rubric from the same guidelines the humans apply
- **C.** Use the same model version that generated the output that is being scored
- **D.** Let the judge revise its own rubric whenever agreement drops
- **E.** Score only the outputs that the assistant flagged as uncertain

39. Selecciona 2

D4 · Evaluación, testing y optimización

A canary is running for a new model version. Which TWO signals should gate the decision to widen the rollout?

- **A.** Task-level scores on the evaluation suite for the canary traffic
- **B.** Guardrail violations and refusal rates compared with the baseline
- **C.** The number of engineers who reviewed the release notes ahead of time
- **D.** The reduction in cost per request the newer version delivers
- **E.** The size of the canary slice relative to total production traffic

40. Selecciona UNA

D4 · Evaluación, testing y optimización

A prompt change scored six points better offline. The team wants production evidence before switching everyone. Which mechanism answers the question they are actually asking?

- **A.** An A/B test splitting live traffic between the current and new prompt
- **B.** A canary that exposes the new prompt to a small slice and watches errors
- **C.** A third offline run of the same suite to confirm the six-point gap
- **D.** A review session where senior staff compare sample outputs side by side

41. Clasifica cada fila

D4 · Evaluación, testing y optimización

For each metric, identify what it tells you about a production system.

#	Escenario	Tu respuesta
1	Share of answers where retrieval found no grounded passage.	
2	Number of formal complaints filed through the support portal.	
3	Average tokens consumed per completed session.	
4	Percentage of requests that returned a server error.	
5	Average relevance score of the passages passed to the model.	

Opciones: leading indicator of quality · lagging indicator of quality · cost signal · availability signal

42. Selecciona UNA

D4 · Evaluación, testing y optimización

A summarisation feature needs a quality metric that can run on every output. What is the sound choice?

- **A.** A rubric scored against the source, checking coverage and faithfulness
- **B.** Length of the summary compared with the length of the source text
- **C.** How closely the summary matches one reference summary word for word
- **D.** How many users clicked to expand the full document after reading

43. Selecciona UNA

D4 · Evaluación, testing y optimización

Prompts change several times a week. Running the full evaluation suite on each change takes two hours. What is the sensible arrangement?

- **A.** A fast subset on every change and the full suite before each release
- **B.** The full suite on every change, since quality cannot be compromised
- **C.** The full suite once a week, regardless of how many changes shipped
- **D.** Manual spot checks per change, with the suite reserved for incidents

44. Selecciona 2

D4 · Evaluación, testing y optimización

An assistant will operate in a regulated insurance context. Which TWO components must the pre-launch evaluation include?

- **A.** A set labelled by experts covering routine and high-risk cases
- **B.** Adversarial cases probing policy violations and prompt injection
- **C.** A comparison of the model against the benchmarks of competing vendors
- **D.** A demonstration to the executive committee with chosen examples
- **E.** A satisfaction survey run with users after the system is launched

45. Selecciona 2

D5 · Gobernanza, seguridad y riesgo

A drafting assistant must never state a delivery date that is not in the logistics system. Which TWO controls enforce that?

- **A.** Block any response containing a date that no lookup call returned
- **B.** Require the model to cite the record its date came from, or omit it
- **C.** Instruct the assistant in the prompt to avoid inventing dates
- **D.** Review a weekly sample of sent messages for invented delivery dates
- **E.** Lower the temperature so the assistant improvises less on its own

46. Selecciona UNA

D5 · Gobernanza, seguridad y riesgo

An agent that books travel has a tool that can cancel existing reservations. Where does the human approval gate belong?

- **A.** Before a cancellation executes, while searches and drafts run freely
- **B.** Before every tool call the agent makes during the booking session
- **C.** After the cancellation, so that a person can reinstate it if it was wrong
- **D.** On a random tenth of sessions, sampled for quality assurance

47. Selecciona UNA

D5 · Gobernanza, seguridad y riesgo

A customer exercises their right to erasure. The company keeps full conversation traces for auditing. How is that resolved at design time?

- **A.** Separate identity from content so the person can be erased from traces
- **B.** Delete every trace touching that customer, accepting the audit gap
- **C.** Refuse the request, since audit obligations override the erasure right
- **D.** Keep the traces and simply stop associating them with any customer

48. Selecciona UNA

D5 · Gobernanza, seguridad y riesgo

A guardrail is blocking legitimate requests along with the harmful ones, and users are complaining. What is the correct response?

- **A.** Measure both error types on labelled cases and tune where the harm is worse
- **B.** Loosen the guardrail until users stop complaining about the blocks
- **C.** Keep it as is: blocking too much is always safer than blocking too little
- **D.** Move the check into the prompt so the model can use its own judgement

49. Selecciona UNA

D5 · Gobernanza, seguridad y riesgo

An architecture routes customer data through three third-party services. What does the risk assessment have to cover?

- **A.** Each provider's data handling, plus what happens when one of them fails
- **B.** The provider with the largest share of the traffic that flows through
- **C.** Only the provider that stores data, since the others merely pass it on
- **D.** The contractual penalties agreed with each of the three providers

50. Selecciona UNA

D5 · Gobernanza, seguridad y riesgo

An automated decision affects whether a person receives a benefit. What does the design owe that person?

- **A.** An explanation of the basis for the decision and a route to human review
- **B.** A notice that the decision involved an automated system in some way
- **C.** The confidence score the model produced when it made the decision
- **D.** The full system prompt, so the criteria applied are fully transparent

51. Selecciona UNA

D5 · Gobernanza, seguridad y riesgo

A subsidiary streams full patient support chats, names included, to a quality-analytics vendor abroad. The data protection officer objects.

- **A.** Strip or pseudonymise personal data before it leaves the boundary
- **B.** Encrypt the stream in transit to the analytics vendor's own endpoint
- **C.** Shorten retention at the vendor from three years down to one year
- **D.** Add a line to the privacy notice saying that chats may be analysed

52. Selecciona UNA

D5 · Gobernanza, seguridad y riesgo

A scholarship-scoring assistant rates applicants from certain school districts lower. Names and demographic fields were already stripped from the input.

- **A.** Proxies still carry the signal: evaluate by segment and remediate
- **B.** The system is behaving correctly, since no protected field was used
- **C.** The base model is at fault and the application layer cannot help it
- **D.** It is a presentation problem: stop showing the score to reviewers

53. Selecciona UNA

D5 · Gobernanza, seguridad y riesgo

A security review finds that a sentence inside an indexed supplier document changed how the agent behaved.

- **A.** Retrieved text is untrusted data and holds no instruction authority
- **B.** Indexes should only contain documents written inside the company
- **C.** Retrieval must stop whenever a document contains an imperative
- **D.** System prompts must be encrypted wherever they are stored at rest

54. Selecciona 2

D6 · Comunicación con stakeholders y ciclo de vida

A discovery workshop is being planned with a client who has already chosen a solution. Which TWO outcomes make the workshop worth running?

- **A.** A written statement of the business problem behind the requested tool
- **B.** Agreement on how success will be measured and at what threshold
- **C.** A signed fixed-price quote for the solution the client asked for
- **D.** A wireframe of the interface the client pictured beforehand
- **E.** A shortlist of the model versions the team intends to benchmark

55. Selecciona UNA

D6 · Comunicación con stakeholders y ciclo de vida

A client asks for a firm launch date before the evaluation dataset exists. What does the architect commit to?

- **A.** A date for the evaluation to be ready, and a launch decision gated on it
- **B.** The requested launch date, with the evaluation work compressed to fit around it
- **C.** No date at all until every acceptance criterion has been measured
- **D.** A launch date for a reduced scope chosen by the delivery team alone

56. Selecciona UNA

D6 · Comunicación con stakeholders y ciclo de vida

A client's risk committee asks to see exactly how the assistant reaches its answers. What does the architect provide?

- **A.** The decision flow, the guardrails and the evaluation results, in their language
- **B.** The system prompt and the model parameters, as the literal answer
- **C.** A statement that large language models are not fully explainable
- **D.** A live demonstration, which shows the behaviour better than documents

57. Selecciona UNA

D6 · Comunicación con stakeholders y ciclo de vida

A pilot exceeded its targets and the sponsor wants to roll it out to the whole organisation next month. What does the architect raise first?

- **A.** What changes at production scale: load, edge cases, support and cost
- **B.** That the pilot results will not hold, so expectations must come down
- **C.** That a second pilot is needed before committing to any rollout
- **D.** The additional budget the wider rollout is going to require

58. Selecciona UNA

D6 · Comunicación con stakeholders y ciclo de vida

During design, a promising approach was evaluated and discarded. Should that appear in the documentation?

- **A.** Yes: recording why it was rejected stops the team revisiting it blindly
- **B.** No: documentation should describe the system as it was finally built
- **C.** Only if the client asks specifically about that particular alternative
- **D.** Only when the rejected option was the more expensive of the two

59. Selecciona UNA

D6 · Comunicación con stakeholders y ciclo de vida

A client wants an SLA that covers answer accuracy the same way it covers uptime. How should the architect respond?

- **A.** Commit to accuracy measured on an agreed test set, reviewed periodically
- **B.** Commit to uptime only, since accuracy cannot be guaranteed by contract
- **C.** Accept the accuracy target and manage it through internal monitoring
- **D.** Propose a penalty regime, which is what an SLA really comes down to

60. Selecciona UNA

D6 · Comunicación con stakeholders y ciclo de vida

A finance director orders the cheapest model for a filing that feeds the regulator. Testing shows it errs materially more often on that task.

- **A.** Show the evidence and let the accountable owner choose informed
- **B.** Use the better model quietly and absorb the difference elsewhere
- **C.** Comply without comment, since the budget belongs to that director
- **D.** Take it over their head to the audit committee for a final ruling

61. Selecciona UNA

D6 · Comunicación con stakeholders y ciclo de vida

A client demands answers in under a second from a pipeline that retrieves, reasons and verifies. Engineering measures three seconds.

- **A.** Show the latency budget, agree a realistic SLA and add streaming
- **B.** Accept the target now and count on later optimisation to close it
- **C.** Remove the retrieval and verification stages so the target is met
- **D.** Promise sub-second timing for demonstrations to the client only

62. Selecciona UNA

D6 · Comunicación con stakeholders y ciclo de vida

Three months after launch the sponsor is losing interest: the monthly report shows tokens, uptime and average latency.

- **A.** Lead with the criteria agreed at discovery, technical data behind
- **B.** Report weekly instead of monthly so the project stays more visible
- **C.** Add further technical indicators so the report looks more complete
- **D.** Drop the report and rely on conversations with the sponsor instead

63. Selecciona UNA

D7 · Productividad del desarrollador

An AI coding assistant needs access to a service that requires an API key. How is that handled?

- **A.** An environment variable the tooling reads, never a file under version control
- **B.** A configuration file kept out of the repository through the ignore list
- **C.** A note in the project documentation explaining where to find the key
- **D.** A key with a short expiry, pasted into the session when it is needed