

Índice

CCAR-F · Examen de práctica

1

CCAR-F · Examen de práctica

60 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 · Claude Code: configuración y workflows

A platform team is enabling AI-assisted development across twelve squads. Which TWO measures scale the practice without scaling the risk?

- **A.** Keep shared configuration and prompt standards under version control
- **B.** Define which commands automated tooling may run and on which branches
- **C.** Grant the tooling broad credentials so squads are never blocked
- **D.** Let each squad publish its own standards independently
- **E.** Adopt every new assistant tool in the week it becomes available

2. Selecciona UNA

D2 · Claude Code: configuración y workflows

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

D5 · Manejo de contexto y fiabilidad

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. Clasifica cada fila

D2 · Claude Code: configuración y workflows

For each situation, identify the practice that applies.

#	Escenario	Tu respuesta
1	Two squads produce different code style from the same assistant.	
2	AI-written changes pass review, yet defects keep reaching production.	
3	An agent begins calling an unexpected tool on a subset of requests.	
4	An unattended pipeline job could run a destructive command.	
5	A new hire should inherit the team's prompt standards on day one.	

Opciones: shared configuration in version control · review from an independent context · inspect the traces first · restrict what automation may execute

5. Selecciona 2

D1 · Arquitectura agéntica y orquestación

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 · Arquitectura agéntica y orquestación

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 · Arquitectura agéntica y orquestación

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 · Arquitectura agéntica y orquestación

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

9. Clasifica cada fila

D1 · Arquitectura agéntica y orquestación

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

10. Selecciona UNA

D1 · Arquitectura agéntica y orquestación

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

11. Selecciona UNA

D1 · Arquitectura agéntica y orquestación

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

12. Selecciona UNA

D1 · Arquitectura agéntica y orquestación

In a four-stage chain, the third stage fails on about one request in twenty. What belongs in the design?

- **A.** A defined behaviour for that failure: retry, fall back, or stop and report
- **B.** A fifth stage that reviews the output and repairs whatever looks wrong
- **C.** A larger model in stage three so the failures become less frequent
- **D.** A rule telling the model to skip stage three when it cannot complete

13. Selecciona UNA

D3 · Prompt engineering y salida estructurada

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

14. Clasifica cada fila

D3 · Prompt engineering y salida estructurada

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

15. Selecciona UNA

D3 · Prompt engineering y salida estructurada

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

16. Selecciona UNA

D5 · Manejo de contexto y fiabilidad

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

17. Selecciona UNA

D4 · Diseño de herramientas e integración MCP

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

18. Selecciona UNA

D4 · Diseño de herramientas e integración MCP

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

19. Selecciona UNA

D5 · Manejo de contexto y fiabilidad

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

20. Selecciona UNA

D4 · Diseño de herramientas e integración MCP

An internal API exposes forty endpoints. The team is deciding how to expose it to the agent. What should drive the tool boundaries?

- **A.** The tasks the agent needs to accomplish, not the endpoints that exist
- **B.** One tool per endpoint, so the whole surface stays available to it
- **C.** A single tool that receives the endpoint name and its parameters
- **D.** The endpoints the backend team considers most stable over time

21. Selecciona UNA

D5 · Manejo de contexto y fiabilidad

An agent platform needs alerting. Which signal deserves a page to the on-call engineer?

- **A.** Tool error rate crossing a threshold sustained over several minutes
- **B.** Any single request that took longer than the median to complete
- **C.** Daily token spend rising above the same day of the previous week
- **D.** A user rating one individual answer as unhelpful in the interface

22. Selecciona UNA

D1 · Arquitectura agéntica y orquestación

A report needs four independent sections researched and then written as one document. How should the work be split across agents?

- **A.** Four subagents research in parallel, one agent writes the document
- **B.** Four subagents each research and write their own section
- **C.** A single agent researches and then writes everything in one long pass
- **D.** Four subagents write, and a fifth merges their four drafts

23. Selecciona UNA

D1 · Arquitectura agéntica y orquestación

An orchestrator dispatches subagents whose number is only known once the problem is read. Which pattern is this?

- **A.** Orchestrator-workers
- **B.** Routing across a fixed menu
- **C.** Parallelization over fixed sections
- **D.** Prompt chaining in a fixed sequence

24. Selecciona 2

D1 · Arquitectura agéntica y orquestación

A subagent receives its assignment from an orchestrator. Which TWO parts of that assignment should a person define in advance, not the orchestrator at runtime?

- **A.** The output format every subagent must return

- **B.** The hard limits and the tools it may use
- **C.** The specific question this subagent must answer
- **D.** The portion of the problem it has been given
- **E.** The order in which it should consult its sources

25. Selecciona 2

D1 · Arquitectura agéntica y orquestación

A team is deciding whether a multi-agent design is worth its overhead. Which TWO conditions make it pay off?

- **A.** The subtasks are independent and do not need to talk to each other
- **B.** Each subtask produces bulk material that is better summarised apart
- **C.** Several subtasks must run in a strict order, each using the last result
- **D.** Two agents need to read and update the same shared state as they go
- **E.** The team wants the architecture to look more capable to the client

26. Selecciona 2

D1 · Arquitectura agéntica y orquestación

Results from parallel subagents come back in incompatible shapes and cannot be combined. Which TWO fixes address the cause?

- **A.** Define one output contract that every subagent must fill in
- **B.** Validate each subagent's return against that contract before use
- **C.** Ask the orchestrator to reformat whatever each subagent sent
- **D.** Increase each subagent's context so it can see all the others' work
- **E.** Run the subagents sequentially so later ones copy the first

27. Selecciona 2

D2 · Claude Code: configuración y workflows

A PreToolUse hook exits with code 2 after a blocked command. Which TWO things happen next?

- **A.** The tool call never executes
- **B.** What the hook wrote to stderr goes back to the model
- **C.** The session ends and has to be restarted by the user
- **D.** PostToolUse still runs so the block can be recorded
- **E.** The model retries the same command automatically

28. Selecciona 2

D2 · Claude Code: configuración y workflows

A monorepo holds twelve applications, each with its own conventions. Which TWO practices keep context usage sane?

- **A.** Put app-specific conventions in a CLAUDE.md inside each app folder
- **B.** Keep only shared, repo-wide conventions in the root CLAUDE.md
- **C.** Concatenate every app's conventions into the single root file for clarity
- **D.** Move all conventions into settings.json so they load on demand
- **E.** Ask developers to paste the relevant conventions at session start

29. Selecciona UNA

D3 · Prompt engineering y salida estructurada

A downstream service parses the model's output as JSON and crashes when a field is missing. What is the correct fix?

- **A.** Enforce the shape with a schema instead of asking for it in the prompt
- **B.** Add a sentence to the prompt insisting the JSON must be complete
- **C.** Wrap the parser in a try block and skip records that fail
- **D.** Post-process the text with a regular expression to patch the gaps

30. Selecciona UNA

D3 · Prompt engineering y salida estructurada

A prompt mixes instructions, a reference document and a user question, and the model keeps confusing them. What is the most direct improvement?

- **A.** Delimit each part with explicit tags so the boundaries are unambiguous
- **B.** Repeat the instructions again after the reference document ends
- **C.** Shorten the reference document until the confusion disappears
- **D.** Raise the temperature so that the model interprets the layout more loosely

31. Selecciona UNA

D3 · Prompt engineering y salida estructurada

An output will be read by a person, not parsed by a program. What follows for its format?

- **A.** Prose is fine; a strict schema adds cost without buying anything
- **B.** A schema should still be enforced for consistency across responses
- **C.** The output should be JSON so it can be reused programmatically later
- **D.** The format should change per request to suit the reader's preference

32. Selecciona UNA

D3 · Prompt engineering y salida estructurada

Few-shot examples were added to a classification prompt and accuracy on one rare category got worse. What is the likely cause?

- **A.** The examples over-represent the common categories and skew the model
- **B.** The examples are too short for the model to learn the whole pattern from
- **C.** Few-shot prompting does not work for classification tasks at all
- **D.** The temperature is too low for the examples to have any influence

33. Selecciona UNA

D3 · Prompt engineering y salida estructurada

An extraction step must return a consistent output shape. Prose instructions describing the format keep producing variations. What works better?

- **A.** Two to four worked input and output examples in the prompt
- **B.** A longer written description of the format
- **C.** A reminder at the end of the prompt to respect the format
- **D.** A lower temperature so the output varies less between runs

34. Selecciona UNA

D3 · Prompt engineering y salida estructurada

A review prompt says 'be conservative and report only high-confidence findings'. Results are inconsistent between runs. What is the defect?

- **A.** The criterion is vague, so each run interprets it differently
- **B.** The prompt is too short for the model to apply judgement
- **C.** The model needs a larger context to be consistent about it
- **D.** Confidence cannot really be expressed in a prompt in any useful way

35. Selecciona UNA

D4 · Diseño de herramientas e integración MCP

A tool wraps an internal API with eleven optional parameters. The agent fills them inconsistently. What is the better design?

- **A.** Expose a task-oriented tool that takes only what the task needs
- **B.** Document all eleven parameters more thoroughly in the description
- **C.** Add validation that rejects calls with unusual parameter combinations
- **D.** Let the agent send free-form text and parse the parameters out of it

36. Selecciona UNA

D4 · Diseño de herramientas e integración MCP

A tool returns an empty list when the underlying service times out. What is wrong with that?

- **A.** The model reads it as 'no results' and answers confidently from nothing
- **B.** The empty list uses more tokens than an explicit error would use
- **C.** Timeouts should always be retried inside the tool without informing anyone
- **D.** Empty lists are valid only when the query itself was malformed

37. Selecciona UNA

D4 · Diseño de herramientas e integración MCP

Two agents from different companies must exchange work without either exposing internal tools. Which mechanism fits?

- **A.** An agent-to-agent protocol, each side mediating its own systems
- **B.** An MCP server published by one of the companies for the other to use
- **C.** A shared tool catalogue registered in both agents at once
- **D.** A direct API integration between the two internal systems

38. Selecciona UNA

D4 · Diseño de herramientas e integración MCP

Two tools share the description 'Retrieves information' and the agent keeps picking the wrong one. What fixes it?

- **A.** Rename each tool and state the concrete data each one returns
- **B.** Add an instruction telling the model to think before choosing
- **C.** Merge both tools into one that takes a mode parameter
- **D.** Remove one of the two tools from the agent's catalogue

39. Selecciona UNA

D4 · Diseño de herramientas e integración MCP

A project's MCP server needs an API token, and the configuration file is committed to the repository. How is that handled?

- **A.** Reference an environment variable from the configuration file
- **B.** Commit the token and rotate it whenever someone leaves
- **C.** Keep the whole configuration out of version control instead
- **D.** Store the token in the project CLAUDE.md so it is documented

40. Selecciona UNA

D4 · Diseño de herramientas e integración MCP

A developer wants to try an MCP server on their own, without affecting teammates. Where does that configuration go?

- **A.** In the user-level configuration in their home directory
- **B.** In the project configuration, disabled by default
- **C.** In a branch of the repository that is never merged
- **D.** In the project configuration, with a comment explaining it

41. Selecciona UNA

D4 · Diseño de herramientas e integración MCP

An agent burns several tool calls each session just discovering what data is available. What reduces that?

- **A.** Expose a resource that maps the available data up front
- **B.** Increase the agent's context so it can remember prior sessions
- **C.** Add a tool that lists every other tool and what it does
- **D.** Instruct the agent to guess the structure before querying

42. Selecciona UNA

D5 · Manejo de contexto y fiabilidad

A long agent session degrades as it goes: earlier facts get dropped and answers drift. What is the sound response?

- **A.** Re-inject the facts that must survive, instead of trusting the history
- **B.** Increase the temperature so the model explores the history more
- **C.** Ask the model to summarise the conversation and continue
- **D.** Start over from scratch each time the drift becomes noticeable

43. Selecciona UNA

D5 · Manejo de contexto y fiabilidad

A retrieval step returns nothing relevant for a user question. What should the system do?

- **A.** Say it has no grounded answer and offer the next step
- **B.** Answer from the model's own knowledge without citing anything
- **C.** Return the closest passages anyway and let the user judge
- **D.** Retry the same query until something relevant comes back

44. Selecciona UNA

D5 · Manejo de contexto y fiabilidad

In a long support conversation, the order number and the amount keep getting lost after the history is summarised. What fixes it?

- **A.** Keep a case-facts block outside the history and resend it every turn
- **B.** Summarise less aggressively so more of the history survives
- **C.** Ask the model to repeat the key facts at the end of each reply
- **D.** Raise the context window so summarisation happens later

45. Selecciona UNA

D5 · Manejo de contexto y fiabilidad

A long input is processed reliably at the beginning and the end, but details in the middle are missed. What is the mitigation?

- **A.** Place the decisive information near the start or the end
- **B.** Repeat the whole input twice so the middle appears elsewhere
- **C.** Increase the temperature so more of the input is considered
- **D.** Split the input into random chunks and process them apart

46. Selecciona UNA

D5 · Manejo de contexto y fiabilidad

A long investigation will span several sessions. How are the findings preserved?

- **A.** Write key findings to a file and reference it in the next session
- **B.** Keep the session open indefinitely so nothing is ever lost
- **C.** Rely on compaction to retain the important parts of the history
- **D.** Repeat the discovery steps at the start of each new session

47. Selecciona 2

D1 · Arquitectura agéntica y orquestación

A coordinator spawns subagents with the Task tool. Which TWO statements about their context are correct?

- **A.** A subagent does not inherit the coordinator's conversation history
- **B.** Everything the subagent needs must be stated in its own prompt
- **C.** Subagents share a memory space that persists across separate Task calls
- **D.** A subagent can read the coordinator's context window on demand
- **E.** Subagents inherit the coordinator's allowed tools automatically

48. Selecciona 2

D1 · Arquitectura agéntica y orquestación

Which TWO situations call for starting a fresh session rather than continuing the current one?

- **A.** The tool results in context are stale and no longer reflect reality
- **B.** A long stretch of unrelated work has degraded the working context
- **C.** The next task needs the same repository and the same conventions
- **D.** A teammate wants to explore an alternative branch of that same task later
- **E.** The conversation is long but every earlier fact is still accurate

49. Clasifica cada fila

D1 · Arquitectura agéntica y orquestación

For each stop_reason value, identify what the loop should do next.

#	Escenario	Tu respuesta
1	tool_use	
2	end_turn	
3	max_tokens	
4	stop_sequence	
5	tool_use returned after appending previous results	

Opciones: execute the requested tools and continue · show the result to the user · raise the output budget and retry · treat the custom stop sequence as the boundary

50. Selecciona 2

D2 · Claude Code: configuración y workflows

Which TWO statements about hook security are accurate?

- **A.** A hook runs locally with the user's own permissions, without asking
- **B.** A hook from an untrusted repository is arbitrary code execution
- **C.** Hooks run inside a sandbox that isolates them from the whole file system
- **D.** Hooks require explicit confirmation from the user on every trigger
- **E.** Hooks can only execute commands that appear in the allow list

51. Selecciona 2

D2 · Claude Code: configuración y workflows

Which TWO facts about permissions in settings.json are correct?

- **A.** A deny rule beats an allow rule from any other configuration layer
- **B.** Rules are written per tool and can use patterns to narrow scope
- **C.** The model reads the permission list and explains its choices from it
- **D.** Allow lists are advisory and the model may override them if needed
- **E.** Permissions are inherited only when the project has no local file

52. Clasifica cada fila

D2 · Claude Code: configuración y workflows

For each need, identify the Claude Code mechanism that fits.

#	Escenario	Tu respuesta
1	The team's preferred naming conventions, which the model should follow.	
2	A guarantee that no destructive migration command ever executes.	

#	Escenario	Tu respuesta
3	Declaring which Bash commands are allowed without a prompt.	
4	A specialised procedure the model should pull in when it applies.	
5	Blocking a tool call and telling the model why it was refused.	

Opciones: CLAUDE.md · hook · settings.json · skill or subagent

53. Selecciona 2

D2 · Claude Code: configuración y workflows

A CLAUDE.md pulls in other files with the @path import syntax. Which TWO statements are correct?

- **A.** Relative paths resolve against the file doing the importing
- **B.** Imports can nest, but only up to a maximum depth of five
- **C.** Imported files load only when the model explicitly asks for them
- **D.** Imports accept a space between the @ and the path for readability
- **E.** Only Markdown files can be imported, never configuration files

54. Selecciona 2

D2 · Claude Code: configuración y workflows

A team moves its conventions from one long CLAUDE.md into .claude/rules/ files with paths frontmatter. Which TWO gains does that produce?

- **A.** Each rule loads only when files matching its patterns are touched
- **B.** Context and tokens are spent only on the conventions in play
- **C.** Rules become enforceable, so the model can no longer ignore any of them
- **D.** All conventions load at session start, which speeds up responses
- **E.** The rules apply to every repository the developer works in

55. Selecciona 2

D2 · Claude Code: configuración y workflows

Which TWO statements about the built-in /compact and /memory commands are correct?

- **A.** Compacting can lose numeric values and dates in the summary
- **B.** What /memory saves reloads automatically in later sessions
- **C.** Compacting preserves the full transcript and only hides it
- **D.** Memory entries apply only to the session where they were written
- **E.** Compacting is required before any session can exceed its window

56. Clasifica cada fila

D2 · Claude Code: configuración y workflows

For each need, identify where the definition belongs.

#	Escenario	Tu respuesta
1	A review checklist the whole team should share.	
2	A personal shortcut one developer uses across every project.	
3	A guarantee that a destructive command never runs.	
4	A noisy investigation routine that should not fill the session.	
5	MCP servers every contributor to this repository needs.	

Opciones: project-level, committed to the repository · user-level, in the home directory
· a hook registered on a tool event · a skill with context: fork

57. Selecciona UNA

D3 · Prompt engineering y salida estructurada

An extraction schema marks every field as required. Documents that lack a field come back with invented values. What is the fix?

- **A.** Make fields that may be absent optional or explicitly nullable
- **B.** Add an instruction telling the model never to invent values
- **C.** Reject any extraction where a required field looks implausible
- **D.** Split the schema so each field is extracted in its own call

58. Selecciona UNA

D3 · Prompt engineering y salida estructurada

A classification enum keeps forcing borderline documents into the closest available category. What design change helps?

- **A.** Add an unclear value, plus an other option with a detail field
- **B.** Add more categories until every observed document fits one
- **C.** Instruct the model to pick the category it considers most likely
- **D.** Raise the temperature so the classification is less rigid

59. Selecciona UNA

D3 · Prompt engineering y salida estructurada

Invoice totals occasionally disagree with the sum of their line items, and nothing downstream notices. Which extraction design surfaces it?

- **A.** Extract the stated total and the calculated total, and flag conflicts
- **B.** Extract only the calculated total, since it is the trustworthy one
- **C.** Extract only the stated total, since it is what the document says
- **D.** Round both values so small differences stop being reported

60. Clasifica cada fila

D4 · Diseño de herramientas e integración MCP

For each need, identify the built-in tool that fits.

#	Escenario	Tu respuesta
1	Find every test file matching a naming pattern across the repository.	
2	Locate which files mention a specific error message.	
3	Load a file in full to trace how a function flows.	
4	Change one unique line inside an existing file.	
5	List all TypeScript files under a components directory.	

Opciones: Glob · Grep · Read · Edit