Anthropic

(Claude Certified Developer Foundations)

Claude Certified Developer Foundations (CCDV-F)

Total: **101 Questions**

Link:

MYEXAM.FR

Prompt instructions alone do not guarantee reliable structured output. Claude may still add prose, change fields, or produce invalid JSON.

C. Run Claude Code with --output-format json and --json-schema, then map validated findings to inline PR comments.

Correct

Using **structured JSON output with a defined schema** provides predictable field names and a machine-readable format. The CI workflow can validate the findings and reliably map each result to the appropriate inline pull request comment.

D. Store previous comments and ask Claude to rewrite them into records.

Incorrect

Reprocessing previous comments introduces unnecessary complexity and still does not guarantee a stable, validated structure for automation.

Final Answer: C

C = Use --output-format json with --json-schema and map the validated findings to inline PR comments — this provides reliable, structured output for CI automation.

Question: 5

Scenario: Multi-Agent Research System
The document analysis subagent receives coordinator-selected PDFs and reports from an approved catalog, but it has a generic URL retrieval tool.
It sometimes follows links inside documents to:
l blogs
l login pages
l duplicate HTML summaries and then cites those pages instead of the approved sources.
What change best addresses this?

A. Keep fetch_url available, but add prompt instructions warning the subagent never to open links found inside documents.

B. Replace fetch_url with a load_document tool that accepts catalog document IDs or approved URLs and validates before fetching.

C. Give the document analysis subagent web search tools too, so it can independently confirm whether linked pages are relevant.

D. Allow fetch_url for any link, then have synthesis discard citations whose domains are not in the approved catalog.

Answer: B

Explanation:

A. Keep fetch_url available, but add prompt instructions warning the subagent never to open links found inside documents.

Incorrect

Prompt instructions can help guide behavior, but they do not technically prevent the subagent from accessing unapproved URLs. A tool-level restriction is more reliable.

Adding a structured **detected_pattern** field makes findings easier to categorize and analyze. The team can determine which patterns — such as missing renewal clauses or informal pricing language — have higher or lower acceptance rates and then improve the system based on evidence.

C. Store only analyst decisions.

Incorrect

Storing only acceptance/rejection decisions loses important information about what pattern generated the finding, making it difficult to identify why certain findings are rejected.

D. Tell Claude to report only findings analysts accept.

Incorrect

This would hide potentially useful findings rather than improving the detection process. It also prevents the team from analyzing rejected findings and understanding systematic errors.

Final Answer: B

B = Add a structured detected_pattern field and analyze acceptance rates by pattern and document type — this provides the data needed to identify systematic QA errors and improve extraction quality.

Question: 8

A multi-agent system uses a coordinator and specialized researchers. How should required MCP server configuration be shared across a team repository?

A. Ask every user to recreate servers manually
B. Store all servers in personal user configuration
C. Combine all integrations into one MCP server
D. Place required shared servers in project-scoped .mcp.json

Answer: D

Explanation:

A. Ask every user to recreate servers manually

Incorrect

Manual recreation is inconsistent and creates unnecessary setup work. Team members may configure servers differently, leading to inconsistent agent behavior.

B. Store all servers in personal user configuration

Incorrect

Personal configuration is tied to an individual user and is not appropriate for sharing required project-level configuration with the entire team.

C. Combine all integrations into one MCP server

Incorrect

Combining unrelated integrations into a single server reduces modularity and is unnecessary. MCP servers

What should happen before modifying files?

A. Implement a simple cache wrapper and iterate later.
B. Ask targeted questions about invariants, invalidation, audit requirements, and failure modes.
C. Provide performance goals and let Claude infer rules.
D. Copy caching patterns from another module.

Answer: B

Explanation:

A. Implement a simple cache wrapper and iterate later.

Incorrect

Implementing immediately could introduce serious issues in authorization, tenant isolation, audit logging, or cache invalidation. These requirements should be clarified first.

B. Ask targeted questions about invariants, invalidation, audit requirements, and failure modes.

Correct

Before modifying a complex legacy system, Claude should clarify the **security and operational constraints**. Targeted questions help identify authorization invariants, tenant boundaries, invalidation rules, audit requirements, and failure scenarios before implementation begins.

C. Provide performance goals and let Claude infer rules.

Incorrect

Performance goals alone do not define critical authorization and security requirements. Letting Claude infer these rules could result in unsafe assumptions.

D. Copy caching patterns from another module.

Incorrect

Another module may have different authorization, tenant, or invalidation requirements. Reusing its pattern without validating the assumptions could introduce defects.

Final Answer: B

B = Ask targeted questions about invariants, invalidation, audit requirements, and failure modes — clarify critical constraints before making changes to a complex legacy authorization system.

Question: 13

A developer creates a multi-agent workflow where one agent researches documentation and another writes implementation plans. How should information flow between agents be handled?

A. Assume all agents share the same conversation context
B. Pass required findings through the coordinator
C. Allow each agent to access every previous message
D. Remove the coordinator role

Incorrect

This relies on the report agent correctly recognizing incomplete input. It is better to enforce the dependency before the report-generation task starts.

C. Check whether synthesis messages contain words like "complete."

Incorrect

Keyword checks are unreliable. A message can say "complete" while still missing required citations or evidence.

D. Block report-generation tasks until synthesis returns a structured citation manifest and success status.

Correct

This creates an explicit **workflow gate**. The report cannot start until the synthesis agent provides a structured citation manifest and confirms successful completion. This ensures the report is based on complete and verifiable evidence.

Final Answer: D

D = Block report generation until the synthesis agent returns a structured citation manifest and success status, preventing incomplete or unverifiable reports from being produced.

Question: 16

A migration assistant sometimes starts dependency analysis before collecting required package manager and framework versions.
What change most reliably enforces the correct first step?

A. Detect whether Claude mentions metadata completion in natural language.
B. Force the metadata extraction tool with tool_choice for the first request.
C. Set tool_choice to any for the first request.
D. Strengthen the system prompt requiring metadata extraction first.

Answer: B

Explanation:

A. Detect whether Claude mentions metadata completion in natural language.

Incorrect

Natural-language detection is unreliable. Claude may mention that metadata is complete without actually calling the required extraction tool or obtaining all necessary information.

B. Force the metadata extraction tool with tool_choice for the first request.

Correct

Using tool_choice explicitly forces the required metadata extraction tool to run first. This provides a deterministic workflow and prevents dependency analysis from starting before the required package manager and framework versions are collected.

C. Set tool_choice to any for the first request.

C. Map failures to customer sentiment levels.

D. Return structured MCP errors with categories, retryability, and user-safe details.

Answer: D

Explanation:

A. Retry every failed tool call twice before escalation.

Incorrect

Not every failure is retryable. Retrying invalid order IDs or policy restrictions wastes resources and can make failures worse.

B. Keep generic failures and let Claude infer retryability.

Incorrect

A generic error message does not provide enough reliable information for the agent to determine whether a failure should be retried, corrected, or escalated.

C. Map failures to customer sentiment levels.

Incorrect

Customer sentiment does not determine the technical recovery strategy. The agent needs actionable information about the actual failure.

D. Return structured MCP errors with categories, retryability, and user-safe details.

Correct

Structured MCP errors give the agent **explicit recovery information**. For example, a timeout can be marked retryable, while an invalid order ID or policy restriction can be marked non-retryable. User-safe details also allow the agent to communicate an appropriate explanation without exposing sensitive backend information.

Final Answer: D

D = Return structured MCP errors containing failure categories, retryability, and user-safe details so the agent can choose the appropriate recovery action.

Question: 21

Scenario: Multi-Agent Research System
You are building a multi-agent research system using the Claude Agent SDK. A coordinator agent delegates to specialized subagents: one searches the web, one analyzes documents, one synthesizes findings, and one generates reports.
Your coordinator's document-analysis subagent depends on an internal library-catalog MCP server. It works in your local Claude Code session, but two teammates who cloned the repository cannot access the catalog tools, so their runs fall back to weaker web searches. Separately, one researcher is testing an experimental notes server that should not affect the team.
What should you change?

A. Document per-user claude mcp add commands in README, and ask each researcher to recreate both servers locally.

B. Place both servers in each user's ~/.claude.json, and rely on onboarding scripts to synchronize changes across machines.

A. Read every repository file before implementing changes.
B. Use plan mode to map relevant flows and compare implementation points before modifications.
C. Directly edit and rely on test failures to guide corrections.
D. Start separate direct sessions for each module and combine changes later.

Answer: B

Explanation:

A. Read every repository file before implementing changes.

Incorrect

Reading the entire repository is inefficient and can overwhelm the analysis with irrelevant information. The goal is to identify the relevant flows and implementation points first.

B. Use plan mode to map relevant flows and compare implementation points before modifications.

Correct

Plan mode allows Claude to **investigate the codebase, trace the cancellation behavior, identify related components, and compare possible implementation points** before making changes. This is especially useful when behavior is distributed across multiple abstractions.

C. Directly edit and rely on test failures to guide corrections.

Incorrect

This approach can lead to repeated incorrect edits and unnecessary regressions. Understanding the architecture first reduces the chance of modifying the wrong layer.

D. Start separate direct sessions for each module and combine changes later.

Incorrect

Separate sessions can lack shared context and may produce inconsistent approaches across modules. A coordinated plan provides a clearer understanding of the complete workflow.

Final Answer: B

B = Use plan mode to map the relevant flows and compare implementation points before making modifications, reducing the risk of editing the wrong abstraction.

Question: 25

A team builds an application where Claude must classify customer requests into predefined categories. Which approach provides the most consistent results?

A. Provide examples of each category and expected classifications
B. Ask Claude to create new categories when needed
C. Use a random classification approach.
D. Remove category definitions from the prompt

Answer: A

Answer: D

Explanation:

A. Repeat broad assignments twice.

Incorrect

Repeating the same broad assignment does not guarantee that the missing Asia-Pacific coverage will be identified or addressed. It may simply produce similar results.

B. Hide unsupported sections.

Incorrect

Hiding unsupported sections reduces transparency and does not solve the underlying coverage gap.

C. Increase coordinator iteration count to a fixed number.

Incorrect

A fixed number of iterations does not ensure that important gaps have been resolved. The system needs to evaluate **what is missing**, not simply perform more iterations.

D. Add a coordinator review loop that checks coverage, delegates follow-ups, and reruns synthesis.

Correct

A coordinator review loop can explicitly **evaluate coverage against the research requirements**, identify gaps such as weak Asia-Pacific coverage, delegate targeted follow-up research, and then rerun synthesis with the new findings.

This creates an adaptive feedback loop rather than relying on repeated broad searches.

Final Answer: D

D = Add a coordinator review loop that checks coverage, delegates targeted follow-ups, and reruns synthesis to close identified research gaps.

Question: 28

A developer wants Claude to review code and return only actionable issues. What prompt design is most effective?

A. Request every possible improvement
B. Define what qualifies as an issue and provide examples.
C. Ask for unlimited suggestions
D. Avoid describing expected output

Answer: B

Explanation:

A. Request every possible improvement

Incorrect

This can produce excessive suggestions, including minor style preferences or changes that are not actually

B. Randomly review a fixed percentage of every extraction category, regardless of document type, field, or prior validation performance.

Incorrect

Uniform sampling ignores known differences in performance. Higher-risk document families and fields may require more focused validation rather than the same review percentage everywhere.

C. Raise the global confidence threshold until the aggregate validation accuracy reaches 99%, then disable review for remaining outputs.

Incorrect

Increasing a global confidence threshold does not address **systematic errors concentrated in specific fields or document types**. High confidence can still be poorly calibrated for certain segments.

D. Use the 97.2% aggregate accuracy to auto-approve all high-confidence outputs, monitoring user complaints for emerging regressions.

Incorrect

This relies on aggregate performance and reactive customer feedback. It does not adequately validate whether high-confidence outputs are reliable within the specific segments where errors are already concentrated.

Final Answer: A

A = Analyze accuracy by document type and field, then automate only validated high-performing segments while keeping higher-risk segments under human review.