
Exam : CCAO-F

**Title : Claude Certified Associate -
Foundations**

Version : DEMO

1. You are a knowledge worker evaluating which Claude feature matches each proposed use case. Which two feature-to-use-case matches are appropriate? Each correct answer presents a complete solution. (Select two.)

- A. Long-context summarization for a quarterly report distilled from many sources.
- B. Long-context summarization for drafting a reply to a single short customer email.
- C. Web search invocation to retrieve the established formula for calculating compound interest.
- D. Tool use to invoke a calculator for rephrasing a product description in simpler language.
- E. Tool use to invoke a calculator for a complex, multi-step pricing calculation.

Answer: A E

Explanation:

Option A appropriately uses Claude's long-context capabilities for a task involving many source documents. A quarterly report may require Claude to compare evidence across financial results, operational updates, customer information, and prior reporting periods. Supplying the relevant sources in a well-organized context allows Claude to synthesize the material while following a consistent report structure. Source boundaries and attribution should remain clear so the final output can be validated.

Option E is also an appropriate match. A calculator or code-execution tool is suitable for complex, multi-step pricing calculations because it provides deterministic computation that can be inspected and tested. Claude can structure the calculation and interpret the result, while the tool performs the arithmetic more reliably.

Option B does not require long-context processing because a single short email can normally be handled directly in chat.

Option C unnecessarily invokes web search for a stable, established formula that does not require current external information.

Option D mismatches the tool to the task: simplifying product language is a writing and audience-adaptation task, not a numerical calculation.

Effective feature selection depends on the actual workload. Large source collections benefit from long-context synthesis, whereas precise arithmetic benefits from deterministic computational tools and explicit result validation.

2. You are a project manager reviewing Claude usage approaches that teammates have proposed for a project workflow.

Which two workflow steps are well-suited for Claude support? Each correct answer presents a complete solution. (Select two.)

- A. approving a final external press release without any human review
- B. issuing legally binding contractual commitments on behalf of the company
- C. deciding promotion outcomes for staff based solely on the model's judgment
- D. synthesizing a literature review across many uploaded research documents
- E. drafting a structured project status update from a set of meeting notes

Answer: D E

Explanation:

Under the objective Identify appropriate and inappropriate use cases, determining suitable application domains hinges on evaluating AI capabilities against risk severity and the presence of human oversight. Claude excels at information extraction, multi-document synthesis, and structured text generation. Synthesizing complex findings from multiple research papers (Option D) and compiling meeting notes

into organized, structured status summaries (Option E) are classic high-leverage knowledge-work tasks that streamline team operations while keeping final outputs subject to team review. Conversely, allowing an AI model to autonomously authorize external corporate press releases without human sign-off (Option A), bind an organization to contractual liabilities (Option B), or determine career-altering personnel promotions (Option C) represents dangerous and inappropriate delegation that violates responsible AI principles, legal compliance, and human accountability standards.

3. When refining a Claude output that does not meet expectations, which iterative refinement step should be performed first?

- A. Submit the revised prompt and capture the new output.
- B. Compare the new output against the original criteria.
- C. Modify the prompt to address each identified weakness.
- D. Identify the specific weaknesses in the current output.

Answer: D

Explanation:

Answer D is correct. Iterative refinement should begin by diagnosing the current output against the original requirements. The reviewer must identify specific weaknesses such as missing information, unsupported claims, incorrect tone, poor organization, excessive length, or failure to follow the requested format. This diagnosis provides evidence for deciding what must change in the prompt.

After identifying the weaknesses, the reviewer can modify the prompt to address them, which corresponds to Option C. The revised prompt is then submitted and the resulting output captured, as described in Option

A. Finally, the new output is compared with the original criteria to determine whether the revision improved performance, as stated in Option

B. The logical sequence is therefore D, C, A, and B.

Beginning with prompt modification before diagnosing the failure encourages random changes and makes it difficult to determine which revision produced an improvement. Similarly, a new output cannot be submitted or evaluated before a targeted revision has been created. Anthropic's guidance recommends defining clear success criteria, evaluating results against those criteria, and using a structured self-correction process: generate a draft, review it, and then refine it based on identified deficiencies. Effective iteration is therefore diagnostic and evidence-based rather than trial and error.

Anthropic prompting best practices

4. You are reviewing Connector setups that colleagues have proposed for various tasks.

Which two Connector setups represent safe and appropriate uses? Each correct answer presents a complete solution. (Select two.)

- A. a Connector that links to the personnel-records system for team leads to support scheduling tasks
- B. a Connector that links to the team's project management workspace under approved scope
- C. a Connector that links to a personal cloud storage account where team members store working drafts alongside personal files
- D. a read-only Connector that links to the public knowledge base for reference content
- E. a Connector that links to the payroll system to allow team members to submit expense reports directly

Answer: B D

Explanation:

Under the objective Manage uploaded knowledge and connectors (e.g., Google Drive, Gmail), establishing data connectors requires adhering to the principle of least privilege, strict scope boundary management, and robust data classification.

Option B is an authorized, safe integration: connecting Claude to an approved, bounded project management workspace with explicit scope controls allows the model to summarize tasks and sprint updates safely.

Option D is an ideal read-only integration, granting access to public documentation and reference repositories where confidential data leakage risks are zero. Conversely, connecting AI tools to core payroll infrastructure (Option

E) or sensitive personnel/HR databases (Option A) introduces severe compliance, privacy, and unauthorized access liabilities. Linking personal cloud storage repositories containing unvetted mixed personal data (Option

C) violates data governance standards by risking unintentional exposure of private files.

5. You are running an optimization pass on a recurring workflow.

Which is the correct order of optimization steps?

- (1) Apply targeted optimization to the bottleneck step.
- (2) Roll out the optimized workflow to the broader team and document the change.
- (3) Identify the bottleneck step using measured time per step and queue length.
- (4) Map the current workflow steps, owners, and decision points.
- (5) Pilot the change with a small group and measure the effect against the prior baseline.

A. 4,3,1,5,2

B. 4,3,2,1,5

C. 4,1,3,5,2

D. 3,4,1,5,2

Answer: A

Explanation:

The correct sequence is 4,3,1,5,2. Optimization must begin by mapping the current workflow, including its steps, owners, inputs, handoffs, and decision points. Without an accurate current-state map, the team cannot reliably associate delays or defects with a particular stage. It should then use measured evidence—such as processing time, waiting time, queue length, rework, or error frequency—to identify the actual bottleneck.

After locating the constraint, the team can apply a targeted improvement to that step. Changing the workflow before identifying the bottleneck, as in Option C, risks optimizing a noncritical activity while leaving the primary delay untouched. The change should next be piloted with a limited group and compared against the previous baseline. This determines whether the intervention improves throughput or quality without creating downstream problems. Only after a successful pilot should the process be rolled out broadly and formally documented.

This sequence is consistent with Anthropic's evaluation approach: define measurable criteria, test changes against representative conditions, and validate performance before production use.

Option D identifies the bottleneck before establishing the workflow map needed to interpret the measurements. Anthropic's evaluation guidance