

Anthropic CCAO-F

Claude Certified Associate - Foundations



For More Information – Visit link below:

<https://www.examsboost.com/>

Product Version

- ✓ Up to Date products, reliable and verified.
- ✓ Questions and Answers in PDF Format.

Latest Version: 6.0

Question: 1

Claude turns a response into a separate Artifact rather than an inline chat reply mainly when the content is:

- A. Substantial, self-contained, and reusable
- B. A response that simply runs past a few paragraphs of prose
- C. Written in a formal, professional tone
- D. Generated during an active Research session

Answer: A

Explanation:

Claude creates a separate Artifact when content is substantial, self-contained, and likely to be reused or referenced later, not simply because it is long or formal. The trigger is reusability, which is why brief reusable items can become Artifacts too.

Question: 2

A Pro-plan user avoids Haiku 4.5 entirely, believing it is only a lesser model for people who cannot access better ones.

Which statement corrects this?

- A. Haiku 4.5 is offered mainly to Free-plan users as an entry-level option
- B. Haiku 4.5 is best reserved for tasks needing the deepest reasoning
- C. Haiku 4.5 produces the highest-quality reasoning of the four
- D. Haiku 4.5 is the deliberate choice for quick, high-volume tasks

Answer: D

Explanation:

Haiku 4.5 is the fastest, lightest model and the right pick for quick, high-volume, low-ambiguity work. It is a purposeful selection on any plan, including paid tiers, not a lesser model reserved for users who cannot reach better ones.

Question: 3

A business user is unsure which model to pick for typical everyday writing and analysis. According to Anthropic's guidance, which model is the recommended default?

- A. Opus 4.8, reserved for complex high-stakes reasoning

- B. Haiku 4.5, the fastest and lightest model
- C. Sonnet 5, the balanced default for most work
- D. Fable 5, the newest and most capable model

Answer: C

Explanation:

Sonnet 5 balances speed and intelligence and handles most everyday writing, analysis, and multi-step tasks, so it is the model a business user should default to when unsure, rather than reaching first for a heavier or lighter option.

Question: 4

For external-facing or sensitive content, Anthropic recommends disclosing AI involvement in addition to which other safeguard?

- A. Qualified professional review before it reaches consumers
- B. Switching to the account's strongest available model
- C. Attaching citations to the generated draft for source traceability
- D. Turning off the model-improvement training toggle

Answer: A

Explanation:

For external-facing or sensitive content, Anthropic recommends disclosing AI involvement in addition to having a qualified professional review the content before it reaches consumers. Disclosure and human review are companion safeguards, not substitutes, and neither is replaced by model choice, citations, or a privacy toggle.

Question: 5

One section of a long Markdown document in an Artifact needs a tone change for a different audience, while the rest is correct.

What is the most appropriate way to revise it?

- A. regenerate the entire document from scratch with a new prompt
- B. start a new conversation and rebuild the whole document
- C. highlight the section and request a targeted edit of that part
- D. retry the original prompt to generate an alternate full version

Answer: C

Explanation:

When only one section of a Markdown Artifact needs a tone change, highlighting that section and requesting a focused edit is the precise, lower-risk technique, preserving everything already correct rather than regenerating the whole document.

Question: 6

A strategist needs both careful reasoning and heavy multi-source gathering for a detailed market-entry analysis.

Which combination does the guidance support for this task?

- A. Neither is needed, since a single Web Search covers everything
- B. Brainstorming alone, since ideation naturally includes reasoning and sources
- C. Research combined with Extended Thinking for reasoning plus source-gathering
- D. Concise Style, which balances depth and brevity for you

Answer: C

Explanation:

A task that demands both deep reasoning and extensive sourcing can pair Extended Thinking with Research, since the two strategies are complementary. A single search, brainstorming, or a Style setting each address only part of the requirement.

Question: 7

Regenerating the same prompt gives you three noticeably different values for a key metric across three tries. You need a single reliable number for a finance summary due to leadership today.

Given the inconsistency across regenerations, what is the most appropriate action?

- A. Pick the value that appeared in two of the three tries
- B. Use whichever value is closest to what you already expected
- C. Average the three values and report the resulting mean figure
- D. Treat the inconsistency as a warning; verify the source data

Answer: D

Explanation:

Three different values for one metric means none can be trusted on its own. Verifying against the source data resolves the figure; selecting the most frequent, the expected, or the average only manufactures false confidence.

Question: 8

A team plans to run a Research session for every routine status question in a daily workflow.

What is the key cost consideration they should weigh?

- A. Research runs on a separate budget that does not touch message limits
- B. Research is unavailable on paid plans for daily use
- C. Research consumes usage quota faster than a normal message
- D. Research lowers the account's usage limit ceiling for the rest of the day

Answer: C

Explanation:

A Research run consumes usage-limit quota faster than an ordinary chat message because it issues many searches and page fetches, so a workflow that would trigger Research on every routine question should budget for that rather than treating it as free.

Question: 9

Profile instructions, project instructions, and a chosen Style relate to each other how on a single reply?

- A. Only the most specific layer takes effect at once
- B. The chosen Style overrides both instruction layers
- C. Project instructions replace profile instructions everywhere
- D. All three apply together within their own scopes

Answer: D

Explanation:

Profile instructions load account-wide and first, project instructions add project-scoped guidance, and a Style shapes tone and format per conversation. All three apply simultaneously within their own scopes; none replaces another, which is why off-tone output can come from a conflict between them.

Question: 10

Anthropic will sign a Business Associate Agreement for HIPAA-eligible work, but that coverage applies to:

- A. Only specific eligible services, not consumer plans
- B. Every plan tier including Free and Pro accounts
- C. Any conversation where a written disclaimer is added
- D. All Team plans automatically, once it is signed

Answer: A

Explanation:

Anthropic will sign a BAA for HIPAA-eligible services, but that agreement covers only specific eligible services and does not automatically extend to Free, Pro, Max, or a standard Team plan. A company-wide agreement therefore does not authorize pasting protected health information into any plan.

Question: 11

A coordinator asks Claude to write one product announcement email from a short brief; the task is easy to describe and not pattern-heavy.

The best-fitting prompting approach is:

- A. Few-shot example pairs, since examples outperform instructions on simple one-off tasks
- B. Zero-shot instructions describing the email's goal, audience, and tone
- C. A chained sequence of five separate drafting prompts
- D. A prompt built from bracketed placeholder variables for each field

Answer: B

Explanation:

Zero-shot instructions are the default for single-turn tasks that are easy to describe and do not hinge on a repeated pattern. A one-off announcement email specified by goal, audience, and tone is a fit, whereas few-shot and chaining add unneeded overhead.

Question: 12

Claude drafts replies inside an existing email routine while a human still sends each one, leaving the surrounding process unchanged.

This integration approach is best described as:

- A. Redesigning the workflow around Claude's capabilities
- B. Automating the process so no human action remains
- C. Replacing the email system with a new platform
- D. Augmenting a single step of the existing process

Answer: D

Explanation:

Augmenting means Claude accelerates one step of a process that otherwise stays intact, such as drafting emails a human still sends, versus a heavier redesign of the whole process around Claude. Many effective integrations simply augment a single step rather than rebuilding the workflow from scratch.

Question: 13

According to Anthropic's responsible-use guidance, how should a business user treat Claude when checking a high-stakes fact?

- A. As one input to scrutinize against an authoritative source
- B. As a singular source of truth requiring no further checking
- C. As accurate whenever the response sounds confident and detailed

D. As reliable once the most capable model is selected

Answer: A

Explanation:

Anthropic guidance directs users never to treat Claude as a single source of truth and to scrutinize any high-stakes advice. The correct posture is to treat a response as one input verified against an authoritative source, since confident tone and capable models do not remove the need to check.

Question: 14

In claude.ai, what does using the Retry control on a specific response do?

- A. Edits the original prompt and discards later messages
- B. Overwrites the previous response, discarding the earlier version
- C. Generates a new answer to the unchanged prompt
- D. Publishes the response as an independently shareable link

Answer: C

Explanation:

Retry keeps the prompt fixed and generates a fresh answer variant, and a version selector lets the user toggle between them, which is the built-in way to compare outputs and spot inconsistency. Editing a message, by contrast, changes the prompt and forks a branch, so the two controls serve different purposes.

Question: 15

A colleague asks Claude to help build a tool that scrapes and compiles individuals' biometric data without their consent.

Under Anthropic's Usage Policy, this request is:

- A. Permitted if a qualified professional reviews the output
- B. Permitted as long as AI involvement is disclosed
- C. A prohibited use banned outright regardless of safeguards
- D. A high-risk use allowed with a human in the loop

Answer: C

Explanation:

Unauthorized data collection and biometric misuse are prohibited outright, unlike high-risk uses that are allowed with safeguards such as human review and disclosure. No amount of review or disclosure converts a prohibited activity into an acceptable one.

Thank You for Trying Our Product

For More Information – **Visit link below:**

<https://www.examsboost.com/>

15 USD Discount Coupon Code:

G74JA8UF

FEATURES

- ✓ **90 Days Free Updates**
- ✓ **Money Back Pass Guarantee**
- ✓ **Instant Download or Email Attachment**
- ✓ **24/7 Live Chat Support**
- ✓ **PDF file could be used at any Platform**
- ✓ **50,000 Happy Customer**



Visit us at: <https://www.examsboost.com/test/ccao-f>