



Vendor: Anthropic

Exam Code: CCAR-P

Exam Name: Claude Certified Architect - Professional

Version: DEMO

QUESTION 1

You are selecting a model for a new production workload and must complete the upstream steps before testing candidate models empirically. Which two steps must be completed BEFORE running a representative sample on a candidate model? (Select two.) Each correct answer presents part of the solution.

- A. Select the lightest model that consistently meets the quality bar across the sample.
- B. Retest the model choice when a new model version becomes available.
- C. Choose a candidate model based on the requirements profile and known capabilities.
- D. Define the quality bar, latency tolerance, and expected volume for the workload.
- E. Sign off the production rollout plan with the platform and security teams.

Answer: CD

Explanation:

Empirical model testing requires a defined decision frame. Option D must occur first because the quality threshold, acceptable latency, expected request volume, and associated cost envelope determine what constitutes a successful result. Without these criteria, a representative-sample test can produce measurements but cannot support a defensible selection. Option C is the next upstream step. Known model capabilities and the workload profile are used to identify a plausible candidate rather than testing models arbitrarily. Anthropic's model-selection guidance recommends evaluating capability, speed, cost, and effort requirements before narrowing the candidate set. Choosing the Right Model

QUESTION 2

You are presenting an architectural decision to a mixed audience that includes an executive sponsor and the engineering leads who will implement the decision. Which presentation strategy best serves both audiences?

- A. Open with a deep dive into low-level implementation details targeted at engineering leads, and stop there without addressing the business outcomes or trade-offs the executive sponsor needs.
- B. Lead with the decision, the business outcomes it serves, and the trade-offs accepted; follow with the technical rationale, alternatives, and implementation implications for the engineering audience.
- C. Skip the rationale, alternatives, and trade-off discussion entirely and simply announce the chosen decision, leaving both audiences without the context needed to implement or validate it.
- D. Present a single undifferentiated narrative that addresses technical and business concerns with equal weight throughout, treating both audiences as requiring the same depth on every section.

Answer: B

Explanation:

Option B uses progressive disclosure to meet the decision needs of both audiences. The executive sponsor first receives the selected direction, intended business outcomes, major constraints, cost or performance implications, and accepted trade-offs. This establishes why the decision matters and whether it remains aligned with strategic objectives. The presentation can then move into the technical rationale required by engineering leads: component responsibilities, integration contracts, model and retrieval choices, security controls, rejected alternatives, implementation dependencies, evaluation gates, operational consequences, and migration requirements. This structure gives engineers enough evidence to implement the decision without forcing the executive audience to interpret unnecessary low-level detail.

QUESTION 3

A solutions architect is analyzing stakeholder feedback collected after the first quarter of a

Claude- powered procurement automation deployment. The feedback includes four statements: (1) "Our procurement team is processing 3x more purchase orders per analyst per day." (2) "We have eliminated the manual data entry role entirely and redeployed those staff to vendor relationship management." (3) "The API integration costs are running 40% over the projected per-transaction budget." (4) "Response latency during end-of-month batch runs is averaging 11 seconds, against our committed 5-second SLA." Which of the following correctly identifies the primary business value pillar each stakeholder statement represents?

- A. Statement 1: Efficiency; Statement 2: Transformation; Statement 3: Solution Cost; Statement 4: Performance SLA
- B. Statement 1: Productivity; Statement 2: Efficiency; Statement 3: Solution Cost; Statement 4: Performance SLA
- C. Statement 1: Transformation; Statement 2: Productivity; Statement 3: Solution Cost; Statement 4: Efficiency
- D. Statement 1: Efficiency; Statement 2: Productivity; Statement 3: Performance SLA; Statement 4: Solution Cost

Answer: A

Explanation:

Statement 1 measures efficiency because the same operating unit--one procurement analyst--now processes three times as many purchase orders per day. It reflects increased throughput and reduced effort per transaction within the existing procurement process.

Statement 2 represents transformation. The organization has not merely accelerated manual data entry; it has eliminated that role from the workflow and redirected employees toward higher-value vendor relationship management. This is a structural redesign of responsibilities and operating practices.

Statement 3 is a solution-cost issue because the measured API integration expense exceeds the approved per-transaction financial model by 40%. The architect must investigate token consumption, model routing, caching, retrieval overhead, tool calls, and integration charges while protecting the required quality level.

Statement 4 directly concerns the performance SLA. The observed 11-second latency is compared with an explicit contractual or operational target of five seconds. This is not simply an efficiency complaint; it is a measurable breach of a defined service-level commitment.

The architect should map each stakeholder statement to its dominant value pillar before recommending remediation because throughput, organizational transformation, cost control, and SLA compliance require different architectural responses. Study Guide references/topics: Business-value pillars; efficiency; transformation; solution cost; performance SLAs; stakeholder-outcome classification.

QUESTION 4

During an architectural review, the security team identifies a risk that adversarial content injected into retrieved documents could manipulate the model's behavior. Which mitigation most directly addresses this threat?

- A. Treat all retrieved content as untrusted input and apply input classifiers with output validation.
- B. Require citations for each claim and constrain responses to source-supported content.
- C. Restrict outbound tool calls to an approved destination allow-list.
- D. Score outputs against a stable adversarial evaluation set on each model-version change.

Answer: A

Explanation:

The scenario describes indirect prompt injection: malicious instructions are embedded in third-party content that Claude processes as retrieved evidence. Option A addresses the threat at the appropriate boundary by classifying retrieved material as untrusted, screening it before use, and

validating the model's output before it can influence downstream actions. Anthropic recommends placing untrusted third-party content in clearly identified tool-result structures, explicitly instructing Claude that retrieved content cannot override system or user instructions, screening tool output for injection attempts, and applying output validation.

QUESTION 5

You are explaining the precedence of Claude Code configuration scopes to the team. Which precedence ordering, from highest to lowest, is correct?

- A. Local -> managed -> user -> command-line arguments -> project
- B. Project -> user -> managed -> local -> command-line arguments
- C. Managed -> command-line arguments -> local -> project -> user
- D. User -> project -> local -> command-line arguments -> managed

Answer: C

Explanation:

The correct general precedence is managed configuration, command-line arguments, local configuration, project configuration, and then user configuration. Managed configuration has the highest authority because it contains centrally administered organizational and compliance controls that ordinary engineers must not be able to override. Command-line arguments provide temporary session-level overrides where the relevant setting permits them. Local settings apply to one user within a specific repository and take precedence over the shared project configuration. Project settings apply to every contributor using that repository and override general user preferences. User settings form the lowest tier because they provide personal defaults across projects and apply only when no higher-precedence source supplies a value. Anthropic documents this order explicitly: managed, command line, local, project, and user. It also notes that permission rules may merge rather than behave as simple scalar overrides, and certain security-sensitive keys have specialized precedence behavior. Claude Code Configuration Scopes The ordering allows personal customization without permitting it to defeat repository conventions or centrally mandated security requirements. Architects must still examine the merge behavior of the individual configuration property rather than assuming every array or permission rule is replaced wholesale.

QUESTION 6

A Claude architect needs to ensure that a security-hardening flag cannot be disabled by any individual engineer after it is set. Which configuration scope correctly enforces this requirement?

- A. User scope (~/.claude/settings.json) on each engineer's machine
- B. Environment variable defined in the CI/CD pipeline only
- C. Managed configuration applied centrally and marked as non-overridable
- D. Project scope (.claude/settings.json) committed to the repository

Answer: C

Explanation:

Managed configuration is designed for organization-wide security and compliance controls that individual users and repositories must not override. Applying the hardening flag centrally establishes an administrative policy boundary and prevents engineers from disabling the control through user, project, or local settings.

Anthropic identifies managed scope as the appropriate location for security policies, non-overridable compliance requirements, and standardized configurations deployed by IT or DevOps. Managed values ordinarily take precedence over command-line arguments and every user-controlled settings scope.

QUESTION 7

You are rolling out a standardized Claude Code configuration to an engineering team and must complete the planning steps before piloting the configuration. Which two steps must be completed BEFORE piloting the configuration with a small group of engineers? (Select two.) Each correct answer presents part of the solution.

- A. Define the project-scope baseline covering Model Context Protocol (MCP) servers, permission rules, and subagents.
- B. Onboard every engineer in the organization to the new configuration through mandatory training sessions.
- C. Roll out the stabilized configuration to additional teams with documentation and a defined support channel.
- D. Identify the team workflows, security boundaries, and which decisions belong to managed configuration versus project scope.
- E. Iterate the configuration based on the pilot findings and stabilize the baseline before broader rollout.

Answer: AD

Explanation:

Before a pilot begins, the architect must understand the workflows being enabled, the associated security boundaries, and the correct ownership of each configuration decision. Option D establishes this foundation. Organization-wide, non-overridable controls belong in managed configuration, whereas repository-specific MCP servers, shared permissions, and team subagents generally belong in project scope.

Option A then converts those decisions into an initial project baseline that the pilot group can use. The baseline should identify approved MCP endpoints, minimum permission rules, denied operations, subagent responsibilities, configuration ownership, and observable success criteria. Anthropic distinguishes managed, user, project, and local scopes according to who they affect and whether they are shared with the team.

QUESTION 8

A business sponsor has requested an AI solution to "improve customer experience." The sponsor cannot articulate which customer journey is failing, which metric reflects the failure, or which decisions the AI should support. The sponsor is asking you to begin design work next week. Which delegation-competency action should you take first?

- A. Recommend that the sponsor revise the request and resubmit it later for evaluation.
- B. Begin prototyping a generic assistant against the broad request before the next deadline.
- C. Propose a fixed scope that you commit to by default based on your own assumptions.
- D. Facilitate a structured discovery to define the failing decision and the target metric.

Answer: D

Explanation:

The request is not yet design-ready because it lacks a defined failure, affected user journey, supported decision, target population, and measurable outcome. The architect should therefore facilitate structured discovery rather than prematurely selecting an AI pattern or model. Discovery should identify where the current customer journey breaks down, who experiences the problem, what decision or task requires assistance, what data is available, and which metric would demonstrate improvement. The team can then establish baseline performance, target thresholds, operational constraints, risks, and explicit non-goals. Anthropic's evaluation guidance states that an LLM application should begin with clearly defined success criteria and a method for measuring them.

QUESTION 9

You are documenting an architectural decision to support future audit and onboarding. Which artifact is the strongest fit?

- A. A slide deck in a presentation folder with no accompanying written rationale.
- B. A code comment in a single file that contains an opinion of one engineer.
- C. An Architecture Decision Record that states the context, the decision, the alternatives considered, the consequences, and the date and authors.
- D. A short verbal note shared during a hallway conversation with no written record.

Answer: C

Explanation:

An Architecture Decision Record is designed to preserve a significant architectural choice in a concise, durable, and reviewable form. The context explains the problem, constraints, and forces that shaped the decision. The decision section records the selected approach. Alternatives demonstrate that other viable options were evaluated, while consequences identify benefits, costs, limitations, operational obligations, and residual risks. The date and authors establish accountability and historical sequence.

This structure supports audit because a reviewer can trace why a model, integration pattern, retrieval architecture, security boundary, or deployment approach was chosen. It also supports onboarding by preventing new team members from interpreting intentional design choices as arbitrary implementation details.

A slide deck may communicate a decision at a meeting but is weak as the sole authoritative record when it lacks written rationale. A code comment is too localized and may not capture system-level implications or rejected alternatives. A verbal conversation creates no durable evidence and depends entirely on individual memory.

The ADR should be version-controlled, linked to affected requirements and implementation artifacts, assigned an owner, and reviewed when its assumptions or operating conditions materially change.

QUESTION 10

You are building an evaluation pipeline for a Claude-based deployment and must complete the specification steps before running the deployment against the dataset. Which two steps must be completed BEFORE running the deployment against the evaluation dataset? (Select two.) Each correct answer presents part of the solution.

- A. Publish the aggregated metrics to a dashboard and gate releases on threshold checks.
- B. Curate and label the evaluation dataset to match the defined slices.
- C. Review failure cases with subject matter experts to refine the scoring rubric.
- D. Define the metrics and slices the framework will report across representative, edge, and adversarial cases.
- E. Score the deployment outputs against the reference labels and aggregate the metrics.

Answer: BD

Explanation:

The evaluation specification must exist before inference begins. Option D defines what the evaluation will measure and how results will be segmented. Representative cases measure normal workload performance, edge cases test boundary behavior, and adversarial cases examine safety, security, or robustness under hostile inputs.

Option B then supplies the correctly curated and labeled dataset corresponding to those slices. Labels, expected outcomes, grading rubrics, provenance, and dataset versions must be

established before the deployment generates outputs. Otherwise, evaluators risk modifying criteria after seeing results and introducing confirmation bias. Anthropic defines an evaluation as an input combined with grading logic used to measure success. Its evaluation guidance distinguishes tasks, trials, graders, traces, outcomes, and the harness that runs and aggregates them.

QUESTION 11

You are configuring tool permissions for a Claude-based assistant. The assistant's defined responsibilities require read access to a knowledge base and write access to a draft queue, and nothing else. Which scoping design best applies least privilege?

- A. Allow access to all tools with read permissions globally and restrict write permissions to the draft-queue tool, without scoping to only the specific tools required for defined tasks.
- B. Per-role allow-list of exactly the read-knowledge-base and write-draft-queue tools, enforced at the orchestration layer.
- C. Allow access to every tool in the catalog and rely on the model to decline tools it should not use.
- D. Disable all tools entirely to achieve a minimal permission surface, accepting that the assistant can no longer perform the read-knowledge-base or write-draft-queue tasks it was designed for.

Answer: B

Explanation:

Least privilege means granting the exact capabilities required for the defined responsibility--no more and no less. Option B permits the assistant to read the approved knowledge base and write only to the draft queue. Enforcing the allow-list at the orchestration layer creates a deterministic authorization boundary outside the model's prompt and reasoning. The rule should also be role-specific so different users or operational roles receive only the tools and data scopes they require. The knowledge-base tool should remain read-only, while the draft-queue tool should be restricted to draft creation rather than external sending, approval, or production publication.

QUESTION 12

You are a solution architect designing a Claude-based assistant with access to 60 internal tools across multiple business domains. Loading every tool definition on every request increases token usage and time to first response. Which design pattern best addresses this issue without sacrificing capability breadth?

- A. Apply progressive tool discovery so a curated initial subset is exposed and additional tools are loaded on demand based on the task.
- B. Use a separate model call to summarize all 60 tool definitions before each user turn.
- C. Increase the maximum context length and load all 60 tool definitions on every request, accepting the higher token cost and latency as necessary for full capability.
- D. Hard-code a fixed set of five tools per request to reduce token usage, regardless of whether those tools are relevant to the current task.

Answer: A

Explanation:

Progressive tool discovery preserves the complete capability catalog while limiting initial context to a curated set of common common common tools and a discovery mechanism. When the task requires an additional capability, Claude searches the catalog and loads only the relevant tool definitions. Anthropic's Tool Search mechanism implements implements this pattern through deferred loading. The context initially contains the search search search tool and any deliberately non-deferred tools; relevant definitions are expanded only when discovered. This reduces token consumption and helps maintain tool-selection accuracy in large catalogs.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



10% Discount Coupon Code: ASTR14