Anthropic

(Claude Certified Architect Professional (CCAR-P))

Claude Certified Architect Professional

Total: **65 Questions**

Link:

reaches the customer. This complements prompt-level guardrails and provides an additional control if the model attempts to generate restricted advice.

Final Answer: D and E

Question: 3

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 every tool in the catalog and rely on the model to decline tools it should not use.

B. Per-role allow-list of exactly the read-knowledge-base and write-draft-queue tools, enforced at the orchestration layer.

C. 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.

D. 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.

Answer: B

Explanation:

A. Allow access to every tool in the catalog and rely on the model to decline tools it should not use.

Incorrect

This does not follow the **least-privilege principle** because the assistant receives access to tools it does not need. Relying on the model to voluntarily avoid those tools is not an adequate security control.

B. Per-role allow-list of exactly the read-knowledge-base and write-draft-queue tools, enforced at the orchestration layer.

Correct

This is the strongest least-privilege design. The assistant receives access to **only the two tools required for its defined responsibilities**. Enforcing the allow-list at the orchestration layer provides a technical permission boundary rather than relying on the model's behavior.

C. 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.

Incorrect

Although disabling all tools minimizes permissions, it also prevents the assistant from performing its **required responsibilities**. Least privilege means granting the minimum permissions necessary to accomplish the intended task, not removing necessary access.

D. 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.

Incorrect

This still grants **unnecessary access** to tools outside the assistant's responsibilities. Least privilege requires permissions to be scoped to the specific capabilities the assistant actually needs.

Incorrect

Maximum reasoning capability is not required for every moderation task. Using a highly capable model for simple, high-volume classification can result in **unnecessary cost and latency**.

Final Answer: B. Haiku, because its latency and cost profile align with high-volume classification workloads that require limited reasoning depth.

Question: 6

You are supporting an engineer whose Claude Code session reports a permission denial when the engineer expected the action to be allowed.
Which resolution step is most appropriate?

A. Inspect the active permission rules across all scopes-managed, command-line, local, project, and user-to identify which rule is denying the action and confirm whether it should be amended or remain denied.

B. Tell the engineer to retry the denied action verbatim repeatedly until the denial stops appearing, without inspecting the active permission rules to determine whether the denial is intentional.

C. Disable the permission system entirely across all scopes to remove the denial, eliminating all tool-pattern and deny-rule controls rather than identifying and amending the specific rule.

D. Grant the engineer unrestricted Bash access to bypass the specific rule causing the denial, removing all tool-pattern constraints rather than amending only the rule in

Answer: A

Explanation:

A. Inspect the active permission rules across all scopes — managed, command-line, local, project, and user — to identify which rule is denying the action and confirm whether it should be amended or remain denied.

Correct

The appropriate approach is to **diagnose the active permission configuration** rather than bypassing the security controls. Reviewing the applicable rules across scopes helps identify the exact rule responsible for the denial and determine whether the rule is intentional or needs to be changed.

B. Tell the engineer to retry the denied action verbatim repeatedly until the denial stops appearing, without inspecting the active permission rules to determine whether the denial is intentional.

Incorrect

Repeated retries do not address the underlying permission configuration. The denial may be intentional, so the correct approach is to **inspect and understand the active rules first**.

C. Disable the permission system entirely across all scopes to remove the denial, eliminating all tool-pattern and deny-rule controls rather than identifying and amending the specific rule.

Incorrect

Disabling the permission system removes important **security controls** and violates the principle of least privilege. Only the specific rule that is incorrectly denying the required action should be considered for amendment.

D. Grant the engineer unrestricted Bash access to bypass the specific rule causing the denial, removing all tool-pattern constraints rather than amending only the rule in question.

Question: 9

A Claude architect is auditing configuration scope assignments.
Which two statements correctly identify an appropriate use of user-scope configuration versus other scopes? (Select two.)

A. Defining MCP server endpoints shared by all contributors to a specific repository.
B. Storing API authentication keys so they are not committed to version control.
C. Enforcing a company-wide policy that disables a feature for all engineers.
D. Persisting personal editor theme preferences that follow an engineer across projects.
E. Saving a preferred Claude response language that applies to all repositories the engineer uses.

Answer: DE

Explanation:

A. Defining MCP server endpoints shared by all contributors to a specific repository.

Incorrect

MCP servers shared by contributors to a particular repository are better suited to **project scope**, because project-level configuration can be shared and maintained with the project rather than tied to one individual engineer.

B. Storing API authentication keys so they are not committed to version control.

Incorrect

Authentication keys and other secrets require **secure secret-management practices**. Simply placing them in user-scope configuration is not the appropriate general solution for protecting sensitive credentials.

C. Enforcing a company-wide policy that disables a feature for all engineers.

Incorrect

Company-wide security or policy requirements should be enforced through **managed/organizational configuration**, rather than user scope. Individual engineers should not be able to override centrally required controls.

D. Persisting personal editor theme preferences that follow an engineer across projects.

Correct

User scope is appropriate for **personal preferences** that should follow an individual engineer across different projects. An editor theme is a good example because it does not need to be shared with the entire team.

E. Saving a preferred Claude response language that applies to all repositories the engineer uses.

Correct

A preferred response language is also an **individual user preference**. User scope allows the preference to persist across the engineer's different repositories and projects without imposing it on other team members.

Final Answer: D and E

different aspects of the decision.
How should you structure the recommendation document?

A. Lead with detailed cost projections across the full multi-year horizon before the rationale section.
B. Lead with the technical architecture diagram and the full component list before any other section.
C. Present the same dense narrative throughout with no stakeholder differentiation in any section.
D. Lead with the architectural decision, then address each stakeholder's primary concerns directly.

Answer: D

Explanation:

A. Lead with detailed cost projections across the full multi-year horizon before the rationale section.

Incorrect

Cost is highly relevant to the CFO, but leading exclusively with detailed financial projections may not address the **legal, technical, and strategic concerns** of the other stakeholders.

B. Lead with the technical architecture diagram and the full component list before any other section.

Incorrect

The technical architecture is particularly important to the CIO, but beginning with a detailed component list may overwhelm stakeholders who primarily need to understand the **decision, rationale, risks, and business impact**.

C. Present the same dense narrative throughout with no stakeholder differentiation in any section.

Incorrect

Different stakeholders have different responsibilities and concerns. A one-size-fits-all narrative may fail to address the **specific information each decision-maker needs**.

D. Lead with the architectural decision, then address each stakeholder's primary concerns directly.

Correct

The strongest approach is to **state the recommendation clearly at the beginning** and then address the key concerns of each stakeholder. For example, the CFO may focus on cost and ROI, the General Counsel on legal risk and compliance, and the CIO on architecture, security, integration, and operational considerations.

This structure makes the recommendation **decision-oriented, concise, and relevant to the entire steering committee**.

Final Answer: D. Lead with the architectural decision, then address each stakeholder's primary concerns directly.

Question: 13

You are investigating an MCP server that fails on first launch but succeeds on subsequent runs. System permission dialogs appeared during the first launch.
Which response is most appropriate?

A. Treat the first-run failure as a permanent fault, replace the MCP server, and do not verify whether subsequent runs succeed.

A. Prompt-level instructions that treat user content as untrusted data, runtime classifiers that detect override attempts, scoped tool permissions that cannot be elevated by user content, and audit logging of attempts.

B. Granting users any privilege level they assert in their message content, on the assumption that cooperative behavior requires honoring self-declared authority without independent verification.

C. Trusting that the model will intrinsically recognize and reject all bypass attempts without prompt-level instructions, runtime classifiers, scoped permissions, or audit logging.

D. Removing all safety restrictions and guardrails to eliminate the attack surface that bypass attempts target, accepting that this makes the assistant unrestricted for all inputs.

Answer: A

Explanation:

A. Prompt-level instructions that treat user content as untrusted data, runtime classifiers that detect override attempts, scoped tool permissions that cannot be elevated by user content, and audit logging of attempts.

Correct

This approach provides **defense in depth**. User-provided claims of authority are treated as untrusted input, runtime controls can identify suspicious override attempts, and scoped permissions prevent the user from escalating privileges simply through prompt content. Audit logging also supports **monitoring, investigation, and incident response**.

B. Granting users any privilege level they assert in their message content, on the assumption that cooperative behavior requires honoring self-declared authority without independent verification.

Incorrect

User-supplied text is not a reliable source of authorization. Allowing users to declare their own privileges creates a serious **privilege-escalation risk**.

C. Trusting that the model will intrinsically recognize and reject all bypass attempts without prompt-level instructions, runtime classifiers, scoped permissions, or audit logging.

Incorrect

A model should not be treated as the **sole security boundary**. Multiple technical and procedural controls are needed because adversarial inputs can sometimes bypass model-level instructions.

D. Removing all safety restrictions and guardrails to eliminate the attack surface that bypass attempts target, accepting that this makes the assistant unrestricted for all inputs.

Incorrect

Removing safety controls does not mitigate the attack; it **eliminates the protections** that are intended to prevent unsafe behavior.

Final Answer: A. Use layered prompt, runtime, permission, and audit controls.

Question: 16

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?

D. Regression tests against a stable reference set of previously known-good outputs.

E. Adversarial tests using prompt-injection and malformed-input cases.

Answer: AE

Explanation:

A. Integration tests that verify end-to-end pipeline behavior across all components.

Correct

Integration tests verify that the different components of the pipeline work correctly **together**. For a sensitive financial-data workflow, this is important because failures can occur at the boundaries between components even when individual components work correctly in isolation.

B. Smoke tests that verify core paths after each deployment.

Incorrect

Smoke tests are useful for quickly checking whether the system's **basic functionality** works after deployment, but they do not provide sufficient coverage for adversarial safety or detailed cross-component correctness.

C. Unit tests targeting only individual prompt-template rendering logic.

Incorrect

Unit tests are valuable for testing isolated logic, but testing only prompt-template rendering does not adequately cover the **end-to-end pipeline** or attacks such as prompt injection and malformed inputs.

D. Regression tests against a stable reference set of previously known-good outputs.

Incorrect

Regression testing is useful for detecting **unexpected changes in output behavior**, but it does not directly provide targeted coverage for adversarial attacks or verify all cross-component interactions.

E. Adversarial tests using prompt-injection and malformed-input cases.

Correct

Adversarial testing directly evaluates how the system behaves under **malicious, unexpected, or malformed inputs**. For a financial-data pipeline, this helps identify safety weaknesses and potential ways attackers could manipulate the system.

Final Answer: A and E

A = Cross-component correctness

E = Safety under attack

Question: 19

You are running a risk assessment on a planned Claude-based deployment and must complete the inventory steps before assessing threats against assets.
Which two steps must be completed BEFORE assessing threats against assets to estimate likelihood and impact? (Select two.) Each correct answer presents part of the solution.

D = Collaboratively revise the SLA when justified by production constraints