QUESTIONS & *ANSWERS*

verified

DBCLOUD EXAMS

EXPERIENICIA Y TECNOLOGÍA

DBCLOUD CERTIFIED

Technology Provider: Anthropic

Exam: Claude Certified Architect – Professional

Number of questions in the database: 71

Exam Version: OCT.2026

Question #1 Topic 1

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.

A. Document the assessment outcome with risks, mitigations, residual risk, and acceptance owners.

B. Recommend mitigations and residual-risk acceptance for the risks that remain after analysis.

C. Identify the assets that the deployment touches, along with the sensitivity of each asset.

D. Recommend mitigations and residual-risk acceptance for risks that remain.

E. Enumerate the threat actors and attack vectors relevant to the deployment.

Correct Answer: CE

Question #2 Topic 1

You are building an ethics-review checklist for deployments supported by artificial intelligence.

Which two checks belong on the list? (Select two.)

Each correct answer presents a complete solution.

A. Confirm that vendor licensing terms permit the planned production use of the model.

B. Verify that outputs do not rely on generalizations about people that the underlying data does not support.

C. Confirm that high-impact decisions retain human accountability rather than being attributed to the model.

D. Restrict ethics review to outputs that exceed a defined model-confidence threshold.

E. Confirm that latency and throughput targets are met across supported user populations.

Correct Answer: BC

https://dbcloud.academy/

Question #3 Topic 1

You are reviewing instrumentation in a multi-agent system.

Which two findings constitute valid observability gaps in the instrumentation? (Select two.) Each correct answer presents a complete solution.

A. Trace spans for each agent step are exported to the shared distributed-tracing backend.

B. Latency and token usage on every span are emitted to the central metrics pipeline.

C. Tool-call payloads and outcomes are recorded with redaction applied to known sensitive fields.

D. Model identity and version on each turn are not recorded with the turn artifacts.

E. Request-scoped correlation identifiers do not propagate across agent and tool calls.

Correct Answer: DE

Question #4 Topic 1

You are supporting an EU-based deployment with GDPR obligations.

Which combination of measures best supports the deployment's GDPR posture?

A. enterprise-tier deployment with a signed Data Processing Addendum, defined data-retention configuration, redaction of personal data not needed for the task, and documented data-subject-rights handling

B. disabling all data-retention configuration, redaction controls, and data-subject-rights handling to simplify day-to-day operations, accepting the resulting GDPR compliance exposure

C. using a personal Claude account tier for processing EU personal data at scale, with no signed Data Processing Addendum and no documented data-subject-rights handling procedure

D. pasting full EU personal data into every prompt to "give Claude complete context" without considering purpose limitation, data minimization, or the organization's Data Processing Addendum obligations

Correct Answer: A

https://dbcloud.academy/

Question #5 Topic 1

You are selecting a pattern for a compliance Q & A assistant that must answer policy questions with citations to the authoritative internal source set. Latency, cost, and audit predictability are prioritized.

Which pattern is the best fit?

A. Augmented LLM with retrieval-augmented generation over the indexed authoritative corpus and citation rendering on each answer.

B. Multi-agent orchestration with a planner, researcher, and writer agent for every query.

C. A pure agent loop with open web-browsing tools to surface the most current policy information, without constraining retrieval to the authoritative internal corpus.

D. A static prompt with the entire policy corpus concatenated into every request.

Correct Answer: A

Question #6 Topic 1

A platform team operates a self-hosted multi-agent system on Kubernetes that orchestrates seven specialized agents for invoice processing. The team spends approximately 40 percent of engineering capacity on infrastructure maintenance, message bus reliability, and agent state recovery. The CFO has asked you to evaluate moving to managed agent infrastructure to reclaim engineering capacity. The security officer requires that all customer financial data remain within an approved network boundary.

Which factor should most heavily influence your recommendation?

A. Whether the current seven agents map cleanly to the patterns supported by managed agents.

B. Whether managed agents reduce per-invoice token costs across the existing processing volume.

C. Whether managed agents support the current bus topology used by the platform team.

D. Whether managed agent data handling satisfies the network boundary required by security.

Correct Answer: D

https://dbcloud.academy/