



Welcome to Cert007 - Your Ultimate IT Certification Partner



➤ Real Exam Questions

➤ Free Updates

➤ Expert Support

➤ Instant Access

➤ Money-Back Guarantee



Visit us at <https://www.cert007.com/> for more information

Exam : AIB-C01

Title : AWS Certified AI Business Strategist

Version : DEMO

1. A wholesaler calculates shipping fees from validated order totals, delivery regions, and membership levels. An approved policy table specifies the fee for every valid combination. The company needs consistent decisions and a record of the policy version used.

Which approach best meets these needs?

- A. Use an AI agent to retrieve the policy and decide each shipping fee.
- B. Apply a versioned business-rule table to the validated fields in each request.
- C. Prompt a language model with the current policy to determine each shipping fee.
- D. Train a classifier on past orders to predict the fee for each new request.

Answer: B

Explanation:

A versioned rule table applies the complete policy directly. Recording the applied policy version alongside each decision provides traceability; learning or inferring a fee adds unnecessary uncertainty when the inputs and required decisions are already defined.

2. A service company wants an assistant to answer questions using maintenance bulletins that change weekly. The collection is too large to include in a single prompt. Technicians need answers linked to the relevant source passages, without retraining a model after each update.

Which approach best meets these requirements?

- A. Include examples of well-written answers in the prompt and use the model's existing knowledge.
- B. Provide an updated summary of all bulletins as the context for each question.
- C. Retrieve relevant passages from the updated collection and generate answers with source references.
- D. Fine-tune the model on each weekly release and answer from its learned knowledge.

Answer: C

Explanation:

Retrieval supplies relevant passages from the current bulletins and allows the assistant to attach source references without retraining after each update. A collection summary can omit needed details, while fine-tuning or style examples alone do not provide this retrieval-and-reference workflow.

3. An AI invoice-processing project avoided \$180,000 in contractor payments during its first year. Implementation cost \$60,000, and total operating costs for that year were \$60,000. These are the only cash benefits and costs. Finance defines cash-based ROI as (cash benefits minus project costs) divided by project costs, multiplied by 100%.

What is the project's first-year cash-based ROI?

- A. 33.3%
- B. 200.0%
- C. 50.0%
- D. 150.0%

Answer: C

Explanation:

Total project costs are $\$60,000 + \$60,000 = \$120,000$. Cash-based ROI is $(\$180,000 - \$120,000) \div \$120,000 \times 100\% = 50.0\%$; dividing benefits by costs without subtracting costs would produce 150.0%, which is not the stated ROI measure.

4. A manufacturer can fund one AI pilot this quarter. Its priority is reducing unplanned equipment downtime. Historical maintenance records are available, but collecting new sensor data would take longer than the quarter. A small team must deliver a narrowly scoped pilot.

Which initiative should the company prioritize?

- A. Pilot invoice coding using existing finance records to reduce accounts-payable processing time.
- B. Pilot maintenance-fault classification using existing records to help technicians route repair requests.
- C. Pilot automated maintenance scheduling across all plants, including integration with their different control systems.
- D. Pilot failure prediction using new sensor data collected after installing monitors across every production line.

Answer: B

Explanation:

A narrowly scoped fault-classification pilot uses available maintenance records and supports the downtime objective by helping route repairs. Sensor-based prediction cannot obtain its required data within the quarter, invoice coding targets a different outcome, and plant-wide scheduling exceeds the intended pilot scope.

5. A procurement team is testing an AI assistant that flags duplicate supplier invoices. Its goals are to prevent duplicate payments without delaying legitimate invoices. Reviewers can verify whether flagged invoices are true duplicates.

Which TWO metrics most directly measure the pilot's intended business outcomes?

Instruction: Select TWO answers.

- A. Total value of confirmed duplicate payments prevented before payment release
- B. Average time the model takes to assign a duplicate-risk score
- C. Percentage of legitimate invoices paid by their contractual due dates
- D. Percentage of flagged invoices that reviewers confirm are actual duplicates
- E. Percentage of submitted invoices that the assistant flags for review

Answer: A, C

Explanation:

Confirmed duplicate payments prevented measure the avoided financial loss, while on-time payment of legitimate invoices measures whether the process preserves payment timeliness. Flag volume, scoring speed, and the precision of flags are useful diagnostic measures, but they do not directly establish these two business outcomes.