KillTest

高品質　より良いサービス

練習問題

https://www.killtest.jp

1年間無料アップデート

Exam : **AIB-C01**

Title : AWS Certified AI Business Strategist

Version : DEMO

1.A wholesale distributor uses AI to recommend reductions in customers' trade-credit limits. Customers receive accurate notices describing AI's role and limitations, with a link to request human review. However, each notice gives only "risk assessment" as the reason for the reduction. Internal reviewers have verified account-specific records showing the information and decision factors used. The distributor wants customers to identify factual errors and meaningfully challenge reductions. Customer communications must protect other customers' information and provide reasons supported by the verified records.
Which action best addresses the remaining responsible AI need? Select ONE.
A. Prioritize transparency: expand the description of the AI workflow and its limitations, and direct customers to the existing review process for individual decisions.
B. Prioritize explainability: give each customer an understandable account of the verified factors behind the reduction and explain how to challenge relevant information during review.
C. Prioritize fairness: compare credit-limit reductions across customer groups, and give customers a summary of those results alongside access to the existing review process.
D. Prioritize explainability: generate an account-specific narrative from general credit policy and the customer's profile, and have staff approve its clarity before including it in the notice.
Answer: B
Explanation:
B addresses the missing explainability. Customers already receive accurate disclosure about AI's involvement and limitations, but they cannot understand why their own credit limit was reduced. An intelligible account grounded in the verified decision records lets them identify errors and challenge the result while protecting other customers' information. More general disclosure or group-level fairness findings would not explain an individual reduction. A generated narrative could sound convincing without reflecting the factors actually used; checking its clarity does not establish its factual connection to the decision.
A - Incorrect: Expanding the workflow description supports transparency, but the existing notice already accurately describes AI's role and limitations. This action still withholds the account-specific reasons customers need to evaluate their reduction. Access to human review alone does not supply that understanding.
B - Correct: This action addresses explainability by connecting an understandable customer-facing account to verified factors used in the individual decision. Explaining how to challenge relevant information makes the account useful for identifying errors and seeking reassessment. The distributor can provide these reasons without revealing other customers' information.
C - Incorrect: Comparing outcomes across customer groups can support fairness assessment. However, a summary of group-level results does not tell an individual customer which information or factors caused the credit-limit reduction. Even acceptable group-level findings would leave the stated explanation gap unresolved.
D - Incorrect: General policy and a customer profile can support a plausible narrative without establishing why this particular reduction occurred. Staff approval of clarity verifies readability, not correspondence with the actual decision factors. Generated prose could help communicate an explanation if grounded in the verified records and checked against them, but this plan omits those requirements.

2.A regional utility is assessing readiness for two AI initiatives. In a four-week internal experiment, AI will group sanitized historical staff support requests to identify training needs. A sponsor has approved the

budget, learning criteria, and stop conditions. The data owner has checked the dataset and permitted its use. Trained employees will review all outputs in a supported test environment, and outputs cannot trigger business actions.
The second initiative would automatically decide whether to disconnect customers for unpaid bills. Account-data accuracy across regions has not been validated, and accountable decision ownership, monitoring, and escalation arrangements have not been agreed.
Which readiness decision is most appropriate under these conditions? Select ONE.
A. Proceed with both initiatives, applying the experiment's sponsorship, staff participation, and technical support as evidence that the utility is ready to adopt AI.
B. Pause both initiatives until the utility has established the data quality, governance, and operational capabilities required for the customer disconnection service.
C. Proceed with the internal experiment within its approved boundaries, and defer the customer service until its data, decision ownership, monitoring, and escalation prerequisites are established.
D. Proceed with the internal experiment and launch the customer service in one region, using customer appeals and weekly reviews to develop its governance and operational arrangements.
Answer: C
Explanation:
Readiness must be assessed against each initiative's intended outcome, scope, and consequences. The internal experiment has an agreed purpose, permitted and checked data, trained reviewers, technical support, and bounded exposure, so useful learning can proceed. The customer service would make consequential decisions while essential data and accountability arrangements remain unresolved. Those gaps justify deferring that service without blocking the supported experiment. Neither transferring readiness from one initiative to another nor reducing the launch population establishes the missing prerequisites.
A - Incorrect: The sponsorship, participation, and support described apply to the internal experiment. They do not establish that customer account data are accurate or that the disconnection service has accountable ownership, monitoring, and escalation. Readiness for the experiment cannot justify proceeding with both.
B - Incorrect: The disconnection service's prerequisites are necessary for its consequential decisions, but the internal experiment does not make those decisions or change customer accounts. Pausing both would prevent bounded learning despite sufficient evidence for the experiment's narrower requirements.
C - Correct: This decision matches each commitment to its readiness evidence. The experiment can proceed within its approved limits, while the disconnection service requires validated account data and established ownership, monitoring, and escalation before launch.
D - Incorrect: A regional launch reduces the number of customers exposed but still permits consequential decisions using account data whose accuracy has not been validated. Customer appeals and weekly reviews do not establish reliable data, accountable ownership, or timely escalation before customers are affected.

3.An appliance repair company needs three outputs from an AI initiative: a numeric estimate of each job's duration in minutes, learned from historical completed jobs, to help schedule appointments; a predefined repair-type label for each new service request; and a newly composed technician handover note grounded in approved job facts. Pilot approval requires checking duration estimates against actual completion times, labels against reviewed examples, and notes for factual support and clarity.

Which proposal best matches these needs while accurately describing the relationship among AI, ML, and GenAI? Select ONE.

A. Use ML classification to assign short, medium, or long duration bands, ML classification to assign repair-type labels, and GenAI to compose handover notes. Treat all three as AI capabilities, and validate duration bands, repair labels, and note quality.

B. Use predictive ML to estimate duration in minutes, ML classification to assign repair-type labels, and GenAI to compose handover notes. Treat all three as AI capabilities, recognize that generative models can also classify, and apply the specified output checks.

C. Use predictive ML to estimate duration in minutes, ML classification to assign repair-type labels, and ML classification to select an existing standard handover note as its output label. Treat all three as AI capabilities, and validate estimates, labels, and selected notes.

D. Use GenAI to draft duration estimates with explanations, ML classification to assign repair-type labels, and GenAI to compose handover notes. Treat all three as AI capabilities, approve duration estimates when their explanations are convincing, and apply the specified label and note checks.

Answer: B

Explanation:

B matches each required result to its primary capability: predictive ML learns patterns from historical jobs to estimate a numeric duration, classification assigns a predefined repair-type label, and GenAI creates a new handover note from supplied facts. These capabilities belong within AI and can overlap; for example, a generative model can perform classification. The proposal also preserves the required evaluation of each output.

A substitutes duration categories for numeric estimates.

C selects an existing note instead of composing the required new content.

D substitutes convincing explanations for evidence that duration estimates agree with actual completion times.

A - Incorrect: Classification is appropriate for the predefined repair types, and GenAI is appropriate for composing the handover notes. However, short, medium, and long are categorical labels. They do not provide the numeric duration estimates in minutes required for appointment scheduling. Validating the bands does not establish that the required numeric predictions are available or accurate.

B - Correct: Predictive ML fits estimating job duration from historical completed jobs. ML classification fits assigning service requests to predefined repair types. GenAI fits composing new handover notes grounded in approved facts. All are AI capabilities, and acknowledging that generative models can also classify preserves their overlapping uses. Checking estimates against completion times, labels against reviewed examples, and notes for support and clarity evaluates the actual business outcomes.

C - Incorrect: The duration prediction and repair-type classification match their respective needs. The handover step does not: assigning an existing standard note as a class label selects previously written content. It does not compose the new note grounded in the current job's approved facts that the company requires. Validating the selected note does not resolve that output mismatch.

D - Incorrect: A generative model is not automatically unsuitable for a workflow involving predictions, but convincing explanations do not establish prediction accuracy. The company explicitly requires duration estimates to be checked against actual completion times. Approving estimates based on persuasive wording therefore fails the pilot criterion, even though the repair-type classification and handover drafting steps fit their purposes.

4.A distributor must introduce AI-assisted classification and routing of standard warranty requests within six weeks. Its budget for implementation and the first 24 months of operation is $90,000. It can assign a part-time business process owner, but no internal developers or AI operations staff during that period. Company policy prohibits suppliers from using warranty-request content to train their own models.
The proposals below have firm delivery estimates and include all required costs for the expected workload. Product pilots have confirmed classification quality and integration with the existing case-management system. Contract terms cannot be changed before the deadline.
Which sourcing plan meets all the stated constraints? Select ONE.
A. Build an internal application using an existing AI service for $45,000 over 24 months. Delivery takes 14 weeks, and internal staff must develop, maintain, and operate the application.
B. Buy Product X with supplier-managed operation and support. Delivery takes five weeks; integration costs $12,000, subscription and usage cost $48,000, and support costs $18,000 over 24 months. Review has approved its data-use, compliance, support, and exit terms.
C. Commission a partner to build and operate a custom application under approved data-use, compliance, support, and exit terms. Delivery takes six weeks; implementation costs $60,000, followed by $42,000 in mandatory operating charges over 24 months.
D. Buy Product Y with supplier-managed operation and support. Delivery takes four weeks; integration costs $6,000, subscription and usage cost $24,000, and support costs $18,000 over 24 months. Its contract permits supplier model training using warranty-request content.
Answer: B
Explanation:
Product X meets every stated constraint: its pilot confirms functional quality and integration, its five-week delivery meets the deadline, and its complete 24-month cost is $78,000, below the $90,000 cap. Supplier-managed operation and support fit the available internal capacity, and the reviewed contractual terms meet the organization's requirements. Internal development misses the deadline and requires unavailable staff. The partner proposal is a feasible delivery approach but costs $102,000 across the required period, exceeding the budget despite its lower implementation fee. Product Y costs only $48,000 but permits a prohibited use of customer content. The decision therefore depends on complete lifecycle fit, including timing, capacity, contractual suitability, and total cost.
A - Incorrect: The $45,000 estimate is within budget, but a 14-week delivery cannot meet the six-week deadline. The plan also depends on internal development and operational capacity that the distributor does not have. Using an existing AI service does not remove ownership of the application.
B - Correct: Product X delivers in five weeks and costs $12,000 + $48,000 + $18,000 = $78,000 over the required period. Its validated capability and integration satisfy the standard business need. Supplier-managed operation and support fit the staffing constraint, and the approved terms satisfy the contractual requirements.
C - Incorrect: The partner supplies the missing delivery and operating capacity and meets the six-week deadline under acceptable terms. However, the mandatory operating charges belong in the budget comparison: $60,000 + $42,000 = $102,000. This exceeds the $90,000 cap, so the lower implementation fee does not make the complete plan affordable.
D - Incorrect: Product Y meets the deadline, fits the staffing constraint, passes the functional pilot, and costs $48,000 in total. However, its fixed contract permits the supplier to train its models using warranty-request content, directly conflicting with company policy. Its lower price cannot compensate for failing a mandatory requirement.

5.An online retailer collects identifiable customer support conversations to resolve service issues. Customers are told that their conversations are used only for that purpose. A marketing team proposes using AI to analyze those conversations and select personalized promotions. The security team confirms that the proposed analysis would use authorized employees, authenticated access, encryption, and effective access controls. No assessment of the appropriateness of the new use has occurred. Company policy requires both appropriate use of personal information and protection against unauthorized access before an AI initiative is approved.

Which recommendation addresses the unresolved concern? Select ONE.

A. Treat this as a security decision; require a separate authenticated workspace and renewed access approval before allowing the campaign analysis.

B. Treat this as a privacy decision; pause campaign analysis until the proposed use is assessed and its conflict with the stated service purpose is resolved.

C. Treat this as a transparency decision; disclose that AI selects promotions and describe its limitations before allowing the campaign analysis.

D. Treat this as an explainability decision; require understandable reasons linking each promotion to conversation content before allowing the campaign analysis.

Answer: B

Explanation:

B is correct because the unresolved issue is whether personal information collected for service resolution may appropriately be used for personalized marketing. Authorized access and effective security controls establish protection against unauthorized access; they do not establish that a secondary use is appropriate. The leader should pause the proposed analysis while the privacy owner assesses the new purpose and resolves its conflict with the stated service-only use. Additional authentication, disclosure of AI involvement, or understandable promotion-selection reasons cannot independently resolve that conflict.

A - Incorrect: A separate authenticated workspace and renewed access approval could address access-related risks. However, authorized users and effective access controls are already established. Additional security controls would not determine whether using support conversations for marketing is appropriate.

B - Correct: Privacy concerns whether personal information is obtained, used, shared, and managed appropriately for its purpose. Marketing analysis introduces a purpose outside the stated service use. Pausing that analysis for privacy assessment addresses the unresolved condition without assuming that secure access authorizes the new use.

C - Incorrect: Disclosing AI involvement and its limitations supports transparency and informed engagement. It does not resolve whether support conversations may appropriately be repurposed for marketing. Approval based only on this disclosure would leave the purpose conflict unaddressed.

D - Incorrect: Understandable reasons for selecting promotions support explainability. A promotion could be fully explainable while relying on personal information used for an inappropriate purpose. Explaining the selection therefore does not resolve the privacy decision.