

CERT-PASS

AWS AWS Certified AI Practitioner AIF-C01

Free Practice Questions Preview

Here are 35 sample questions to help you get started. Unlock the full exam to access all 1100+ questions with detailed explanations.

Question 1 : Domain 1: Fundamentals of AI and ML

[Easy] Matching-style question: A cloud practitioner at a regional bank is validating a study map before an architecture workshop. Which option correctly matches every item to its primary purpose?

- A. Amazon Transcribe = text to speech; Amazon Polly = speech to text; Amazon Comprehend = sentiment and entity analysis
- B. Amazon Transcribe = speech to text; Amazon Polly = sentiment and entity analysis; Amazon Comprehend = text to speech
- C. Amazon Transcribe = sentiment and entity analysis; Amazon Polly = text to speech; Amazon Comprehend = speech to text
- D. Amazon Transcribe = speech to text; Amazon Polly = text to speech; Amazon Comprehend = sentiment and entity analysis**

Answer: D

Option D correctly maps all three items. The distractors swap closely related concepts or services, which is a common exam trap.

Question 2 : Domain 1: Fundamentals of AI and ML

[Easy] A governance checkpoint at a payment processor confirms that the organization has intermittent prediction traffic and wants to avoid managing endpoint capacity. The solution should remain easy to govern. Which option is the strongest recommendation?

- A. Use a nightly job even when users are waiting.
- B. Use only a data warehouse.
- C. Provision a permanently oversized endpoint.
- D. Evaluate serverless inference.**

Answer: D

Serverless inference can suit intermittent traffic when the team wants reduced capacity management. Why the strongest distractor fails: A permanently oversized endpoint may create unnecessary cost.

Question 3 : Domain 1: Fundamentals of AI and ML

[Easy] A global wholesaler needs to predict the numeric amount of electricity demand for the next hour. The choice must address the business requirement directly. Which action should the team take next?

- A. Use classification.
- B. Use Amazon Lex.
- C. Use clustering.
- D. Use regression.**

Answer: D

Regression predicts a continuous numeric value. Why the strongest distractor fails: Classification predicts categories rather than a numeric forecast.

Question 4 : Domain 1: Fundamentals of AI and ML

[Easy] During a design review, the AI program lead for an online retailer explains that the organization needs to identify sentiment and entities in customer reviews. The organization needs an approach that can be evaluated with measurable outcomes. Which recommendation is the best fit?

- A. Use Amazon Rekognition.
- B. Use Amazon Transcribe.
- C. Use Amazon Comprehend.**
- D. Use Amazon Polly.

Answer: C

Amazon Comprehend performs NLP tasks such as sentiment and entity detection. Why the strongest distractor fails: Amazon Rekognition analyzes images and video rather than review text.

Question 5 : Domain 1: Fundamentals of AI and ML

[Medium] During a design review, the application owner for a manufacturing group explains that the organization must minimize false positives when flagging legitimate transactions for manual review. The design must minimize unnecessary operational overhead. Which answer best addresses the requirement?

- A. Prioritize precision.**
- B. Prioritize recall only.
- C. Use ROUG
- E. D. Use token count as the primary metric.

Answer: A

Precision measures how many predicted positives are actually positive, which matters when false positives are costly. Why the strongest distractor fails: Recall emphasizes finding positives and can tolerate more false positives.

Question 6 : Domain 1: Fundamentals of AI and ML

[Medium] A data analyst is advising an automotive supplier. The organization needs an AI system that can reason about a goal, use approved tools, and complete a multi-step workflow. The team needs a production-ready recommendation. Which option is the strongest recommendation?

- A. Use an agentic AI approach.
- B. Use a relational database index only.
- C. Use text-to-speech synthesis only.
- D. Use a static classification threshold only.

Answer: A

Agentic AI is relevant when the system must plan, use tools, and orchestrate steps toward a goal. Why the strongest distractor fails: A classifier can predict a label but does not orchestrate a multi-step workflow.

Question 7 : Domain 1: Fundamentals of AI and ML

[Medium] For a regulated internal workflow, a property-management company has support tickets labeled with the correct routing team and wants to learn a routing rule from examples. The reviewers want a clear reason for the recommendation. Which action should the team take next?

- A. Use unsupervised learning.
- B. Use reinforcement learning.
- C. Use supervised learning.
- D. Use text-to-speech.

Answer: C

Supervised learning uses labeled examples to predict a known target. Why the strongest distractor fails: Unsupervised learning is used when target labels are absent.

Question 8 : Domain 1: Fundamentals of AI and ML

[Hard] An insurance company must make a regulated credit-risk prediction and needs a narrowly scoped, explainable output. The reviewers want a clear reason for the recommendation. The reviewers want the answer that best handles the stated tradeoff rather than a merely plausible alternative. Which recommendation is the best fit? Select the option containing the TWO best actions.

- A. 1. Evaluate a traditional ML model before choosing a foundation model. 2. Compare capabilities, latency, cost, and compliance constraints.
- B. 1. Always use the largest foundation model. 2. Skip explainability because the workload is regulated.
- C. 1. Evaluate a traditional ML model before choosing a foundation model. 2. Always use the largest foundation model.
- D. 1. Use an image-generation model. 2. Compare capabilities, latency, cost, and compliance constraints.

Answer: A

Option A contains the complementary actions. Traditional ML can be a better fit for a narrow predictive problem with explainability constraints. The paired distractors include an unsuitable or incomplete action. A foundation model is not automatically the best choice for every AI use case.

Question 9 : Domain 1: Fundamentals of AI and ML

[Hard] For a multi-region business process, a university needs to extract printed text, forms, and table structure from scanned invoices. The recommendation must control cost without sacrificing the stated requirement. The reviewers want the answer that best handles the stated tradeoff rather than a merely plausible alternative. Which answer best addresses the requirement?

- A. Use Amazon Rekognition only.
- B. Use Amazon Textract.**
- C. Use Amazon Personalize.
- D. Use Amazon Transcribe.

Answer: B

Amazon Textract extracts text, forms, and tables from documents. Why the strongest distractor fails: Amazon Rekognition is not the specialized document-extraction service.

Question 10 : Domain 1: Fundamentals of AI and ML

[Hard] A food-delivery platform is evaluating whether an AI assistant creates business value after deployment. The recommendation must control cost without sacrificing the stated requirement. The reviewers want the answer that best handles the stated tradeoff rather than a merely plausible alternative. Which option is the strongest recommendation?

- A. Track business metrics such as ROI, customer feedback, and cost per user.**
- B. Track only the number of subnets.
- C. Track only the number of IAM groups.
- D. Track only model file size.

Answer: A

Business metrics show whether the workload achieves its intended outcome. Why the strongest distractor fails: Infrastructure counts do not establish business value.

Question 11 : Domain 1: Fundamentals of AI and ML

[Easy] During a design review, the security reviewer for a customer-support outsourcer explains that the organization must process large document requests that can take several minutes and return results when processing is complete. The team needs a production-ready recommendation. Which action should the team take next?

- A. Use Amazon Polly.
- B. Use asynchronous inference.**
- C. Use a synchronous low-latency endpoint only.
- D. Use clustering.

Answer: B

Asynchronous inference fits long-running requests where callers can retrieve results later. Why the strongest distractor fails: A synchronous endpoint can time out or tie up resources for long-running jobs.

Question 12 : Domain 1: Fundamentals of AI and ML

[Easy] During a design review, the operations manager for an online retailer explains that the organization needs to predict whether a loan application belongs to an approve or review category. The team wants to avoid an overly complex design. Which recommendation is the best fit?

- A. Use clustering.
- B. Use regression.
- C. Use classification.**
- D. Use text-to-speech.

Answer: C

Classification predicts discrete categories. Why the strongest distractor fails: Regression predicts a numeric value rather than a category.

Question 13 : Domain 1: Fundamentals of AI and ML

[Easy] A governance checkpoint at a travel marketplace confirms that the organization needs to translate product descriptions into multiple languages. The team needs a production-ready recommendation. Which answer best addresses the requirement?

- A. Use Amazon Comprehend.
- B. Use Amazon Lex.
- C. Use Amazon Translate.**
- D. Use Amazon Personalize.

Answer: C

Amazon Translate provides machine translation. Why the strongest distractor fails: Amazon Comprehend analyzes text but is not the managed translation service.

Question 14 : Domain 1: Fundamentals of AI and ML

[Medium] A utilities provider is planning a regulated internal workflow. It has deployed an ML model and wants to detect deteriorating production behavior over time. The recommendation must control cost without sacrificing the stated requirement. Which option is the strongest recommendation?

- A. Measure only the number of IAM users.
- B. Assume the initial evaluation remains valid forever.
- C. Monitor performance and retrain when needed.**
- D. Remove the evaluation dataset.

Answer: C

MLOps includes ongoing monitoring and retraining as data and behavior change. Why the strongest distractor fails: Deployment does not eliminate drift risk.

Question 15 : Domain 1: Fundamentals of AI and ML

[Medium] Before a production rollout, a global wholesaler asks what differentiates generative AI from a traditional predictive model. The reviewers want a clear reason for the recommendation. Which action should the team take next? Select the option containing the TWO best actions.

- A. 1. Generative AI can create new content such as text, images, audio, or code. 2. Generative AI can only assign predefined labels.
- B. 1. Generative AI can create new content such as text, images, audio, or code. 2. Validate the choice against the business requirement and representative examples.**
- C. 1. Generative AI can only assign predefined labels. 2. Generative AI eliminates the need for evaluation.
- D. 1. Generative AI is the same as encryption. 2. Validate the choice against the business requirement and representative examples.

Answer: B

Option B contains the complementary actions. Generative AI produces new content rather than only predicting a label or numeric value. The paired distractors include an unsuitable or incomplete action. Assigning predefined labels describes classification, not the distinguishing capability of GenAI.

Question 16 : Domain 1: Fundamentals of AI and ML

[Medium] As part of a seasonal campaign with unpredictable demand, a healthcare network has historical claims records with an approved fraud or non-fraud outcome for each record. The recommendation must control cost without sacrificing the stated requirement. Which recommendation is the best fit?

- A. Treat the outcomes as encryption keys.
- B. Treat the records as unlabeled because they are historical.
- C. Treat the records as image data.
- D. Treat the records as labeled data.**

Answer: D

The known target outcomes make the dataset labeled. Why the strongest distractor fails: Historical age does not determine whether data is labeled.

Question 17 : Domain 1: Fundamentals of AI and ML

[Medium] For a seasonal campaign with unpredictable demand, a university must calculate a tax amount from a complete deterministic rule table and always return the exact same result for the same inputs. The team wants the simplest suitable option. Which answer best addresses the requirement?

- A. Increase model temperature.
- B. Use a conventional rules-based implementation.**
- C. Use clustering.
- D. Train a generative AI model.

Answer: B

A deterministic rules engine is a better fit when the outcome is fully specified. Why the strongest distractor fails: A generative model introduces variability and unnecessary cost.

Question 18 : Domain 1: Fundamentals of AI and ML

[Medium] For a proof of concept that may later scale, a software vendor needs to detect labels and objects in uploaded product images. The recommendation must control cost without sacrificing the stated requirement. Which option is the strongest recommendation?

- A. Use Amazon Translate.
- B. Use Amazon Textract.
- C. Use Amazon Rekognition.**
- D. Use Amazon Comprehend.

Answer: C

Amazon Rekognition analyzes images and video. Why the strongest distractor fails: Amazon Textract focuses on extracting text and document structure.

Question 19 : Domain 1: Fundamentals of AI and ML

[Hard] As part of a nightly back-office process, a telecom operator evaluates an imbalanced classifier and wants a metric that balances precision and recall. The choice must address the business requirement directly. The reviewers want the answer that best handles the stated tradeoff rather than a merely plausible alternative. Which action should the team take next?

- A. Use token count.
- B. Use only accuracy.
- C. Use BLEU.
- D. Use the F1 score.**

Answer: D

F1 balances precision and recall and is useful when class imbalance makes accuracy misleading. Why the strongest distractor fails: Accuracy can hide poor performance on a rare but important class.

Question 20 : Domain 1: Fundamentals of AI and ML

[Hard] During a design review, the solutions architect for an online retailer explains that the organization must return a fraud-risk prediction while a card payment is being authorized. The team needs a production-ready recommendation. The reviewers want the answer that best handles the stated tradeoff rather than a merely plausible alternative. Which recommendation is the best fit?

- A. Use a weekly batch job.
- B. Use Amazon Translate.

C. Use real-time inference.

D. Use clustering.

Answer: C

Real-time inference is appropriate when an immediate prediction is required for an interactive decision. Why the strongest distractor fails: A batch job cannot meet the authorization latency requirement.

Question 21 : Domain 1: Fundamentals of AI and ML

[Easy] A governance checkpoint at a travel marketplace confirms that the organization is training a system to choose actions and improve its behavior based on rewards and penalties. The solution should remain easy to govern. Which answer best addresses the requirement?

A. Use regression.

B. Use reinforcement learning.

C. Use a static S3 lifecycle rule.

D. Use clustering.

Answer: B

Reinforcement learning learns action policies from reward signals. Why the strongest distractor fails: Clustering discovers groups and does not learn from rewards.

Question 22 : Domain 1: Fundamentals of AI and ML

[Easy] A healthcare network is planning a nightly back-office process. It needs to generate spoken audio from written accessibility content. The reviewers want a clear reason for the recommendation. Which option is the strongest recommendation? Select the option containing the TWO best actions.

A. 1. Use Amazon Transcribe. 2. Use Amazon Textract.

B. 1. Use Amazon Polly. 2. Use Amazon Transcribe.

C. 1. Use Amazon Polly. 2. Verify that the AWS service matches the required data type and outcome.

D. 1. Use Amazon Comprehend. 2. Verify that the AWS service matches the required data type and outcome.

Answer: C

Option C contains the complementary actions. Amazon Polly converts text into lifelike speech. The paired distractors include an unsuitable or incomplete action. Amazon Transcribe converts speech to text.

Question 23 : Domain 1: Fundamentals of AI and ML

[Easy] A governance checkpoint at an automotive supplier confirms that the organization wants to use an existing foundation model without managing model-serving infrastructure. The team wants to avoid an overly complex design. Which action should the team take next?

A. Use S3 Glacier as the inference API.

B. Self-host every model by default.

C. Use a managed API service such as Amazon Bedrock.

D. Build a custom neural network before evaluating managed services.

Answer: C

A managed API reduces infrastructure management when existing model capabilities meet the need. Why the strongest distractor fails: Self-hosting can add operational complexity that the requirement does not justify.

Question 24 : Domain 1: Fundamentals of AI and ML

[Easy] A global wholesaler needs to identify the concept that commonly uses multilayer neural networks to learn representations from large datasets. The recommendation must control cost without sacrificing the stated requirement. Which recommendation is the best fit?

A. A rules engine.

B. Deep learning.

C. Batch inference.

D. Clustering only.

Answer: B

Deep learning commonly uses multilayer neural networks. Why the strongest distractor fails: Batch inference is a serving pattern, not a learning approach.

Question 25 : Domain 1: Fundamentals of AI and ML

[Medium] The next decision for a media company concerns this need: the organization is cataloging relational order records, customer emails, and product photos for an AI project. The team needs a production-ready recommendation. Which answer best addresses the requirement?

A. Treat relational records as unstructured data.

B. Treat relational records as structured data and emails and photos as unstructured data.

C. Treat product photos as labeled data by definition.

D. Treat all three sources as structured data.

Answer: B

Relational data follows a schema; free text and images generally do not. Why the strongest distractor fails: A stored image is not automatically labeled merely because it is a file.

Question 26 : Domain 1: Fundamentals of AI and ML

[Medium] The next decision for a logistics provider concerns this need: the organization wants to group similar news articles without creating topic labels in advance. The team needs a production-ready recommendation. Which option is the strongest recommendation?

- A. Use classification.
- B. Use regression.
- C. Use clustering.**
- D. Use Amazon Polly.

Answer: C

Clustering discovers groups in unlabeled data. Why the strongest distractor fails: Classification needs predefined labels.

Question 27 : Domain 1: Fundamentals of AI and ML

[Medium] An architecture workshop for a logistics provider identifies this requirement: the organization needs a conversational interface that recognizes intent and captures values such as appointment date. The design must minimize unnecessary operational overhead. Which action should the team take next?

- A. Use Amazon Lex.**
- B. Use Amazon Polly only.
- C. Use Amazon Textract.
- D. Use Amazon Macie.

Answer: A

Amazon Lex supports conversational interfaces with intents and slots. Why the strongest distractor fails: Amazon Polly can synthesize speech but does not manage conversational intent.

Question 28 : Domain 1: Fundamentals of AI and ML

[Medium] A governance checkpoint at an automotive supplier confirms that the organization must catch as many truly fraudulent transactions as possible, even if analysts review some false alarms. The team wants to avoid an overly complex design. Which recommendation is the best fit?

- A. Use average prompt length.
- B. Use BLEU.
- C. Prioritize precision only.
- D. Prioritize recall.**

Answer: D

Recall measures how many actual positives are found. Why the strongest distractor fails: Precision alone does not ensure that most true fraud cases are detected.

Question 29 : Domain 1: Fundamentals of AI and ML

[Hard] The application owner at an online retailer is comparing plausible alternatives because the organization can generate inventory forecasts overnight for the next business day. The organization needs an approach that can be evaluated with measurable outcomes. The reviewers want the answer that best handles the stated tradeoff rather than a merely plausible

alternative. Which answer best addresses the requirement? Select the option containing the TWO best actions.

A. 1. Use batch inference. 2. Use real-time inference for each record.

B. 1. Use batch inference. 2. Match the serving pattern to response-time and traffic requirements.

C. 1. Use real-time inference for each record. 2. Use Amazon Polly.

D. 1. Use reinforcement learning. 2. Match the serving pattern to response-time and traffic requirements.

Answer: B

Option B contains the complementary actions. Batch inference is suitable when predictions can be processed together on a schedule. The paired distractors include an unsuitable or incomplete action. Real-time inference adds unnecessary responsiveness and may add cost for an overnight requirement.

Question 30 : Domain 1: Fundamentals of AI and ML

[Hard] Matching-style question: During an exam-readiness review, a payment processor must verify three closely related concepts. Which option correctly matches every item to its primary purpose?

A. classification = predict a numeric value; regression = predict a category; clustering = discover groups without labels

B. classification = discover groups without labels; regression = predict a numeric value; clustering = predict a category

C. classification = predict a category; regression = discover groups without labels; clustering = predict a numeric value

D. classification = predict a category; regression = predict a numeric value; clustering = discover groups without labels

Answer: D

Option D correctly maps all three items. The distractors swap closely related concepts or services, which is a common exam trap.

Question 31 : Domain 1: Fundamentals of AI and ML

[Easy] For a latency-sensitive interactive service, a software vendor needs to convert recorded customer calls into written transcripts. The team wants the simplest suitable option. Which action should the team take next?

A. Use Amazon Polly.

B. Use Amazon Transcribe.

C. Use Amazon Translate.

D. Use Amazon Rekognition.

Answer: B

Amazon Transcribe converts speech to text. Why the strongest distractor fails: Amazon Polly performs the opposite direction: text to speech.

Question 32 : Domain 1: Fundamentals of AI and ML

[Easy] Ordering-style question: The product manager for a nonprofit organization needs to document a safe implementation sequence. Which option puts the activities in the most appropriate order?

- A. Collect and prepare data -> monitor and retrain when needed -> deploy -> evaluate -> train
- B. train -> evaluate -> deploy -> monitor and retrain when needed -> Collect and prepare data
- C. monitor and retrain when needed -> deploy -> evaluate -> train -> Collect and prepare data
- D. Collect and prepare data -> train -> evaluate -> deploy -> monitor and retrain when needed**

Answer: D

Option D uses a sensible sequence. The distractors move evaluation, controls, or deployment steps into an unsafe or ineffective order.

Question 33 : Domain 1: Fundamentals of AI and ML

[Easy] A governance checkpoint at a payment processor confirms that the organization needs to identify the concept that commonly uses multilayer neural networks to learn representations from large datasets. The organization needs an approach that can be evaluated with measurable outcomes. Which answer best addresses the requirement?

- A. A rules engine.
- B. Batch inference.
- C. Clustering only.
- D. Deep learning.**

Answer: D

Deep learning commonly uses multilayer neural networks. Why the strongest distractor fails: Batch inference is a serving pattern, not a learning approach.

Question 34 : Domain 1: Fundamentals of AI and ML

[Easy] As part of a nightly back-office process, a utilities provider is cataloging relational order records, customer emails, and product photos for an AI project. The recommendation must control cost without sacrificing the stated requirement. Which option is the strongest recommendation?

- A. Treat relational records as structured data and emails and photos as unstructured data.**
- B. Treat product photos as labeled data by definition.
- C. Treat relational records as unstructured data.
- D. Treat all three sources as structured data.

Answer: A

Relational data follows a schema; free text and images generally do not. Why the strongest distractor fails: A stored image is not automatically labeled merely because it is a file.

Question 35 : Domain 1: Fundamentals of AI and ML

[Medium] As part of a regulated internal workflow, a telecom operator wants to group similar news articles without creating topic labels in advance. The recommendation must control cost without sacrificing the stated requirement. Which action should the team take next?

- A. Use clustering.
- B. Use Amazon Polly.
- C. Use regression.
- D. Use classification.

Answer: A

Clustering discovers groups in unlabeled data. Why the strongest distractor fails: Classification needs predefined labels.

Unlock All 1100+ Questions

Get the complete Q&A package with detailed explanations, topic analytics, and exam-accurate practice.

From €29.00

Visit: <https://cert-pass.com/exams/aws-aws-certified-ai-practitioner-aif-c01>

CERT-PASS

cert-pass.com

© 2026 Cert-Pass. This material is for personal use only. Do not distribute.