

# 2026 AIP-C01–100% Free Examinations Actual Questions | Reliable AWS Certified Generative AI Developer - Professional Reliable Exam Question

PAGE 1 OF 4

MARKS ALONE WON'T GET YOU SUCCESS IN THE REAL WORLD... **PRACTICAL SKILLS WILL!**

Companies & Firms don't pay for what you studied... They pay for what you can DO practically.

EVERY YEAR, THOUSANDS OF CA STUDENTS FACE THE SAME PROBLEM:

- Knows concepts... but don't know **practical work** (GST / ITR / TDS)
- Cleared exams... but struggle in **real-time office work**
- Articleship completed... but still **not confident** to handle work independently
- Exposure limited to **only few areas** in articleship or small firms
- Unable to answer **practical questions** in interviews
- Gap in articleship or **not up-to-date** with latest compliances

COMPLETE PRACTICAL COVERAGE OF REAL CA OFFICE WORK

- GST RETURNS:** Registration, GSTR-1, 3B, 9, Reconciliation & Real time filing
- INCOME TAX RETURNS:** ITR 1 to ITR 7, Salary, Business, Deductions, E-filing & more
- TDS RETURNS:** 24Q, 26Q, 27Q, Challan & Returns Filing, Correction
- TALLY PRIME:** Accounting, Inventory, GST Integration, Reports & MIS
- ADVANCED EXCEL:** Formulas, Pivot Tables, Dashboards, MIS & Data Analysis
- TAX AUDIT:** u/s 44AB, 3CD, Audit Report, Disclosure Req.
- B/S FINALISATION:** Trial Balance, Adjustments, Final Accounts, Schedules & Notes
- PF & ESI FILINGS:** Registration, Monthly Returns, Challan & Compliance
- MIS REPORTS:** MIS Preparation, Variance Analysis, Management Reports
- PROJECT REPORTS:** Financial Projections, Ratio Analysis, Presentation & more
- FIRM REGISTRATION:** Proprietorship, Partnership, LLP, FSSAI, MSME, GST Registration
- OTHER PRACTICALS:** Bank Reconciliation, E-way Bill, ROC Filings, Practical Assignments & more

In Interviews, They Don't Ask What You Studied... They Ask What You Can DO Practically.

At **CAPITAL TRAINERS HYDERABAD**, you learn exactly what companies & firms expect from Day 1.

- Practical Exposure
- Real Time Work
- Expert Faculty
- Career Support

Next Page: Find your situation & see how we make you job ready.

**CAPITAL TRAINERS HYDERABAD**  
India's #1 Practical Taxation Training Institute  
Learn Practically. Work Confidently.

- PRACTICAL TRAINING
- EXPERT FACULTIES
- JOB READY SKILLS
- 100% PRACTICAL EXPOSURE

BONUS!!! Download part of Exams4sures AIP-C01 dumps for free: <https://drive.google.com/open?id=1t0em4FEIqXxVvmlZxbm093W924xy1mFRx>

We offer you free demo for AIP-C01 pdf dumps. You can check out the questions quality and usability of our training material before you buy. Amazon AIP-C01 questions are written to the highest standards of technical accuracy with accurate answers. If you prepare for your exams using Exams4sures AIP-C01 practice torrent, it is easy to succeed for your certification in the first attempt. Besides, we offer the money refund policy, in case of failure, you can ask for full refund.

Our company guarantees this pass rate from various aspects such as content and service on our AIP-C01 exam questions. We have hired the most authoritative professionals to compile the content Of the AIP-C01 study materials. And we offer 24/7 service online to help you on all kinds of the problems about the AIP-C01 learning guide. Of course, we also consider the needs of users, our AIP-C01 exam questions hope to help every user realize their dreams.

>> Examinations AIP-C01 Actual Questions <<

**Latest AIP-C01 – 100% Free Examinations Actual Questions | AIP-C01 Reliable Exam Question**

Our company provides three different versions to choose for our customers. The software version of our AIP-C01 exam question has a special function that this version can simulate test-taking conditions for customers. If you feel very nervous about exam, we think it is very necessary for you to use the software version of our AIP-C01 Guide Torrent. By simulating actual test-taking conditions, we believe that you will relieve your nervousness before examination. So hurry to buy our AIP-C01 test questions, it will be very helpful for you to pass your AIP-C01 exam and get your certification.

## Amazon AWS Certified Generative AI Developer - Professional Sample Questions (Q33-Q38):

### NEW QUESTION # 33

A company uses Amazon Bedrock to implement a Retrieval Augmented Generation (RAG)-based system to serve medical information to users. The company needs to compare multiple chunking strategies, evaluate the generation quality of two foundation models (FMs), and enforce quality thresholds for deployment.

Which Amazon Bedrock evaluation configuration will meet these requirements?

- A. Set up a pipeline that uses multiple retrieve-only evaluation jobs to assess retrieval quality. Create separate evaluation jobs for both FMs that use Amazon Nova Pro as the LLM-as-a-judge model. Evaluate based on faithfulness and citation precision metrics.
- B. Create a retrieve-only evaluation job that uses a supported version of Anthropic Claude Sonnet as the evaluator model. Configure metrics for context relevance and context coverage. Define deployment thresholds in a separate CI/CD pipeline.
- C. Create a retrieve-and-generate evaluation job that uses custom precision-at-k metrics and an LLM-as-a-judge metric with a scale of 1-5. Include each chunking strategy in the evaluation dataset. Use a supported version of Anthropic Claude Sonnet to evaluate responses from both FMs.
- D. Create a separate evaluation job for each chunking strategy and FM combination. Use Amazon Bedrock built-in metrics for correctness and completeness. Manually review scores before deployment approval.

**Answer: C**

Explanation:

Option B is the correct evaluation configuration because it enables end-to-end assessment of both retrieval and generation quality while supporting direct comparison of chunking strategies and foundation models.

Amazon Bedrock evaluation jobs are designed to support RAG workflows by evaluating how well retrieved context supports accurate and high-quality model outputs.

A retrieve-and-generate evaluation job evaluates the complete RAG pipeline, not just retrieval. This is essential for medical information use cases, where both the relevance of retrieved content and the correctness of generated responses directly impact user safety and trust. Including multiple chunking strategies in the evaluation dataset allows side-by-side comparison under identical prompts and conditions.

Custom precision-at-k metrics measure how effectively the retrieval component surfaces relevant chunks, while an LLM-as-a-judge metric provides qualitative scoring of generated responses. Using a numeric scale enables consistent, repeatable evaluation and supports automated quality gates. Amazon Bedrock supports LLM-based evaluators to score dimensions such as accuracy, completeness, and relevance.

Using the same evaluator model to assess outputs from both FMs ensures consistent scoring and eliminates evaluator bias. This configuration allows the company to define quantitative thresholds that must be met before deployment, enabling automated promotion through CI/CD pipelines.

Option A evaluates retrieval only and cannot assess generation quality. Option C introduces manual review, which does not scale and delays deployment. Option D separates retrieval and generation evaluation, making it harder to correlate chunking strategies with final output quality.

Therefore, Option B best meets the requirements for systematic evaluation, comparison, and quality enforcement in an Amazon Bedrock-based RAG system.

### NEW QUESTION # 34

An elevator service company has developed an AI assistant application by using Amazon Bedrock. The application generates elevator maintenance recommendations to support the company's elevator technicians.

The company uses Amazon Kinesis Data Streams to collect the elevator sensor data.

New regulatory rules require that a human technician must review all AI-generated recommendations. The company needs to establish human oversight workflows to review and approve AI recommendations. The company must store all human technician review decisions for audit purposes.

Which solution will meet these requirements?

- A. Configure Amazon EventBridge rules with custom event patterns to route AI recommendations to human technicians for review. Create AWS Glue jobs to process human technician approval queues. Use Amazon ElastiCache to cache all human technician review decisions.
- B. Create an AWS Glue workflow that has a human approval step. After the human technician review, integrate the application with an AWS Lambda function that calls the SendTaskSuccess API. Store all human technician review decisions in Amazon DynamoDB.
- C. Create a custom approval workflow by using AWS Lambda functions and Amazon SQS queues for human review of AI recommendations. Store all review decisions in Amazon DynamoDB for audit purposes.
- D. Create an AWS Step Functions workflow that has a human approval step that uses the waitForResource API to pause execution. After a human technician completes a review, use an AWS Lambda function to call the SendTaskSuccess API with the approval decision. Store all review decisions in Amazon DynamoDB.

**Answer: D**

Explanation:

AWS Step Functions provides native support for human-in-the-loop workflows, making it the best fit for regulatory oversight requirements. The waitForResource integration pattern is explicitly designed to pause a workflow until an external actor—such as a human reviewer—completes a task.

In this architecture, AI-generated recommendations are sent to a human technician for review. The workflow pauses execution using a task token. Once the technician approves or rejects the recommendation, an AWS Lambda function calls SendTaskSuccess or SendTaskFailure, allowing the workflow to continue deterministically.

This approach ensures full auditability, as Step Functions records every state transition, timestamp, and execution path. Storing review outcomes in Amazon DynamoDB provides durable, queryable audit records required for regulatory compliance.

Option A requires custom orchestration and lacks native workflow state management. Option C incorrectly uses AWS Glue, which is not designed for approval workflows. Option D uses caching instead of durable audit storage and introduces unnecessary complexity.

Therefore, Option B is the AWS-recommended, lowest-risk, and most auditable solution for mandatory human review of AI outputs.

### NEW QUESTION # 35

A company is developing a customer communication platform that uses an AI assistant powered by an Amazon Bedrock foundation model (FM). The AI assistant summarizes customer messages and generates initial response drafts.

The company wants to use Amazon Comprehend to implement layered content filtering. The layered content filtering must prevent sharing of offensive content, protect customer privacy, and detect potential inappropriate advice solicitation. Inappropriate advice solicitation includes requests for unethical practices, harmful activities, or manipulative behaviors.

The solution must maintain acceptable overall response times, so all pre-processing filters must finish before the content reaches the FM.

Which solution will meet these requirements?

- A. Use toxicity detection with thresholds configured to 0.5 for all categories. Use parallel processing for both prompt safety classification and personally identifiable information (PII) detection with entity redaction. Apply Amazon CloudWatch alarms to filter metrics.
- B. Deploy a multi-stage process. Configure the process to use prompt safety classification first, then toxicity detection on safe prompts only, and finally personally identifiable information (PII) detection in streaming mode. Route flagged messages through Amazon EventBridge for human review.
- C. Use parallel processing with asynchronous API calls. Use toxicity detection for offensive content. Use prompt safety classification for inappropriate advice solicitation. Use personally identifiable information (PII) detection without redaction.
- D. Use custom classification to build an FM that detects offensive content and inappropriate advice solicitation. Apply personally identifiable information (PII) detection as a secondary filter only when messages pass the custom classifier.

**Answer: A**

Explanation:

Option D best satisfies all functional, performance, and governance requirements while minimizing architectural complexity. The requirement explicitly states that all filtering must complete before content reaches the foundation model, which rules out asynchronous or streaming-based approaches that could delay enforcement.

Amazon Comprehend supports toxicity detection, prompt safety classification, and PII detection with entity redaction as managed capabilities. Running these filters in parallel ensures low end-to-end latency, which is essential for customer-facing communication platforms. Parallel execution avoids the cumulative latency that would be introduced by sequential pipelines.

Toxicity detection identifies offensive or abusive content early. Prompt safety classification detects requests for unethical, harmful, or

manipulative advice, which directly addresses inappropriate advice solicitation requirements. PII detection with entity redaction ensures that customer privacy is preserved before data is sent to the FM, preventing sensitive information from being processed or echoed in generated responses.

Configuring thresholds allows fine-grained control over sensitivity while maintaining acceptable false-positive rates. Using CloudWatch metrics and alarms enables continuous monitoring of filtering behavior and intervention rates without adding custom routing or human review pipelines that would slow responses.

Option A lacks PII redaction. Option B introduces unnecessary model-building complexity and delayed PII checks. Option C adds sequential latency and introduces human review routing, which violates the response-time requirement.

Therefore, Option D provides the most robust, performant, and AWS-aligned layered content filtering solution.

### NEW QUESTION # 36

A company is implementing a serverless inference API by using AWS Lambda. The API will dynamically invoke multiple AI models hosted on Amazon Bedrock. The company needs to design a solution that can switch between model providers without modifying or redeploying Lambda code in real time. The design must include safe rollout of configuration changes and validation and rollback capabilities.

Which solution will meet these requirements?

- A. Configure an Amazon API Gateway REST API to route requests to separate Lambda functions. Hardcode each Lambda function to a specific model provider. Switch the integration target manually.
- B. Store the active model provider in AWS Systems Manager Parameter Store. Configure a Lambda function to read the parameter at runtime to determine which model to invoke.
- C. Store the active model provider in AWS AppConfig. Configure a Lambda function to read the configuration at runtime to determine which model to invoke.
- D. Store the active model provider in a JSON file hosted on Amazon S3. Use AWS AppConfig to reference the S3 file as a hosted configuration source. Configure a Lambda function to read the file through AppConfig at runtime to determine which model to invoke.

**Answer: C**

Explanation:

Option B is the correct solution because AWS AppConfig is specifically designed to support dynamic configuration management with safe rollout, validation, and rollback, which are explicit requirements in the scenario.

By storing the active model provider configuration in AWS AppConfig, the company can switch between Amazon Bedrock model providers in real time without redeploying Lambda code. AppConfig supports deployment strategies such as canary releases, linear rollouts, and immediate deployments, allowing safe and controlled changes. If a configuration causes issues, AppConfig supports automatic rollback, reducing operational risk.

AWS AppConfig also supports schema validation, ensuring that configuration values such as model identifiers, provider names, or inference parameters are valid before being applied. This prevents misconfiguration from impacting production workloads.

Option A uses Parameter Store, which lacks native rollout strategies, validation, and automated rollback, making it unsuitable for safe real-time switching. Option C requires manual routing changes and code coupling, increasing operational overhead and deployment risk. Option D introduces unnecessary complexity by hosting configuration files in Amazon S3 when AppConfig already supports native hosted configurations.

Therefore, Option B provides the most robust, scalable, and low-maintenance solution for dynamic model switching in a serverless Amazon Bedrock inference architecture.

### NEW QUESTION # 37

A company uses an AI assistant application to summarize the company's website content and provide information to customers. The company plans to use Amazon Bedrock to give the application access to a foundation model (FM).

The company needs to deploy the AI assistant application to a development environment and a production environment. The solution must integrate the environments with the FM. The company wants to test the effectiveness of various FMs in each environment. The solution must provide product owners with the ability to easily switch between FMs for testing purposes in each environment.

Which solution will meet these requirements?

- A. Create one AWS CDK application for the production environment. Configure the application to invoke the Amazon Bedrock FMs by using the `aws_bedrock.ProvisionedModel.fromProvisionedModelArn()` method. Create a pipeline in AWS CodePipeline. Configure the pipeline to deploy to the production environment by using an AWS CodeBuild deploy action. For the development environment, manually recreate the resources by referring to the production application code.

- B. Create one AWS CDK application. Create multiple pipelines in AWS CodePipeline. Configure each pipeline to have its own settings for each FM. Configure the application to invoke the Amazon Bedrock FMs by using the `aws_bedrock.ProvisionedModel.fromProvisionedModelArn()` method.
- C. Create one AWS CDK application. Configure the application to invoke the Amazon Bedrock FMs by using the `aws_bedrock.FoundationModel.fromFoundationModelId()` method. Create a pipeline in AWS CodePipeline that has a deployment stage for each environment that uses AWS CodeBuild deploy actions.
- D. Create a separate AWS CDK application for each environment. Configure the applications to invoke the Amazon Bedrock FMs by using the `aws_bedrock.FoundationModel.fromFoundationModelId()` method. Create a separate pipeline in AWS CodePipeline for each environment.

**Answer: C**

Explanation:

Option C best satisfies the requirement for flexible FM testing across environments while minimizing operational complexity and aligning with AWS-recommended deployment practices. Amazon Bedrock supports invoking on-demand foundation models through the `FoundationModel` abstraction, which allows applications to dynamically reference different models without requiring dedicated provisioned capacity. This is ideal for experimentation and A/B testing in both development and production environments. Using a single AWS CDK application ensures infrastructure consistency and reduces duplication.

Environment-specific configuration, such as selecting different foundation model IDs, can be externalized through parameters, context variables, or environment-specific configuration files. This allows product owners to easily switch between FMs in each environment without modifying application logic.

A single AWS CodePipeline with distinct deployment stages for development and production is an AWS best practice for multi-environment deployments. It enforces consistent build and deployment steps while still allowing environment-level customization. AWS CodeBuild deploy actions enable automated, repeatable deployments, reducing manual errors and improving governance. Option A increases complexity by introducing multiple pipelines and relies on provisioned models, which are not necessary for FM evaluation and experimentation. Provisioned throughput is better suited for predictable, high-volume production workloads rather than frequent model switching.

Option B creates unnecessary operational overhead by duplicating CDK applications and pipelines, making long-term maintenance more difficult.

Option D directly conflicts with infrastructure-as-code best practices by manually recreating development resources, which increases configuration drift and reduces reliability.

Therefore, Option C provides the most flexible, scalable, and AWS-aligned solution for testing and switching foundation models across development and production environments.

## NEW QUESTION # 38

.....

Success in the AWS Certified Generative AI Developer - Professional (AIP-C01) certification exam helps people update their skills. Many aspirants don't find updated Amazon AIP-C01 practice test questions and fail the final test. This failure in the Amazon AIP-C01 Exam leads to a loss of money and time. If you are also planning to attempt the AWS Certified Generative AI Developer - Professional (AIP-C01) exam and are confused about where to prepare yourself for it then you are at the right place.

**AIP-C01 Reliable Exam Question:** <https://www.exams4sures.com/Amazon/AIP-C01-practice-exam-dumps.html>

Our experts strive hard provide you with the latest AIP-C01 exam questions, Amazon Examinations AIP-C01 Actual Questions And you are capable for your job, The Exams4sures AIP-C01 Reliable Exam Question created exceptional approaches of study, Exams4sures AIP-C01 PDF files prepared specially for the 100% success in Amazon certifications, Amazon offers a free demo version for you to verify the authenticity of the Amazon AIP-C01 exam prep material before buying it.

Perl for the Web, Tired of losing money in the market, Our experts strive hard provide you with the Latest AIP-C01 Exam Questions, And you are capable for your job.

The Exams4sures created exceptional approaches of study, Exams4sures AIP-C01 PDF files prepared specially for the 100% success in Amazon certifications, Amazon offers a free demo version for you to verify the authenticity of the Amazon AIP-C01 exam prep material before buying it.

## 100% Pass Amazon - Trustable AIP-C01 - Examinations AWS Certified Generative AI Developer - Professional Actual Questions

- 2026 Pass-Sure 100% Free AIP-C01 – 100% Free Examinations Actual Questions | AWS Certified Generative AI

Developer - Professional Reliable Exam Question ☐ Search for ➤ AIP-C01 ☐ and download exam materials for free through ⇒ [www.verifieddumps.com](http://www.verifieddumps.com) ⇐ ☐ AIP-C01 Latest Exam Labs

- 100% Pass Quiz AIP-C01 - Valid Examinations AWS Certified Generative AI Developer - Professional Actual Questions ☐ ☐ Open “[www.pdfvce.com](http://www.pdfvce.com)” and search for [ AIP-C01 ] to download exam materials for free ☐ Reliable AIP-C01 Test Book
- Free PDF Quiz Useful Amazon - AIP-C01 - Examinations AWS Certified Generative AI Developer - Professional Actual Questions ☐ Search for 【 AIP-C01 】 and download exam materials for free through ✓ [www.prepawayexam.com](http://www.prepawayexam.com) ☐ ✓ ☐ AIP-C01 Preparation
- AIP-C01 Exam Dumps Pdf!! AIP-C01 Exam Labs ☐ AIP-C01 Online Test ☐ Search for 《 AIP-C01 》 and obtain a free download on ☐ [www.pdfvce.com](http://www.pdfvce.com) ☐ ☐ Dumps AIP-C01 Questions
- Exact Inside Examinations AIP-C01 Actual Questions Questions and Answers ☐ Search for ☐ AIP-C01 ☐ and download it for free on ▶ [www.troytecdumps.com](http://www.troytecdumps.com) ◀ website ☐ AIP-C01 Exam Labs
- AIP-C01 Preparation ☐ Latest AIP-C01 Exam Simulator ☐ Latest AIP-C01 Exam Simulator ☐ Download { AIP-C01 } for free by simply searching on 「 [www.pdfvce.com](http://www.pdfvce.com) 」 ☐ Best AIP-C01 Preparation Materials
- Pass Guaranteed 2026 Amazon - AIP-C01 - Examinations AWS Certified Generative AI Developer - Professional Actual Questions ☐ Download { AIP-C01 } for free by simply entering ➡ [www.dumpsmaterials.com](http://www.dumpsmaterials.com) ☐ website ☐ AIP-C01 Preparation
- Valid AIP-C01 Practice Materials ☐ Reliable AIP-C01 Braindumps Files ☐ Reliable AIP-C01 Braindumps Files ☐ Search for ➡ AIP-C01 ☐ ☐ ☐ and obtain a free download on [ [www.pdfvce.com](http://www.pdfvce.com) ] ☐ Best AIP-C01 Preparation Materials
- Exact Inside Examinations AIP-C01 Actual Questions Questions and Answers ☐ Open ✨ [www.examdiscuss.com](http://www.examdiscuss.com) ☐ ✨ ☐ and search for ☐ AIP-C01 ☐ to download exam materials for free ☐ Dumps AIP-C01 Questions
- Latest AIP-C01 Exam Simulator ☐ Valid Real AIP-C01 Exam ☐ AIP-C01 Exam Question ☐ Search for 【 AIP-C01 】 and obtain a free download on 【 [www.pdfvce.com](http://www.pdfvce.com) 】 ☐ Reliable AIP-C01 Braindumps Files
- 100% Pass Quiz AIP-C01 - Valid Examinations AWS Certified Generative AI Developer - Professional Actual Questions ☐ ☐ Easily obtain free download of ➤ AIP-C01 ☐ by searching on ⇒ [www.vceengine.com](http://www.vceengine.com) ⇐ ☐ New AIP-C01 Exam Guide
- [socialbuzzmaster.com](http://socialbuzzmaster.com), [finnianqbib160186.illawiki.com](http://finnianqbib160186.illawiki.com), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [lilliosoz373003.bloggadores.com](http://lilliosoz373003.bloggadores.com), [directoryarmy.com](http://directoryarmy.com), [nettiegsbg494311.activablog.com](http://nettiegsbg494311.activablog.com), [bookmarkyourpage.com](http://bookmarkyourpage.com), [nikolaskozi222329.nizarblog.com](http://nikolaskozi222329.nizarblog.com), [wiishlist.com](http://wiishlist.com), [divisionmidway.org](http://divisionmidway.org), Disposable vapes

P.S. Free 2026 Amazon AIP-C01 dumps are available on Google Drive shared by Exams4sures: <https://drive.google.com/open?id=1t0em4FEIQxXVmlZxbmf93W924xy1mFRx>