



**TRUSTED, ACCURATE AND
RELIABLE!**

**The most comprehensive IT certification
preparation materials in the industry!**

All rights reserved. No part of this document may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other non-commercial uses permitted by copyright law. Unauthorized copying, reselling, or distribution of this document is strictly prohibited and may result in legal action.

<https://www.virtulearner.com>
support@virtulearner.com

Amazon

AIP-C01

AWS Certified Generative AI

Developer - Professional

Exam

QUESTION: 1

A retail company has a generative AI (GenAI) product recommendation application that uses Amazon Bedrock. The application suggests products to customers based on browsing history and demographics. The company needs to implement fairness evaluation across multiple demographic groups to detect and measure bias in recommendations between two prompt approaches. The company wants to collect and monitor fairness metrics in real time. The company must receive an alert if the fairness metrics show a discrepancy of more than 15% between demographic groups. The company must receive weekly reports that compare the performance of the two prompt approaches.

Which solution will meet these requirements with the LEAST custom development effort?

- A. Configure an Amazon CloudWatch dashboard to display default metrics from Amazon Bedrock API calls.
Create custom metrics based on model outputs. Set up Amazon EventBridge rules to invoke AWS lambda functions that perform post-processing analysis on model responses and publish custom fairness metrics.
- B. Create the two prompt variants in Amazon Bedrock Prompt Management. Use Amazon Bedrock Flows to deploy the prompt variants with defined traffic allocation. Configure Amazon Bedrock guardrails that have content filters to monitor demographic fairness. Set up Amazon CloudWatch alarms on the GuardrailContentSource dimension that use InvocationsIntervened metrics to detect recommendation discrepancy threshold violations.
- C. Set up Amazon SageMaker Clarify to analyze model outputs. Publish fairness metrics to Amazon CloudWatch. Create CloudWatch composite alarms that combine SageMaker Clarify bias metrics with Amazon Bedrock latency metrics to provide a comprehensive fairness evaluation dashboard.
- D. Create an Amazon Bedrock model evaluation job to compare fairness between the two prompt variants.
Enable model invocation logging in Amazon CloudWatch. Set up CloudWatch alarms for InvocationsIntervened metrics with a dimension for each demographic group.

Answer(s): C

Explanation:

Amazon SageMaker Clarify provides built-in bias and fairness evaluation across demographic groups without requiring you to build custom scoring logic. Clarify can compute and publish fairness metrics to Amazon CloudWatch for near-real-time monitoring, where CloudWatch alarms can alert when group-to-group metric deltas exceed the 15% threshold. Clarify also produces periodic bias analysis outputs that can be used to generate weekly comparative reporting for the two prompt approaches with minimal additional implementation.

QUESTION: 2

A finance company is developing an AI assistant to help clients plan investments and manage their portfolios. The company identifies several high-risk conversation patterns such as requests for specific stock recommendations or guaranteed returns. High-risk conversation patterns could lead to regulatory violations if the company cannot implement appropriate controls.

The company must ensure that the AI assistant does not provide inappropriate financial advice, generate content about competitors, or make claims that are not factually grounded in the

company's approved financial guidance. The company wants to use Amazon Bedrock Guardrails to implement a solution.

Which combination of steps will meet these requirements? (Choose three.)

- A. Add the high-risk conversation patterns to a denied topics guardrail.
- B. Configure a content filter guardrail to filter prompts that contain the high-risk conversation patterns.
- C. Configure a content filter guardrail to filter prompts that contain competitor names.
- D. Add the names of competitors as custom word filters. Set the input and output actions to block.
- E. Set a low grounding score threshold.
- F. Set a high grounding score threshold.

Answer(s): A, D, F

Explanation:

Adding high-risk financial requests as denied topics ensures the assistant blocks conversations that could result in regulatory violations or inappropriate advice. Custom word filters with competitor names and block actions prevent the model from generating or responding with competitor-related content. Setting a high grounding score threshold forces responses to stay closely aligned with approved, trusted financial guidance, reducing the risk of unsupported or non-factual claims.

QUESTION: 3

A company has deployed an AI assistant as a React application that uses AWS Amplify, an AWS AppSync GraphQL API, and Amazon Bedrock Knowledge Bases. The application uses the GraphQL API to call the Amazon Bedrock RetrieveAndGenerate API for knowledge base interactions. The company configures an AWS Lambda resolver to use the RequestResponse invocation type.

Application users report frequent timeouts and slow response times. Users report these problems more frequently for complex questions that require longer processing.

The company needs a solution to fix these performance issues and enhance the user experience.

Which solution will meet these requirements?

- A. Use AWS Amplify AI Kit to implement streaming responses from the GraphQL API and to optimize client- side rendering.
- B. Increase the timeout value of the Lambda resolver. Implement retry logic with exponential backoff.
- C. Update the application to send an API request to an Amazon SQS queue. Update the AWS AppSync resolver to poll and process the queue.
- D. Change the RetrieveAndGenerate API to the InvokeModelWithResponseStream API. Update the application to use an Amazon API Gateway WebSocket API to support the streaming response.