DBCLOUD
EXAMS

QUESTIONS & ANSWERS

verified



Technology Provider: Amazon

Exam: AWS Certified Generative AI Developer - Professional AIP-C01

Number of questions in the database: 129

Exam Version: NOV.2026

Topic 1

Question #1Topic 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.

Correct Answer: C

https://dbcloud.academy/

Question #2Topic 1

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.

Correct Answer: ADF

https://dbcloud.academy/

Question #3Topic 1

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.

Correct Answer: A

https://dbcloud.academy/

Question #4Topic 1

An ecommerce company operates a global product recommendation system that needs to switch between multiple foundation models (FM) in Amazon Bedrock based on regulations, cost optimization, and performance requirements. The company must apply custom controls based on proprietary business logic, including dynamic cost thresholds, AWS Region-specific compliance rules, and real-time A/B testing across multiple FMs. The system must be able to switch between FMs without deploying new code. The system must route user requests based on complex rules including user tier, transaction value, regulatory zone, and real-time cost metrics that change hourly and require immediate propagation across thousands of concurrent requests.

Which solution will meet these requirements?

A. Deploy an AWS Lambda function that uses environment variables to store routing rules and Amazon Bedrock FM IDs. Use the Lambda console to update the environment variables when business requirements change. Configure an Amazon API Gateway REST API to read request parameters to make routing decisions.

B. Deploy Amazon API Gateway REST API request transformation templates to implement routing logic based on request attributes. Store Amazon Bedrock FM endpoints as REST API stage variables. Update the variables when the system switches between models.

C. Configure an AWS Lambda function to fetch routing configurations from the AWS AppConfig Agent for each user request. Run business logic in the Lambda function to select the appropriate FM for each request. Expose the FM through a single Amazon API Gateway REST API endpoint.

D. Use AWS Lambda authorizers for an Amazon API Gateway REST API to evaluate routing rules that are stored in AWS AppConfig. Return authorization contexts based on business logic. Route requests to model-specific Lambda functions for each Amazon Bedrock FM.

Correct Answer: C

https://dbcloud.academy/

Question #5Topic 1

A company is developing an internal generative AI (GenAI) assistant that uses Amazon Bedrock to summarize corporate documents for multiple business units. The GenAI assistant must generate responses in a consistent format that includes a document summary, classification of business risks, and terms that are flagged for review. The GenAI assistant must adapt the tone of responses for each user's business unit, such as legal, human resources, or finance. The GenAI assistant must block hate speech, inappropriate topics, and sensitive information such as personal health information.

The company needs a solution to centrally manage prompt variants across business units and teams. The company wants to minimize ongoing orchestration efforts and maintenance for post-processing logic. The company also wants to have the ability to adjust content moderation criteria for the GenAI assistant over time.

Which solution will meet these requirements with the LEAST maintenance overhead?

A. Use Amazon Bedrock Prompt Management to configure reusable templates and business unit-specific prompt variants. Apply Amazon Bedrock guardrails that have category filters and sensitive term lists to block prohibited content.

B. Use Amazon Bedrock Prompt Management to define base templates. Enforce business unit-specific tone by using system prompt variables. Configure Amazon Bedrock guardrails to apply audience-based threshold tuning. Manage the guardrails by using an internal administration API.

C. Use Amazon Bedrock with business unit-based instruction injection in API calls. Store response formatting rules in Amazon DynamoDB. Use AWS Step functions to validate responses. Use Amazon Comprehend to apply content filters after the GenAI assistant generates responses.

D. Use Amazon Bedrock with custom prompt templates that are stored in Amazon DynamoDB. Create one AWS Lambda function to select business unit-specific prompts. Create a second Lambda function to call Amazon Comprehend to filter prohibited content from responses.

Correct Answer: A

https://dbcloud.academy/

Question #6Topic 1

A financial services company is building a customer support application that retrieves relevant financial regulation documents from a database based on semantic similarities to user queries. The application must integrate with Amazon Bedrock to generate responses. The application must be able to search documents that are in English, Spanish, and Portuguese. The application must filter documents by metadata such as publication date, regulatory agency, and document type.

The database stores approximately 10 million document embeddings. To minimize operational overhead, the company wants a solution that minimizes management and maintenance effort. The application must provide low-latency responses for real-time customer interactions.

Which solution will meet these requirements?

A. Use Amazon OpenSearch Serverless to provide vector search capabilities and metadata filtering. Connect to Amazon Bedrock Knowledge Bases to enable Retrieval Augmented Generation (RAG) capabilities that use an Anthropic Claude foundation model (FM).

B. Deploy an Amazon Aurora PostgreSQL database with the pgvector extension. Define tables to store embeddings and metadata. Use SQL queries to perform similarity searches. Send retrieved documents to Amazon Bedrock to generate responses.

C. Use Amazon S3 Vectors to configure a vector index and non-filterable metadata fields. Integrate S3 Vectors with Amazon Bedrock to enable Retrieval Augmented Generation (RAG) capabilities.

D. Set up an Amazon Neptune Analytics graph database. Configure a vector index that has appropriate dimensionality to store document embeddings. Use Amazon Bedrock to perform graph-based retrieval and to generate responses.

Correct Answer: A

https://dbcloud.academy/

Question #7Topic 1

A medical company is building a generative AI (GenAI) application that uses RAG to provide evidence-based medical information. The application uses Amazon OpenSearch Service to retrieve vector embeddings. Users report that searches frequently miss results that contain exact medical terms and acronyms and return too many semantically similar but irrelevant documents. The company needs to improve retrieval quality and maintain low end user latency, even as the document collection grows to millions of documents.

Which solution will meet these requirements with the LEAST operational overhead?

A. Configure hybrid search by combining vector similarity with keyword matching to improve semantic understanding and exact term and acronym matching.

B. Increase the dimensions of the vector embeddings from 384 to 1536. Use a post-processing AWS Lambda function to filter out irrelevant results after retrieval.

C. Replace OpenSearch Service with Amazon Kendra. Use query expansion to handle medical acronyms and terminology variants during pre-processing.

D. Implement a two-stage retrieval architecture in which initial vector search results are re-ranked by an ML model that is hosted on Amazon SageMaker AI.

Correct Answer: A

Question #8Topic 1

A company runs a generative AI (GenAI)-powered summarization application in an application AWS account that uses Amazon Bedrock. The application architecture includes an Amazon API Gateway REST API that forwards requests to AWS Lambda functions that are attached to private VPC subnets. The application summarizes sensitive customer records that the company stores in a governed data lake in a centralized data storage account. The company has enabled Amazon S3, Amazon Athena, and AWS Glue in the data storage account.

The company must ensure that calls that the application makes to Amazon Bedrock use only private connectivity between the company's application VPC and Amazon Bedrock. The company's data lake must provide fine-grained column-level access across the company's AWS accounts.

Which solution will meet these requirements?

A. In the application account, create interface VPC endpoints for Amazon Bedrock runtimes. Run Lambda functions in private subnets. Use IAM conditions on inference and data-plane policies to allow calls only to approved endpoints and roles. In the data storage account, use AWS Lake Formation LF-tag-based access control to create table and column-level cross-account grants.

B. Run Lambda functions in private subnets. Configure a NAT gateway to provide access to Amazon Bedrock and the data lake. Use S3 bucket policies and ACLs to manage permissions. Export AWS CloudTrail logs to Amazon S3 to perform weekly reviews.

https://dbcloud.academy/

C. Create a gateway endpoint only for Amazon S3 in the application account. Invoke Amazon Bedrock through public endpoints. Use database-level grants in AWS Lake Formation to manage data access. Stream AWS CloudTrail logs to Amazon CloudWatch Logs. Do not set up metric filters or alarms.

D. Use VPC endpoints to provide access to Amazon Bedrock and Amazon S3 in the application account. Use only IAM path-based policies to manage data lake access. Send AWS CloudTrail logs to Amazon CloudWatch Logs. Periodically create dashboards and allow public fallback for cross-Region reads to reduce setup time.

Correct Answer: A

Question #9Topic 1

A media company must use Amazon Bedrock to implement a robust governance process for AI-generated content. The company needs to manage hundreds of prompt templates. Multiple teams use the templates across multiple AWS Regions to generate content. The solution must provide version control with approval workflows that include notifications for pending reviews. The solution must also provide detailed audit trails that document prompt activities and consistent prompt parameterization to enforce quality standards.

Which solution will meet these requirements?

A. Configure Amazon Bedrock Studio prompt templates. Use Amazon CloudWatch to create dashboards that display prompt usage metrics. Store the approval status of content in Amazon DynamoDB. Use AWS Lambda functions to enforce approvals.

B. Use Amazon Bedrock Prompt Management to implement version control. Configure AWS CloudTrail for audit logging. Use IAM policies to control approval permissions. Create parameterized prompt templates by specifying variables.

C. Use AWS Step Functions to create an approval workflow. Store prompts as documents in Amazon S3. Use tags to implement version control. Use Amazon EventBridge to send notifications.

D. Deploy Amazon SageMaker Canvas with prompt templates that are stored in Amazon S3. Use AWS CloudFormation to implement version control. Use AWS Config to enforce approval policies.

Correct Answer: B

https://dbcloud.academy/