AIP-C01 Brain Dumps & Reliable AIP-C01 Test Materials

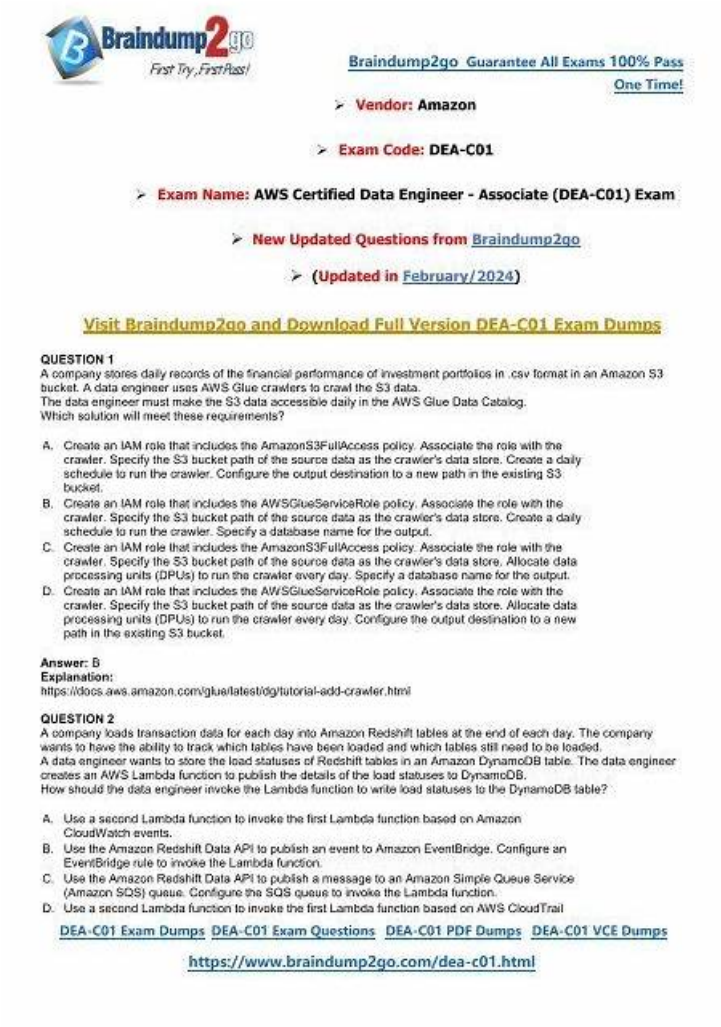

Braindump2go Guarantee All Exams 100% Pass One Time!

- Vendor: Amazon
- Exam Code: DEA-C01
- Exam Name: AWS Certified Data Engineer - Associate (DEA-C01) Exam
- New Updated Questions from Braindump2go
- (Updated in February/2024)

Visit Braindump2go and Download Full Version DEA-C01 Exam Dumps

QUESTION 1
A company stores daily records of the financial performance of investment portfolios in .csv format in an Amazon S3 bucket. A data engineer uses AWS Glue crawlers to crawl the S3 data.
The data engineer must make the S3 data accessible daily in the AWS Glue Data Catalog.
Which solution will meet these requirements?

A. Create an IAM role that includes the AmazonS3FullAccess policy. Associate the role with the crawler. Specify the S3 bucket path of the source data as the crawler's data store. Create a daily schedule to run the crawler. Configure the output destination to a new path in the existing S3 bucket.
B. Create an IAM role that includes the AWSGlueServiceRole policy. Associate the role with the crawler. Specify the S3 bucket path of the source data as the crawler's data store. Create a daily schedule to run the crawler. Specify a database name for the output.
C. Create an IAM role that includes the AmazonS3FullAccess policy. Associate the role with the crawler. Specify the S3 bucket path of the source data as the crawler's data store. Allocate data processing units (DPUs) to run the crawler every day. Specify a database name for the output.
D. Create an IAM role that includes the AWSGlueServiceRole policy. Associate the role with the crawler. Specify the S3 bucket path of the source data as the crawler's data store. Allocate data processing units (DPUs) to run the crawler every day. Configure the output destination to a new path in the existing S3 bucket.

Answer: B
Explanation:
https://docs.aws.amazon.com/glue/latest/dg/tutorial-add-crawler.html

QUESTION 2
A company loads transaction data for each day into Amazon Redshift tables at the end of each day. The company wants to have the ability to track which tables have been loaded and which tables still need to be loaded.
A data engineer wants to store the load statuses of Redshift tables in an Amazon DynamoDB table. The data engineer creates an AWS Lambda function to publish the details of the load statuses to DynamoDB.
How should the data engineer invoke the Lambda function to write load statuses to the DynamoDB table?

A. Use a second Lambda function to invoke the first Lambda function based on Amazon CloudWatch events.
B. Use the Amazon Redshift Data API to publish an event to Amazon EventBridge. Configure an EventBridge rule to invoke the Lambda function.
C. Use the Amazon Redshift Data API to publish a message to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the SQS queue to invoke the Lambda function.
D. Use a second Lambda function to invoke the first Lambda function based on AWS CloudTrail

DEA-C01 Exam Dumps DEA-C01 Exam Questions DEA-C01 PDF Dumps DEA-C01 VCE Dumps

https://www.braindump2go.com/dea-c01.html

What's more, part of that PracticeTorrent AIP-C01 dumps now are free: https://drive.google.com/open?id=1v7O7Wms6CAA4x8FuitDX9wF4EMV5WbTQ

With AIP-C01 practice test questions you can not only streamline your exam Amazon AIP-C01 exam preparation process but also feel confident to pass the challenging AIP-C01 Exam easily. One of the top features of Amazon AIP-C01 valid dumps is their availability in different formats.

Amazon AIP-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	• Foundation Model Integration, Data Management, and Compliance: This domain covers designing GenAI architectures, selecting and configuring foundation models, building data pipelines and vector stores, implementing retrieval mechanisms, and establishing prompt engineering governance.
Topic 2	• Testing, Validation, and Troubleshooting: This domain covers evaluating foundation model outputs, implementing quality assurance processes, and troubleshooting GenAI-specific issues including prompts, integrations, and retrieval systems.

Topic 3	• Implementation and Integration: This domain focuses on building agentic AI systems, deploying foundation models, integrating GenAI with enterprise systems, implementing FM APIs, and developing applications using AWS tools.
Topic 4	• Operational Efficiency and Optimization for GenAI Applications: This domain encompasses cost optimization strategies, performance tuning for latency and throughput, and implementing comprehensive monitoring systems for GenAI applications.
Topic 5	• AI Safety, Security, and Governance: This domain addresses input • output safety controls, data security and privacy protections, compliance mechanisms, and responsible AI principles including transparency and fairness.

Reliable Amazon AIP-C01 Test Materials & AIP-C01 Test Vce

Amazon AIP-C01 valid test cram will help you to get your AIP-C01 certification. It will be a breeze to get your AIP-C01 certification with the help of the PracticeTorrent AIP-C01 pdf vce. We will help whenever you need: 24*7 dedicated email and chat support are available. Besides, we ensure you a flawless shopping experience by Paypal. You can get passed by our latest & updated AIP-C01 Preparation material.

Amazon AWS Certified Generative AI Developer - Professional Sample Questions (Q102-Q107):

NEW QUESTION # 102
A large ecommerce company has deployed a foundation model (FM) to generate product descriptions. The company ' s engineering team monitors technical metrics such as token usage, latency, and error rates by using Amazon CloudWatch. The company ' s marketing team tracks business metrics such as conversion rates and revenue impact in its own systems. The company needs a unified observability solution that correlates technical performance with business outcomes. The solution must provide automatic alerts to stakeholders when operational metrics indicate degradation. The solution must provide comprehensive visibility across both technical and business metrics. Which solution will meet these requirements?

- A. Stream CloudWatch metrics to Amazon S3 by using CloudWatch metric streams. Create Amazon QuickSight dashboards to visualize the combined technical metrics and business metrics. Set up Amazon EventBridge rules to send notifications to stakeholders when metrics exceed predefined thresholds.
- B. Use Amazon Managed Grafana to visualize technical metrics from CloudWatch with business metrics from external sources. Configure Amazon Managed Grafana alerts to invoke AWS Lambda functions.
  Configure the Lambda functions to remediate issues automatically when metrics exceed predefined thresholds.
- C. Configure CloudWatch custom dashboards that integrate operational metrics with imported business metrics. Set up CloudWatch composite alarms with anomaly detection. Use Amazon SNS to create alarm actions to notify stakeholders when correlated metrics indicate performance issues.
- D. Create CloudWatch dashboards that include technical metrics and imported business metrics. Configure CloudWatch composite alarms that combine technical data and business data. Use Amazon SNS to set up notifications to stakeholders.

Answer: C

NEW QUESTION # 103
A retail company is using Amazon Bedrock to develop a customer service AI assistant. Analysis shows that
70% of customer inquiries are simple product questions that a smaller model can effectively handle. However,
30% of inquiries are complex return policy questions that require advanced reasoning.
The company wants to implement a cost-effective model selection framework to automatically route customer inquiries to appropriate models based on inquiry complexity. The framework must maintain high customer satisfaction and minimize response latency.
Which solution will meet these requirements with the LEAST implementation effort?

- A. Create separate Amazon Bedrock endpoints for simple and complex inquiries. Implement a rule-based routing system

based on keyword detection. Use on-demand pricing for the smaller model and provisioned throughput for the larger model.

- B. Use Amazon Bedrock intelligent prompt routing to automatically analyze inquiries. Route simple product inquiries to smaller models and route complex return policy inquiries to more capable larger models.
- C. Implement a single-model solution that uses an Amazon Bedrock mid-sized foundation model (FM) with on-demand pricing. Include special instructions in model prompts to handle both simple and complex inquiries by using the same model.
- D. Create a multi-stage architecture that uses a small foundation model (FM) to classify the complexity of each inquiry. Route simple inquiries to a smaller, more cost-effective model. Route complex inquiries to a larger, more capable model. Use AWS Lambda functions to handle routing logic.

Answer: B

Explanation:
Option B is the correct solution because it leverages native Amazon Bedrock intelligent prompt routing, which is specifically designed to reduce cost and complexity in multi-model GenAI architectures. Intelligent prompt routing automatically analyzes incoming prompts and selects the most appropriate foundation model based on prompt characteristics and complexity-without requiring custom classification logic or orchestration code.
This approach directly meets the requirement for least implementation effort. The company does not need to deploy additional Lambda functions, maintain routing rules, or manage separate classification stages. Routing decisions are handled by Bedrock, which simplifies architecture and reduces operational risk.
By routing the majority (70%) of simple product inquiries to smaller, lower-cost models, the company minimizes inference cost and latency. More complex return policy inquiries are automatically routed to larger models that provide better reasoning capabilities, preserving response quality and customer satisfaction.
Because routing is handled inline by Bedrock, response latency remains low compared to multi-stage architectures that require an additional classification model call before inference. This is critical for customer service scenarios where responsiveness directly impacts satisfaction.
Option A introduces additional inference steps and custom logic. Option C increases cost by overusing a mid- sized model for all queries. Option D relies on brittle keyword rules and increases operational overhead through endpoint management.
Therefore, Option B delivers the optimal balance of cost efficiency, performance, and simplicity for dynamic model selection in Amazon Bedrock.

NEW QUESTION # 104
A company is using Amazon Bedrock to build a customer-facing AI assistant that handles sensitive customer inquiries. The company must use defense-in-depth safety controls to block sophisticated prompt injection attacks. The company must keep audit logs of all safety interventions. The AI assistant must have cross- Region failover capabilities.
Which solution will meet these requirements?

- A. Configure Amazon Bedrock guardrails with custom content filters and word filters set to high.Configure cross-Region guardrail replication for failover. Store logs in AWS CloudTrail for compliance auditing.
- B. Deploy Amazon Comprehend custom classifiers to detect prompt injection attacks. Use Amazon API Gateway request validation. Use CloudWatch Logs to capture intervention events.
- C. Configure Amazon Bedrock guardrails with content filters set to high. Use AWS WAF to block suspicious inputs. Use AWS CloudTrail to log API calls.
- D. Configure Amazon Bedrock guardrails with content filters set to high to protect against prompt injection attacks. Use a guardrail profile to implement cross-Region guardrail inference. Use Amazon CloudWatch Logs with custom metrics to capture detailed guardrail intervention events.

Answer: D

Explanation:
Option A provides the most complete, AWS-native defense-in-depth solution for protecting against prompt injection attacks while meeting audit and resiliency requirements. Amazon Bedrock guardrails are designed specifically to enforce safety policies on both user inputs and model outputs, including protections against prompt injection and jailbreak attempts.
Setting content filters to high increases sensitivity to malicious or manipulative inputs. Guardrail profiles allow the same guardrail configuration to be applied consistently across multiple Regions, enabling cross- Region inference and failover without configuration drift. This directly satisfies the requirement for regional resilience.
Amazon CloudWatch Logs captures detailed guardrail intervention events, including when content is blocked, modified, or flagged. Custom metrics derived from these logs enable fine-grained auditing, alerting, and reporting on safety enforcement actions. This provides a more detailed audit trail of safety interventions than API-level logs alone.
Option B adds WAF protection but lacks detailed guardrail intervention logging. Option C introduces additional services and custom logic that increase complexity and may miss model-specific injection patterns.

Option D references replication concepts that are not aligned with Bedrock guardrail operational models and relies on word filters, which are insufficient against sophisticated prompt injection techniques.
Therefore, Option A best meets the requirements for layered protection, auditability, and cross-Region resilience using managed Amazon Bedrock safety controls.

NEW QUESTION # 105
A company is building a serverless application that uses AWS Lambda functions to help students around the world summarize notes. The application uses Anthropic Claude through Amazon Bedrock. The company observes that most of the traffic occurs during evenings in each time zone. Users report experiencing throttling errors during peak usage times in their time zones.
The company needs to resolve the throttling issues by ensuring continuous operation of the application. The solution must maintain application performance quality and must not require a fixed hourly cost during low traffic periods.
Which solution will meet these requirements?

- A. Create custom Amazon CloudWatch metrics to monitor model errors. Set up a failover mechanism to redirect invocations to a backup AWS Region when the errors exceed a specified threshold.
- B. Enable invocation logging in Amazon Bedrock. Monitor key metrics such as Invocations, InputTokenCount, OutputTokenCount, and InvocationThrottles. Distribute traffic across cross-Region inference endpoints.
- C. Enable invocation logging in Amazon Bedrock. Monitor InvocationLatency, InvocationClientErrors, and InvocationServerErrors metrics. Distribute traffic across multiple versions of the same model.
- D. Create custom Amazon CloudWatch metrics to monitor model errors. Set provisioned throughput to a value that is safely higher than the peak traffic observed.

Answer: B

Explanation:
Option C is the correct solution because it resolves throttling while preserving performance and avoiding fixed costs during low-traffic periods. Amazon Bedrock supports on-demand inference with usage-based pricing, making it well suited for applications with time-zone-dependent traffic spikes.
Throttling during peak hours typically occurs when inference requests exceed available regional capacity.
Cross-Region inference allows Amazon Bedrock to automatically distribute requests across multiple AWS Regions, reducing contention and preventing throttling without requiring reserved or provisioned capacity.
This approach ensures continuous operation while maintaining low latency for users in different geographic locations.
Invocation logging and native metrics such as InvocationThrottles, InputTokenCount, and OutputTokenCount provide visibility into usage patterns and capacity constraints. Monitoring these metrics enables teams to validate that traffic distribution is working as intended and that performance remains consistent during peak periods.
Option A introduces fixed hourly costs by relying on provisioned throughput, which directly violates the requirement to avoid unnecessary spend during low-traffic periods. Option B introduces regional failover complexity and reactive behavior instead of proactive load distribution. Option D does not address the root cause of throttling, as distributing traffic across model versions within the same Region does not increase available capacity.
Therefore, Option C best aligns with AWS Generative AI best practices for scalable, cost-efficient, global serverless applications.

NEW QUESTION # 106
A company has a recommendation system. The system's applications run on Amazon EC2 instances. The applications make API calls to Amazon Bedrock foundation models (FMs) to analyze customer behavior and generate personalized product recommendations.
The system is experiencing intermittent issues. Some recommendations do not match customer preferences.
The company needs an observability solution to monitor operational metrics and detect patterns of operational performance degradation compared to established baselines. The solution must also generate alerts with correlation data within 10 minutes when FM behavior deviates from expected patterns.
Which solution will meet these requirements?

- A. Use Amazon OpenSearch Service with the Observability plugin. Ingest model metrics and logs by using Amazon Kinesis. Create custom Piped Processing Language (PPL) queries to analyze model behavior patterns. Establish operational dashboards to visualize anomalies in real time.
- B. Implement AWS X-Ray to trace requests through the application components. Enable CloudWatch Logs Insights for error pattern detection. Set up AWS CloudTrail to monitor all API calls to Amazon Bedrock. Create custom dashboards in Amazon QuickSight.
- C. Configure Amazon CloudWatch Container Insights for the application infrastructure. Set up CloudWatch alarms for

latency thresholds. Add custom metrics for token counts by using the CloudWatch embedded metric format. Create CloudWatch dashboards to visualize the data.

- D. Enable Amazon CloudWatch Application Insights for the application resources. Create custom metrics for recommendation quality, token usage, and response latency by using the CloudWatch embedded metric format with dimensions for request types and user segments. Configure CloudWatch anomaly detection on the model metrics. Establish log pattern analysis by using CloudWatch Logs Insights.

Answer: D

Explanation:
Option C best satisfies the requirements because it combines application-aware observability, metric baselining, anomaly detection, and correlated alerting using fully managed AWS services with minimal operational overhead. Amazon CloudWatch Application Insights is designed to automatically monitor application health by analyzing metrics, logs, and events across EC2-based workloads. This aligns directly with the need to detect intermittent performance issues and deviations from expected behavior.
By publishing custom metrics using the CloudWatch embedded metric format, the application can track generative AI-specific signals such as recommendation quality indicators, token usage, request volume, and response latency from Amazon Bedrock foundation model calls. Adding dimensions such as request type or user segment enables fine-grained visibility into which workloads or customer groups are impacted when recommendation quality degrades.
A critical requirement is detecting degradation compared to established baselines and generating alerts within 10 minutes. CloudWatch anomaly detection automatically builds statistical models of normal behavior for time-series metrics and flags deviations without requiring manually tuned thresholds. This capability is well suited for monitoring foundation model behavior, which can vary subtly over time. When anomalies are detected, CloudWatch alarms can trigger notifications with contextual metric data quickly, meeting the alerting requirement.
CloudWatch Logs Insights complements the metric-based view by enabling log pattern analysis and correlation. Engineers can query application logs and model response logs to identify recurring error patterns or shifts in output behavior that explain why recommendations no longer align with user preferences.
Application Insights further correlates metrics and logs to surface probable root causes, reducing mean time to resolution.
The other options lack one or more critical elements. Option A focuses on infrastructure-level metrics without baseline anomaly detection. Option B emphasizes tracing and auditing but does not provide automated performance deviation analysis. Option D offers flexibility but requires significantly more development and operational effort than a native CloudWatch-based solution.

NEW QUESTION # 107
.....

- Exam AIP-C01 questions and answers □ Open ➠ www.troytecdumps.com □ enter 【 AIP-C01 】 and obtain a free download □Latest AIP-C01 Study Plan
- 100% AIP-C01 Exam Coverage □ New AIP-C01 Test Pattern ✉ AIP-C01 Test Valid □ Search for □ AIP-C01 □ and download it for free on □ www.pdfvce.com □ website ↕AIP-C01 Practice Guide
- Get Ready for AIP-C01 with Amazon's Realistic Exam Questions and Accurate Answers □ Open website □ www.testkingpass.com □ and search for 《 AIP-C01 》 for free download □Latest AIP-C01 Exam Pdf
- myportal.utt.edu.tt, 0001.yygame.tw, palangshim.com, lms.mastertecheducation.in, secureedges.com, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

BTW, DOWNLOAD part of PracticeTorrent AIP-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1v7O7Wms6CAA4x8FuitDX9wF4EMV5WbTQ