有效的最新AIP-C01題庫資源，最新的考試題庫幫助妳快速通過AIP-C01考試

2026 Fast2test最新的AIP-C01 PDF版考試題庫和AIP-C01考試問題和答案免費分享：https://drive.google.com/open?id=1wsm2wq-rDyhqHn79D5JBwQ331mot61eb

目前，考生報考 Amazon 認證最多的科目：AIP-C01。選擇 AIP-C01 考古題準備考試只是一種方式，優點在于快速有效的幫助考生通過考試。缺點就是缺乏實踐，實踐是在平時的工作之余可以勤加練習。如果決定參加 AIP-C01 認證考試并通過考試，拿到屬于自己的 Amazon 的 AIP-C01 認證是當務之急。而 AIP-C01 考古題可以幫助你在準備考試時節省很多的時間，順利通過考試。

選擇捷徑、使用技巧是為了更好地獲得成功。如果你想獲得一次就通過AIP-C01認證考試的保障，那麼Fast2test的AIP-C01考古題是你唯一的、也是最好的選擇。這絕對是一個讓你禁不住讚美的考古題。你不可能找到比它更好的考試相關的資料了。這個考古題可以讓你更準確地瞭解考試的出題點，從而讓你更有目的地學習相關知識。另外，如果你實在沒有準備考試的時間，那麼你只需要記好這個考古題裏的試題和答案。因為這個考古題包括了真實考試中的所有試題，所以只是這樣你也可以通過考試。

>> 最新AIP-C01題庫資源 <<

新版Amazon AIP-C01題庫 - AIP-C01權威認證

我們Fast2test Amazon的AIP-C01考題按照相同的教學大綱，其次是實際的AIP-C01認證考試，我們也在不斷升級我們的培訓資料，使你在第一時間得到最好和最新的資訊。當你購買我們AIP-C01的考試培訓材料，你所得到的培訓資料有長達一年的免費更新期，你可以隨時延長更新訂閱時間，讓你有更久的時間來準備考試。

Amazon AIP-C01 考試大綱：

主題	簡介
主題 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.
主題 2	• 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.
主題 3	• 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.
主題 4	• 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.

主題 5	• 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.

最新的 Amazon Professional AIP-C01 免費考試真題 (Q97-Q102):

問題 #97
A book publishing company wants to build a book recommendation system that uses an AI assistant. The AI assistant will use ML to generate a list of recommended books from the company's book catalog. The system must suggest books based on conversations with customers.
The company stores the text of the books, customers' and editors' reviews of the books, and extracted book metadata in Amazon S3. The system must support low-latency responses and scale efficiently to handle more than 10,000 concurrent users.
Which solution will meet these requirements?

- A. Use Amazon Bedrock Knowledge Bases to generate embeddings. Store the embeddings as a vector store in Amazon OpenSearch Service. Create an AWS Lambda function that queries the knowledge base. Configure Amazon API Gateway to invoke the Lambda function when handling user requests.
- B. Use Amazon Bedrock Knowledge Bases to generate embeddings. Store the embeddings as a vector store in Amazon DynamoDB. Create an AWS Lambda function that queries the knowledge base.
  Configure Amazon API Gateway to invoke the Lambda function when handling user requests.
- C. Create an Amazon Kendra GenAI Enterprise Edition index that uses the S3 connector to index the book catalog data stored in Amazon S3. Configure built-in FAQ in the Kendra index. Develop an AWS Lambda function that queries the Kendra index based on user conversations. Deploy Amazon API Gateway to expose this functionality and invoke the Lambda function.
- D. Use Amazon SageMaker AI to deploy a pre-trained model to build a personalized recommendation engine for books. Deploy the model as a SageMaker AI endpoint. Invoke the model endpoint by using Amazon API Gateway.

答案：A

解題說明：
Option A best meets the requirements because it directly implements a Retrieval Augmented Generation pattern for conversational recommendations using managed Amazon Bedrock capabilities and a scalable vector store. The company's source data already resides in Amazon S3, which aligns naturally with Amazon Bedrock Knowledge Bases ingestion workflows. A knowledge base can ingest book text, reviews, and metadata, generate embeddings using a supported embedding model, and persist those vectors in a purpose- built vector backend such as Amazon OpenSearch Service. This enables semantic retrieval that is well suited to conversation-driven intent, where user prompts are often descriptive and do not map cleanly to keyword filters.
The requirement to suggest books based on conversations implies the system must interpret natural language context and retrieve relevant passages, reviews, and metadata to ground the recommendation. Knowledge Bases provide managed orchestration for embedding creation and retrieval, which reduces development effort compared to building custom embedding pipelines. OpenSearch Service provides scalable vector search and k- nearest neighbors style similarity retrieval, which supports low-latency responses when properly indexed and sized.
For scaling to more than 10,000 concurrent users, the API layer design in option A is a common AWS pattern: Amazon API Gateway provides a managed front door with throttling and request handling, while AWS Lambda scales horizontally with demand and can invoke the knowledge base retrieval operations. This separates compute scaling from the vector store scaling and helps keep latency predictable under load.
Option B is not the best choice because DynamoDB is not the standard native vector store target for Amazon Bedrock Knowledge Bases in this context and would introduce additional implementation complexity around vector indexing and similarity search behavior. Option C requires substantial ML lifecycle work, model hosting, tuning, and continuous iteration to achieve quality recommendations at scale. Option D provides strong enterprise search, but it focuses on retrieval and FAQs rather than a managed RAG recommendation workflow grounded in embeddings and conversational context for generative responses.

問題 #98
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. 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.
- D. 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.

答案：C

解題說明：
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.

問題 #99
A university is building an AI-powered application that includes several sub-applications. The sub- applications include AI assistants, assignment graders, and internal analytics applications. The university is defining and testing multiple prompts by using various foundation models (FMs). The university wants to compare variants of each prompt and choose the variant that yield outputs that are best-suited for specified use cases. The university requires a version control solution for the prompts. The university must be able to test prompt variations and collect audit trails for prompt changes and usage. The solution must also maintain consistency while allowing the prompts to integrate into the main application. Which combination of solutions will meet these requirements with the LEAST operational overhead? (Select TWO.)

- A. Configure Amazon Bedrock intelligent prompt routing.
- B. Use Amazon Bedrock Flows to create workflows that combine FMs and AWS services.
- C. Configure AWS Config to record prompt changes. Use AWS CloudTrail to track prompt usage.
- D. Use Amazon Bedrock Prompt Management to create versioned prompts. Include parameterized variables for each use case.
- E. Store prompts in Amazon S3. Use AWS Step Functions to orchestrate the model interactions and service integrations.

答案：B,D

解題說明：
Amazon Bedrock Prompt Management is the purpose-built service for prompt lifecycle management. it provides native version control , allowing developers to test and compare variants side-by-side. Use of parameterized variables ensures that a single prompt structure can be consistently reused across different sub-applications (assistants vs. graders) while still being tailored to the specific context. To " integrate into the main application " with minimal overhead, Amazon Bedrock Flows provide a managed orchestration layer.
Flows allow developers to link managed prompts with AWS services (like knowledge bases or Lambda functions) without writing complex state-machine logic in Step Functions (Option B). This combination ensures consistent, auditable, and easily deployable prompt assets across the university ' s diverse use cases.

問題 #100
A financial services company is deploying a generative AI (GenAI) application that uses Amazon Bedrock to assist customer service representatives to provide personalized investment advice to customers. The company must implement a comprehensive governance solution that follows responsible AI practices and meets regulatory requirements.
The solution must detect and prevent hallucinations in recommendations. The solution must have safety controls for customer interactions. The solution must also monitor model behavior drift in real time and maintain audit trails of all prompt-response pairs for regulatory review. The company must deploy the solution within 60 days. The solution must integrate with the company's existing compliance dashboard and respond to customers within 200 ms.
Which solution will meet these requirements with the LEAST operational overhead?

- A. Configure Amazon Bedrock guardrails to apply custom content filters and toxicity detection. Use Amazon Bedrock Model Evaluation to detect hallucinations. Store prompt-response pairs in Amazon DynamoDB to capture audit trails and set a TTL. Integrate Amazon CloudWatch custom metrics with the existing compliance dashboard.
- B. Use Amazon Bedrock Agents and Amazon Bedrock Knowledge Bases to ground responses. Use Amazon Bedrock Guardrails to enforce content safety. Use Amazon OpenSearch Service to store and index prompt-response pairs. Integrate OpenSearch Service with Amazon QuickSight to create compliance reports and to detect model behavior drift.
- C. Deploy Amazon Bedrock and use AWS PrivateLink to access the application securely. Use AWS Lambda functions to implement custom prompt validation. Store prompt-response pairs in an Amazon S3 bucket and configure S3 Lifecycle policies. Create custom Amazon CloudWatch dashboards to monitor model performance metrics.
- D. Use Amazon SageMaker Model Monitor to detect model behavior drift. Use AWS WAF to filter content. Store customer interactions in an encrypted Amazon RDS database. Use Amazon API Gateway to create custom HTTP APIs to integrate with the compliance dashboard.

答案： A

解題說明：
Option A is the correct solution because it uses native Amazon Bedrock governance and evaluation capabilities to meet regulatory, performance, and deployment timeline requirements with the least operational overhead.
Amazon Bedrock guardrails provide built-in safety controls that enforce responsible AI policies directly during inference. Custom content filters and toxicity detection protect customer interactions and prevent disallowed investment guidance patterns without requiring custom application logic. Guardrails operate inline and are optimized for low latency, which helps meet the strict 200 ms response-time requirement.
Hallucination detection is addressed through Amazon Bedrock Model Evaluation, which supports automated evaluation at scale using LLM-as-a-judge techniques. This enables the company to detect factual inaccuracies and policy violations systematically, without building custom evaluation pipelines or requiring extensive human review. Evaluation outputs can be surfaced as metrics.
Storing all prompt-response pairs in Amazon DynamoDB provides a low-latency, highly scalable audit store that aligns with financial regulatory requirements. Using TTL enforces data retention policies automatically, reducing compliance risk and storage overhead.
Amazon CloudWatch custom metrics integrate seamlessly with existing compliance dashboards, allowing near-real-time monitoring of safety interventions, hallucination rates, and drift indicators. CloudWatch anomaly detection can be applied to these metrics to surface behavior changes quickly.
Option B relies on custom Lambda logic and S3-based auditing, increasing latency and operational complexity. Option C introduces additional services that increase setup time and may exceed the 60-day deployment window. Option D uses non-Bedrock-native monitoring and adds unnecessary infrastructure layers.
Therefore, Option A provides the most complete, compliant, and low-overhead governance solution for a regulated GenAI financial services application.

問題 #101
A company wants to select a new FM for its AI assistant. A GenAI developer needs to generate evaluation reports to help a data

scientist assess the quality and safety of various foundation models FMs. The data scientist provides the GenAI developer with sample prompts for evaluation. The GenAI developer wants to use Amazon Bedrock to automate report generation and evaluation. Which solution will meet this requirement?

- A. Combine the sample prompts into a single JSON document. Create an Amazon Bedrock knowledge base from the document. Create an Amazon Bedrock evaluation job that uses the retrieval and response generation evaluation type. Specify an Amazon S3 bucket as the output. Run an evaluation job for each FM.
- B. Combine the sample prompts into a single JSONL document. Store the document in an Amazon S3 bucket. Create an Amazon Bedrock evaluation job that uses a judge model. Specify the S3 location as input and Amazon QuickSight as output. Run an evaluation job for each FM and select the FM as the evaluator.
- C. Combine the sample prompts into a single JSONL document. Store the document in an Amazon S3 bucket. Create an Amazon Bedrock evaluation job that uses a judge model. Specify the S3 location as input and a different S3 location as output. Run an evaluation job for each FM and select the FM as the generator.
- D. Combine the sample prompts into a single JSON document. Create an Amazon Bedrock knowledge base with the document. Write a prompt that asks the FM to generate a response to each sample prompt.
  Use the RetrieveAndGenerate API to generate a report for each model.

答案：C

解題說明：
Option B is correct because it uses the managed evaluation capability in Amazon Bedrock that is intended specifically for comparing foundation models using a consistent prompt set and producing structured results with minimal custom tooling. In a Bedrock evaluation workflow, you provide an input dataset of prompts, typically in JSON Lines format so each line represents one evaluation record. Storing the JSONL file in Amazon S3 allows Bedrock to read the dataset at scale and write standardized evaluation outputs back to S3 for downstream analysis, sharing, and retention.
The key requirement is to assess both quality and safety across multiple models. A Bedrock evaluation job can use a judge model to score the generated outputs against defined criteria. This approach supports repeatable, apples-to-apples comparisons because the same judge model and scoring rubric can be applied to every candidate foundation model. The candidate models are configured as generators, meaning each evaluation job run uses one selected FM to produce answers for the same prompt set, and the judge model evaluates those answers. That matches the requirement to generate evaluation reports that help a data scientist select the best FM.
Option A does not use Bedrock evaluation jobs, and a knowledge base plus RetrieveAndGenerate is a RAG pattern, not an evaluation framework. It would produce responses but not standardized scoring and reporting suitable for model selection. Option C is incorrect because Bedrock evaluation outputs are delivered to S3, not directly to a BI destination, and selecting the candidate FM as the evaluator conflicts with the intended pattern of using a stable judge model. Option D misuses knowledge bases and retrieval evaluation types when the requirement is prompt-based model assessment rather than evaluating retrieval quality.

問題 #102
.....

- 高質量的最新AIP-C01題庫資源和資格考試中的領先供應平臺＆有效的AIP-C01：AWS Certified Generative AI Developer - Professional □ 來自網站➽ www.newdumpspdf.com □打開並搜索▹ AIP-C01 ◃免費下載新版AIP-C01考古題
- 專業的最新AIP-C01題庫資源和資格考試領先提供商和可信賴的新版AIP-C01題庫 □ 到➤ www.newdumpspdf.com □搜索➥ AIP-C01 □輕鬆取得免費下載AIP-C01考古題更新
- AIP-C01題庫資料 □ AIP-C01認證 □ AIP-C01證照考試 □ 在□ www.newdumpspdf.com □網站上查找✔ AIP-C01 □✔□的最新題庫新版AIP-C01考古題
- 有利的最新AIP-C01題庫資源，最新的學習資料幫助妳快速通過AIP-C01考試 □ 透過“ www.kaoguti.com ”搜索➽ AIP-C01 □免費下載考試資料新版AIP-C01考古題
- hashnode.com, ycs.instructure.com, zenwriting.net, www.stes.tyc.edu.tw, training.emecbd.com, spanishatjuans.com, portfolium.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, Disposable vapes

P.S. Fast2test在Google Drive上分享了免費的、最新的AIP-C01考試題庫：https://drive.google.com/open?id=1wsm2wq-rDyhqHn79D5JBwQ331mot61eb