


AIP-C01在線考題 - AIP-C01認證考試解析



所有的Amazon職員都知道，AIP-C01認證考試的資格是不容易拿到的。但是，參加AIP-C01認證考試獲得資格又是提升自己能力以及更好地證明自己的價值的途徑，所以不得不選擇。那麼，難道沒有一個簡單的方法可以讓大家更容易地通過Amazon認證考試嗎？當然有了。Fast2test的考古題就是一個最好的方法。Fast2test有你需要的所有資料，絕對可以滿足你的要求。你可以到Fast2test的網站了解更多的信息，找到你想要的考試資料。

Amazon AIP-C01 考試大綱：

| 主題 | 簡介 |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 主題 1 | <ul style="list-style-type: none">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 | <ul style="list-style-type: none">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 3 | <ul style="list-style-type: none">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 | <ul style="list-style-type: none">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. |

主題 5

- 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.

>> AIP-C01在線考題 <<

AIP-C01認證考試解析，AIP-C01最新試題

Amazon是一個成功的公司，提供各種認證和考試。通過AIP-C01考試是其中的核心要求。將帶來一個新的前沿，對你的職業道路起著如此重要的角色。AIP-C01認證考試的考題按照相同的教學大綱，其次是實際的Amazon的AIP-C01認證考試，我們也是不斷的升級我們的培訓資料，你得到的所有產品高達1年的免費更新，你也可以隨時延長更新訂閱時間，你將得到更多的時間來充分準備考試。

最新的 Amazon Professional AIP-C01 免費考試真題 (Q21-Q26):

問題 #21

A software company is using Amazon Q Business to build an AI assistant that allows employees to access company information and personal information by using natural language prompts. The company stores this information in an Amazon S3 bucket.

Each department in the company has a dedicated prefix in the S3 bucket. Each object name includes the S3 prefix of the department that it belongs to. Each department can belong to only a single group in AWS IAM Identity Center. Each employee belongs to a single department.

The company configures Amazon Q Business to access data stored in an S3 bucket as a data source. The company needs to ensure that the AI assistant respects access controls based on the user's IAM Identity Center group membership.

Which solution will meet this requirement with the LEAST operational overhead?

- A. Create a single JSON file named `acl.json` at the top level of the S3 bucket. Add access control entries that map each department's S3 prefix to its corresponding IAM Identity Center group. Indicate the location of the JSON file in the Access Control section of the data source settings.
- B. Create a JSON file named `acl.json` in each department folder. In each file, create access control entries that specify the IAM Identity Center group that should have access to that department's data. Indicate the location of the JSON file in the Access Control section of the data source settings.
- C. For each IAM Identity Center group, create a separate permissions set that denies access to all prefixes in the S3 bucket. Add a `StringNotEquals` condition key to the permissions set for each group that specifies the department each group is associated with. Attach the permissions sets to the Identity Center groups.
- D. Create a metadata file named `metadata.json` at the top level of the S3 bucket. Add an `AccessControlList` object to the file that specifies the S3 path of each department's prefix. Specify the IAM Identity Center group that should have access to each department's prefix. Reference the file location in the data source metadata settings.

答案：A

解題說明：

Option B is the correct solution because Amazon Q Business natively supports access control lists (ACLs) for S3 data sources using a single, centralized JSON file that maps S3 prefixes to IAM Identity Center groups.

This approach directly aligns with the company's data organization model, where each department's data is stored under a distinct S3 prefix and each employee belongs to exactly one department group.

Using a single `acl.json` file at the bucket root minimizes operational overhead by centralizing access control logic in one location. Administrators can update department mappings without touching individual folders or changing IAM permissions, which simplifies governance and reduces the risk of configuration drift. Amazon Q Business automatically evaluates the user's IAM Identity Center group membership at query time and filters accessible documents accordingly.

Option A increases operational complexity by requiring a separate ACL file in every department folder, which becomes difficult to maintain as departments or prefixes change. Option C attempts to enforce access using IAM permissions sets, but Amazon Q Business access control for S3 data sources is not designed to be managed through IAM condition logic and would significantly increase complexity. Option D introduces a custom metadata structure that is not the supported mechanism for Amazon Q Business access enforcement.

Therefore, Option B provides the cleanest, most scalable, and AWS-recommended solution for enforcing department-based access control with the least operational effort.

問題 #22

A company is building a generative AI (GenAI) application that produces content based on a variety of internal and external data sources. The company wants to ensure that the generated output is fully traceable.

The application must support data source registration and enable metadata tagging to attribute content to its original source. The application must also maintain audit logs of data access and usage throughout the pipeline.

Which solution will meet these requirements?

- A. Use AWS Lake Formation to catalog data sources and control access. Apply metadata tags directly in Amazon S3. Use AWS CloudTrail to monitor API activity.
- B. Store data in Amazon S3 and use object tagging for attribution. Use AWS Glue Data Catalog to manage schema information. Use AWS CloudTrail to log access to S3 buckets.
- C. Use AWS Glue Data Catalog to register and tag data sources. Use Amazon CloudWatch Logs to monitor access patterns and application behavior.
- D. Use AWS Glue Data Catalog to register all data sources. Apply metadata tags to attribute data sources. Use AWS CloudTrail to log access and activity across services.

答案: D

解題說明:

Option D is the correct solution because it directly satisfies all three core requirements: data source registration, metadata-based attribution, and end-to-end audit logging, while remaining service-agnostic and scalable across internal and external data sources. The AWS Glue Data Catalog is the AWS-native service for registering datasets and managing metadata centrally. It supports structured registration of diverse data sources and enables consistent tagging that can be used to attribute generated content back to its original source. This is essential for GenAI applications that combine multiple datasets and must provide traceability for outputs. Metadata tags applied within the Glue Data Catalog ensure a consistent attribution framework that downstream systems—such as Retrieval Augmented Generation (RAG) pipelines or evaluation systems—can reference without embedding attribution logic directly in application code. This improves maintainability and governance.

AWS CloudTrail provides immutable audit logs of API activity across AWS services, including data access, metadata changes, and pipeline interactions. CloudTrail logs are critical for compliance and regulatory review because they capture who accessed which data, when, and through which service. This satisfies the requirement to maintain audit logs "throughout the pipeline," not just at storage or application layers.

Option A introduces Lake Formation, which is primarily intended for fine-grained data lake permissions and is not required solely for traceability. Option B relies on CloudWatch Logs, which does not provide authoritative audit logging across services. Option C limits audit scope to S3 access and does not register or govern all data sources comprehensively.

Therefore, Option D provides the most complete and least intrusive solution for traceable, auditable GenAI data pipelines.

問題 #23

A financial technology company is using Amazon Bedrock to build an assessment system for the company's customer service AI assistant. The AI assistant must provide financial recommendations that are factually accurate, compliant with financial regulations, and conversationally appropriate. The company needs to combine automated quality evaluations at scale with targeted human reviews of critical interactions.

What solution will meet these requirements?

- A. Create an Amazon Lex bot to manage customer service interactions. Configure AWS Lambda functions to check responses against a static compliance database. Configure intents that call the Lambda functions. Add an additional intent to collect end-user reviews.
- B. Configure a pipeline in which financial experts manually score all responses for accuracy, compliance, and conversational quality. Use Amazon SageMaker notebooks to analyze results to identify improvement areas.
- C. Configure Amazon CloudWatch to monitor response patterns from the AI assistant. Configure CloudWatch alerts for potential compliance violations. Establish a team of human evaluators to review flagged interactions.
- D. Configure Amazon Bedrock evaluations that use Anthropic Claude Sonnet as a judge model to assess response accuracy and appropriateness. Configure custom Amazon Bedrock guardrails to check responses for compliance with financial policies. Add Amazon Augmented AI (Amazon A2I) human reviews for flagged critical interactions.

答案: D

解題說明:

Option B meets the requirement to combine scalable automated evaluation with targeted human oversight using managed AWS GenAI capabilities. Amazon Bedrock evaluations enable systematic, repeatable quality assessment across large volumes of interactions. Using an LLM-as-a-judge approach with a strong evaluator model such as Anthropic Claude Sonnet allows the

company to automatically score outputs for dimensions like factual accuracy, conversational appropriateness, and policy alignment. This directly supports "automated quality evaluations at scale" without building custom scoring models. However, financial recommendations add higher risk because regulatory compliance requires additional enforcement beyond general quality scoring. Amazon Bedrock guardrails provide a dedicated policy enforcement layer that can block or intervene when responses violate compliance constraints. Guardrails are particularly important for preventing disallowed financial guidance patterns and ensuring consistent behavior across deployments.

The requirement also calls for "targeted human reviews of critical interactions." Amazon Augmented AI (A2I) is a managed human review service that supports routing specific items to human reviewers based on rules or confidence thresholds. In this design, the system can automatically send only high-risk or policy-flagged interactions to qualified financial experts for review, keeping human effort focused where it matters most while maintaining scale.

Option A is not scalable because it requires manual review of all responses. Option C relies on static rules and end-user feedback, which is insufficient for regulatory compliance and factual accuracy assurance. Option D provides monitoring but not structured quality evaluation or policy enforcement.

Therefore, Option B provides the most complete, AWS-aligned solution for scalable evaluation plus human oversight in a regulated financial context.

問題 #24

A GenAI developer is building a Retrieval Augmented Generation (RAG)-based customer support application that uses Amazon Bedrock foundation models (FMs). The application needs to process 50 GB of historical customer conversations that are stored in an Amazon S3 bucket as JSON files. The application must use the processed data as its retrieval corpus. The application's data processing workflow must extract relevant data from customer support documents, remove customer personally identifiable information (PII), and generate embeddings for vector storage. The processing workflow must be cost-effective and must finish within 4 hours.

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

- A. Use AWS Lambda and Amazon Comprehend to process files in parallel, remove PII, and call Amazon Bedrock APIs to generate vectors. Configure Lambda concurrency limits and memory settings to optimize throughput.
- B. Deploy an Amazon EMR cluster that runs Apache Spark with user-defined functions (UDFs) that call Amazon Comprehend to detect PII. Use Amazon Bedrock APIs to generate vectors. Store outputs in Amazon Aurora PostgreSQL with the pgvector extension.
- C. Implement a data processing pipeline that uses AWS Step Functions to orchestrate a workload that uses Amazon Comprehend to detect PII and Amazon Bedrock to generate embeddings. Directly integrate the workflow with Amazon OpenSearch Serverless to store vectors and provide similarity search capabilities.
- D. Create an AWS Glue ETL job to run PII detection scripts on the data. Use Amazon SageMaker Processing to run the HuggingFaceProcessor to generate embeddings by using a pre-trained model. Store the embeddings in Amazon OpenSearch Service.

答案: C

解題說明:

Comprehensive and Detailed 250 to 350 words of Explanation From AWS Generative AI concepts and services documents: Option D is the best solution because it delivers a fully managed, scalable pipeline with minimal infrastructure management while meeting the 50 GB and 4-hour constraint. AWS Step Functions provides a serverless orchestration layer that can coordinate parallel processing steps, retries, and error handling without managing clusters or tuning long-running compute.

Using Amazon Comprehend for PII detection fulfills the requirement to remove customer PII in a managed and consistent way. Step Functions can coordinate Comprehend calls at scale and route sanitized outputs into the embedding step. Generating embeddings with Amazon Bedrock keeps the entire workflow within AWS managed services, eliminates the need to maintain custom embedding models, and supports consistent vector representations for downstream retrieval.

Direct integration with Amazon OpenSearch Serverless provides a low-operations vector store that can handle large-scale indexing and similarity search without cluster sizing, node maintenance, or shard management.

This aligns strongly with the requirement for least operational overhead and supports growth beyond the initial 50 GB corpus. Step Functions can batch and parallelize ingestion into OpenSearch Serverless to meet the 4-hour completion goal in a cost-effective manner by controlling concurrency, chunk sizes, and failure handling.

Option A can be difficult and costly at this scale because Lambda concurrency and per-invocation overhead can become complex to tune for 50 GB within 4 hours. Option B introduces SageMaker Processing and embedding model management, increasing operational complexity. Option C requires EMR cluster management and tuning, which is the opposite of minimal overhead.

Therefore, Option D is the most operationally efficient, scalable, and managed approach to build the required PII-sanitized embedding pipeline for a RAG corpus.

問題 #25

A company has a customer service application that uses Amazon Bedrock to generate personalized responses to customer inquiries. The company needs to establish a quality assurance process to evaluate prompt effectiveness and model configurations across updates. The process must automatically compare outputs from multiple prompt templates, detect response quality issues, provide quantitative metrics, and allow human reviewers to give feedback on responses. The process must prevent configurations that do not meet a predefined quality threshold from being deployed.

Which solution will meet these requirements?

- A. Use Amazon Bedrock evaluation jobs to compare model outputs by using custom prompt datasets.
Configure AWS CodePipeline to run the evaluation jobs when prompt templates change. Configure CodePipeline to deploy only configurations that exceed the predefined quality threshold.
- B. Use AWS Lambda functions to create an automated testing framework that samples production traffic and routes duplicate requests to the updated model version. Use Amazon Comprehend sentiment analysis to compare results. Block deployment if sentiment scores decrease.
- C. Set up Amazon CloudWatch alarms to monitor response latency and error rates from Amazon Bedrock.
Use Amazon EventBridge rules to notify teams when thresholds are exceeded. Configure a manual approval workflow in AWS Systems Manager.
- D. Create an AWS Lambda function that sends sample customer inquiries to multiple Amazon Bedrock model configurations and stores responses in Amazon S3. Use Amazon QuickSight to visualize response patterns. Manually review outputs daily.
Use AWS CodePipeline to deploy configurations that meet the quality threshold.

答案： A

解題說明：

Option B is the correct solution because Amazon Bedrock evaluation jobs are purpose-built to assess prompt effectiveness, model behavior, and response quality in a repeatable and automated manner. Evaluation jobs support both quantitative metrics and LLM-based judgment, making them suitable for detecting subtle response quality regressions that simple sentiment or latency metrics cannot capture.

By using custom prompt datasets, the company can consistently test multiple prompt templates and model configurations against the same inputs. This enables accurate comparison across updates and eliminates variability introduced by live traffic sampling. Amazon Bedrock evaluation jobs also support structured scoring outputs, which can be used to enforce objective quality thresholds.

Integrating evaluation jobs directly into AWS CodePipeline ensures that quality checks are automatically triggered whenever prompt templates or configurations change. This creates a gated deployment workflow in which only configurations that meet or exceed the predefined quality threshold are promoted. This directly satisfies the requirement to prevent low-quality configurations from being deployed.

Human reviewers can be incorporated by reviewing evaluation results and scores produced by the jobs, enabling informed feedback without manual data collection. Option A and D rely on custom frameworks and indirect quality signals, increasing complexity and reducing reliability. Option C focuses on operational health rather than response quality.

Therefore, Option B provides the most robust, scalable, and AWS-aligned quality assurance process for Amazon Bedrock-based applications.

問題 #26

.....

我們承諾，使用我們Fast2test Amazon的AIP-C01的考試培訓資料，確保你在你的第一次嘗試中通過測試，如果你準備考試使用我們Fast2test Amazon的AIP-C01考試培訓資料，我們保證你通過，如果沒有通過測試，我們給你退還購買的全額退款，送你一個相同價值的免費產品。

AIP-C01認證考試解析: <https://tw.fast2test.com/AIP-C01-premium-file.html>

- 信任www.vcesoft.com中的授權的AIP-C01在線考題是通過AWS Certified Generative AI Developer - Professional的有效方式 □ ⇒ www.vcesoft.com ← 網站搜索 ➔ AIP-C01 □ 並免費下載AIP-C01考試
- 最優質的的AIP-C01在線考題和資格考試中的領先供應商和完整的AIP-C01認證考試解析 □ 《 www.newdumpspdf.com 》提供免費* AIP-C01 □ * □ 問題收集AIP-C01認證題庫
- 獲取更新AIP-C01在線考題 - 全部在www.vcesoft.com □ 透過 ➔ www.vcesoft.com □ □ □ 搜索《 AIP-C01 》免費下載考試資料AIP-C01考試備考經驗
- AIP-C01認證題庫 □ AIP-C01真題 □ 新版AIP-C01考古題 □ 請在《 www.newdumpspdf.com 》網站上免費下載【 AIP-C01 】題庫新版AIP-C01考古題
- 最新AIP-C01考題 □ AIP-C01考古題 □ AIP-C01學習指南 □ 在 ➔ www.newdumpspdf.com ← 網站下載免費“ AIP-C01 ”題庫收集AIP-C01考古題

