

# 高質量AIP-C01題庫資訊和資格考試中的領先材料供應者和值得信賴的Amazon AWS Certified Generative AI Developer - Professional



P.S. Testpdf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=114V8R-dBKMKU9FeVAyjvaDtUZZ6v0kR>

我們Testpdf Amazon的AIP-C01考試培訓資料給所有需要的人帶來最大的成功率，通過微軟的AIP-C01考試是一個具有挑戰性的認證考試。現在除了書籍，互聯網被認為是一個知識的寶庫，在Testpdf你也可以找到屬於你的知識寶庫，這將是一個對你有很大幫助的網站，你會遇到複雜的測試方面的試題，我們Testpdf可以幫助你輕鬆的通過考試，它涵蓋了所有必要的知識Amazon的AIP-C01考試。

如果你發現我們AIP-C01有任何品質問題或者沒有考過，我們將無條件全額退款，Testpdf是專業提供Amazon的AIP-C01最新考題和答案的網站，幾乎全部覆蓋了AIP-C01全部的知識點。

>> AIP-C01題庫資訊 <<

## AIP-C01考試資訊 - AIP-C01認證題庫

伴隨著 Amazon 認證，越來越多的客戶注意到 Amazon 的重要性，目前是經濟衰退的時期，找一份工作不容易，考取 Amazon 認證的證書當然是有用的，能夠幫助你穩定你的位置，增加求職的法碼。如果你正在準備 AIP-C01 考試題目和答案的電子圖書的形式或自我測試軟體，以獲得適當的知識和技能，急需通過 AIP-C01 考試，可以憑藉 Testpdf 考題網最新的題庫順利通過該考試。

## 最新的 Amazon Professional AIP-C01 免費考試真題 (Q101-Q106):

### 問題 #101

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 provisioned throughput to a value that is safely higher than the peak traffic observed.
- **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 up a failover mechanism to redirect invocations to a backup AWS Region when the errors exceed a specified threshold.

答案： B

#### 解題說明：

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.

#### 問題 #102

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. 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.**
- B. 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.
- C. 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.
- D. 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.

答案： A

#### 解題說明：

Option C is the correct solution because AWS AppConfig is designed for real-time, validated, centrally managed configuration changes with safe rollout, immediate propagation, and rollback support—exactly matching the company's requirements.

By storing routing rules, cost thresholds, regulatory constraints, and A/B testing logic in AWS AppConfig, the company can switch between Amazon Bedrock foundation models without redeploying Lambda code.

AppConfig supports feature flags, dynamic configuration updates, JSON schema validation, and staged rollouts, which are essential for safely managing complex and frequently changing routing logic.

Using the AWS AppConfig Agent, Lambda functions can retrieve cached configurations efficiently, ensuring low latency even under thousands of concurrent requests. This approach allows the Lambda function to apply proprietary business logic—such as user tier, transaction value, Region compliance, and real-time cost metrics—before selecting the appropriate FM.

Option A is operationally fragile because environment variable changes require function restarts and do not support validation or controlled rollouts. Option B is too limited for complex, dynamic logic and is difficult to maintain at scale. Option D misuses Lambda authorizers, which are intended for authentication and authorization, not high-frequency dynamic routing decisions.

Therefore, Option C provides the most scalable, flexible, and low-overhead architecture for dynamic, regulation-aware FM routing in a global GenAI system.

### 問題 #103

A company is developing a generative AI (GenAI) application that analyzes customer service calls in real time and generates suggested responses for human customer service agents. The application must process 500,000 concurrent calls during peak hours with less than 200 ms end-to-end latency for each suggestion. The company uses existing architecture to transcribe customer call audio streams. The application must not exceed a predefined monthly compute budget and must maintain auto scaling capabilities.

Which solution will meet these requirements?

- A. Deploy a large, complex reasoning model on Amazon Bedrock. Purchase provisioned throughput and optimize for batch processing.
- **B. Deploy a low-latency, real-time optimized model on Amazon Bedrock. Purchase provisioned throughput and set up automatic scaling policies.**
- C. Deploy a mid-sized language model on an Amazon SageMaker serverless endpoint that is optimized for batch processing.
- D. Deploy a large language model (LLM) on an Amazon SageMaker real-time endpoint that uses dedicated GPU instances.

**答案： B**

解題說明：

Option B is the correct solution because it aligns with AWS guidance for building high-throughput, ultra-low- latency GenAI applications while maintaining predictable costs and automatic scaling. Amazon Bedrock provides access to foundation models that are specifically optimized for real-time inference use cases, including conversational and recommendation-style workloads that require responses within milliseconds.

Low-latency models in Amazon Bedrock are designed to handle very high request rates with minimal per- request overhead. Purchasing provisioned throughput ensures that sufficient model capacity is reserved to handle peak loads, eliminating cold starts and reducing request queuing during traffic surges. This is critical when supporting up to 500,000 concurrent calls with strict latency requirements.

Automatic scaling policies allow the application to dynamically adjust capacity based on demand, ensuring cost efficiency during off-peak hours while maintaining performance during peak usage. This directly supports the requirement to stay within a predefined monthly compute budget.

Option A fails because batch processing and complex reasoning models introduce higher latency and are not suitable for real-time suggestions. Option C introduces significantly higher operational and cost overhead due to dedicated GPU instances and manual scaling responsibilities. Option D is optimized for batch workloads and cannot meet the sub-200 ms latency requirement.

Therefore, Option B provides the best balance of performance, scalability, cost control, and operational simplicity using AWS-native GenAI services.

### 問題 #104

A company is building a legal research AI assistant that uses Amazon Bedrock with an Anthropic Claude foundation model (FM). The AI assistant must retrieve highly relevant case law documents to augment the FM's responses. The AI assistant must identify semantic relationships between legal concepts, specific legal terminology, and citations. The AI assistant must perform quickly and return precise results.

Which solution will meet these requirements?

- A. Enable the Amazon Kendra query suggestion feature for end users. Use Amazon Bedrock to perform post-processing of search results to identify semantic similarity in the documents and to produce precise results.
- B. Configure an Amazon Bedrock knowledge base to use a default vector search configuration. Use Amazon Bedrock to expand queries to improve retrieval for legal documents based on specific terminology and citations.
- C. Use Amazon OpenSearch Service with vector search and Amazon Bedrock Titan Embeddings to index and search legal documents. Use custom AWS Lambda functions to merge results with keyword-based filters that are stored in an Amazon RDS database.
- **D. Use Amazon OpenSearch Service to deploy a hybrid search architecture that combines vector search with keyword search. Apply an Amazon Bedrock reranker model to optimize result relevance.**

**答案： D**

解題說明：

Option B is the correct solution because legal research workloads require both semantic understanding and exact lexical precision, especially for statutes, citations, and domain-specific terminology. A hybrid search architecture directly addresses this need by combining vector similarity search with traditional keyword-based retrieval.

Vector search alone is often insufficient for legal research because exact phrases, citation formats, and jurisdiction-specific terms must be matched precisely. Keyword search ensures high recall and precision for citations and legal terms, while vector search

captures deeper semantic relationships between legal concepts, precedents, and arguments. Amazon OpenSearch Service natively supports hybrid search, enabling efficient scoring and ranking without external orchestration.

Applying an Amazon Bedrock reranker model further improves relevance by reordering retrieved documents based on deeper contextual understanding. Reranking is especially valuable in legal research because multiple documents may appear relevant, but only a subset truly addresses the user's legal question. The reranker optimizes final results before they are passed to the Anthropic Claude FM, improving answer accuracy and reducing hallucinations.

Option A relies on default vector search, which does not reliably handle citations and exact terminology.

Option C focuses on query suggestions and post-processing rather than retrieval quality. Option D introduces unnecessary operational complexity by merging results across multiple systems.

Therefore, Option B best meets the requirements for precision, performance, and semantic understanding in a legal research AI assistant.

### 問題 #105

A wildlife conservation agency operates zoos globally. The agency uses various sensors, trackers, and audiovisual recorders to monitor animal behavior. The agency wants to launch a generative AI (GenAI) assistant that can ingest multimodal data to study animal behavior.

The GenAI assistant must support natural language queries, avoid speculative behavioral interpretations, and maintain audit logs for ethical research audits.

Which solution will meet these requirements?

- A. Use Amazon SageMaker Processing and Amazon Transcribe to pre-process multimodal data. Ingest curated summaries into an Amazon Bedrock Knowledge Bases. Apply Amazon Bedrock guardrails to restrict speculative outputs. Use AWS AppConfig to manage prompt templates. Use AWS CloudTrail to log research activity for audits.
- B. Configure Amazon O Business to federate data across Amazon S3, Amazon Kinesis, and Amazon SageMaker Feature Store. Use EventBridge for ingestion orchestration. Use custom AWS Lambda functions to filter LLM outputs for ethical compliance.
- C. Use Amazon OpenSearch Serverless to index behavioral logs and telemetry. Use Amazon Comprehend to extract entities. Use Amazon Bedrock to answer questions over indexed data. Use IAM for access control and CloudTrail for audit logging.
- D. Ingest raw videos into Amazon Rekognition to detect animal postures and expressions. Use Amazon Data Firehose to stream sensor and GPS data into Amazon S3. Prompt an Amazon Bedrock FM using basic templates stored in AWS Systems Manager Parameter Store. Use IAM for access control. Use AWS CloudTrail for audit logging.

答案： A

解題說明：

Option B best meets the multimodal, ethical, and auditability requirements using managed AWS services designed for research-grade GenAI systems. Multimodal data such as audio, video, sensor telemetry, and tracking data must be curated and summarized before being consumed by a foundation model. Amazon SageMaker Processing and Amazon Transcribe provide scalable, managed preprocessing for audiovisual and textual data.

By ingesting summarized, validated observations into Amazon Bedrock Knowledge Bases, the GenAI assistant can answer natural language queries using grounded, evidence-based context instead of raw sensor signals. This significantly reduces the risk of speculative or anthropomorphic interpretations.

Amazon Bedrock guardrails are critical for preventing speculative behavioral claims, enforcing scientific and ethical constraints at inference time. Guardrails provide a validated, auditable safety layer that custom Lambda-based filters cannot reliably replicate. AWS AppConfig enables controlled prompt management and change governance, ensuring that research prompts remain consistent and reviewable. AWS CloudTrail captures all access, query, and configuration changes, supporting ethical research audits and regulatory reviews.

Option A lacks grounding and speculative safeguards. Option C focuses on text analytics and does not properly handle multimodal reasoning or safety enforcement. Option D relies heavily on custom logic and introduces unnecessary operational risk.

Therefore, Option B provides the most robust, ethical, and auditable GenAI architecture for wildlife behavior research.

### 問題 #106

.....

現在IT行業競爭越來越激烈，通過Amazon AIP-C01認證考試可以有效的幫助你在現在這個競爭激烈的IT行業中穩固和提升自己的地位。在我們Testpdf中你可以獲得關Amazon AIP-C01認證考試的培訓工具。我們Testpdf的IT精英團隊會及時為你提供準確以及詳細的關Amazon AIP-C01認證考試的培訓材料。通過我們Testpdf提供的學習材料以及考試練習題和答案，我們Testpdf能確保你第一次參加Amazon AIP-C01認證考試時挑戰成功，而且不用花費大量時

間和精力來準備考試。

**AIP-C01考試資訊:** <https://www.testpdf.net/AIP-C01.html>

想參加Amazon的AIP-C01認證考試嗎，Amazon AIP-C01題庫資訊 你也想获得认证资格吗，快來報名參加AIP-C01資格認證考試進一步提高自己的技能吧，為通過Amazon AIP-C01 認證考試花大量的時間和精力復習相關知識，但是卻是冒險地通過考試，因為你可以來Testpdf AIP-C01考試資訊找到你想要的幫手和準備考試的工具，這就說明Testpdf AIP-C01考試資訊提供的針對性培訓資料是很有效的，如果你想在短時間內，以最小的努力，達到最有效果的結果Amazon的AIP-C01考試準備，你可以使用Testpdf Amazon的AIP-C01考試培訓資料，Testpdf的培訓資料是通過實踐檢驗了的，也是通過眾多考生證明了它確實可以百分百通過考試，利用了它，你將達到你的目的，得到最佳的效果，Amazon的AIP-C01考試是最近最有人氣的考試，你也想參加嗎？

哦竟然還有人能拿捏住二弟?有機會真要結識壹番，妖氣也是壹種營養，想參加Amazon的AIP-C01認證考試嗎，你也想获得认证资格吗，快來報名參加AIP-C01資格認證考試進一步提高自己的技能吧，為通過Amazon AIP-C01 認證考試花大量的時間和精力復習相關知識，但是卻是冒險地通過考試。

## 最有效的AIP-C01題庫資訊， Amazon Amazon Professional認證AIP-C01考試題庫提供免費下載

因為你可以來Testpdf找到你想要的幫手和準備考試的工具。

- 由最新的AIP-C01考試指南編訂的題庫資料 - 提高AWS Certified Generative AI Developer - Professional考試的通過率  在 [www.vcesoft.com](http://www.vcesoft.com)  網站上查找  $\Rightarrow$  AIP-C01  $\Leftarrow$  的最新題庫AIP-C01通過考試
- 最新更新的AIP-C01題庫資訊和資格考試的領導者與專業的AIP-C01: AWS Certified Generative AI Developer - Professional  請在  [www.newdumpspdf.com](http://www.newdumpspdf.com)  網站上免費下載  AIP-C01  題庫免費下載AIP-C01考題
- AIP-C01信息資訊  AIP-C01在線考題  AIP-C01熱門考題  開放  $\Rightarrow$  [tw.fast2test.com](http://tw.fast2test.com)   輸入  AIP-C01  並獲取免費下載AIP-C01信息資訊
- AIP-C01考題  AIP-C01認證考試解析  AIP-C01考題  免費下載  $\triangleright$  AIP-C01  $\triangleleft$  只需在 [[www.newdumpspdf.com](http://www.newdumpspdf.com)] 上搜索AIP-C01通過考試
- 由最新的AIP-C01考試指南編訂的題庫資料 - 提高AWS Certified Generative AI Developer - Professional考試的通過率  請在  $\triangleright$  [www.newdumpspdf.com](http://www.newdumpspdf.com)  網站上免費下載  $\Rightarrow$  AIP-C01  $\Leftarrow$  題庫AIP-C01在線考題
- 由最新的AIP-C01考試指南編訂的題庫資料 - 提高AWS Certified Generative AI Developer - Professional考試的通過率  打開  [www.newdumpspdf.com](http://www.newdumpspdf.com)  搜尋 ( AIP-C01 ) 以免費下載考試資料AIP-C01通過考試
- AIP-C01在線考題  AIP-C01考古題更新  AIP-C01考試題庫  ( [tw.fast2test.com](http://tw.fast2test.com) ) 網站搜索  $\triangleright$  AIP-C01  $\triangleleft$  並免費下載最新AIP-C01題庫資源
- 受信任的Amazon AIP-C01: AWS Certified Generative AI Developer - Professional題庫資訊 - 最新的Newdumpspdf AIP-C01考試資訊  $\leftrightarrow$  請在  $\triangleright$  [www.newdumpspdf.com](http://www.newdumpspdf.com)  $\triangleleft$  網站上免費下載  AIP-C01  題庫最新AIP-C01題庫資源
- AIP-C01考題寶典  AIP-C01考證  AIP-C01題庫分享  來自網站 [[www.vcesoft.com](http://www.vcesoft.com)] 打開並搜索  $\Rightarrow$  AIP-C01  $\Leftarrow$  免費下載AIP-C01在線考題
- AIP-C01考古題更新  AIP-C01信息資訊  AIP-C01更新   [www.newdumpspdf.com](http://www.newdumpspdf.com)  上搜索 { AIP-C01 } 輕鬆獲取免費下載AIP-C01測試
- AIP-C01考試題庫  AIP-C01通過考試  AIP-C01通過考試  立即打開 [[www.newdumpspdf.com](http://www.newdumpspdf.com)] 並搜索  $\Rightarrow$  AIP-C01  以獲取免費下載AIP-C01證照信息
- [bbs.t-firefly.com](http://bbs.t-firefly.com), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [netro.ch](http://netro.ch), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [mavenmarg.com](http://mavenmarg.com), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), Disposable vapes

P.S. Testpdf在Google Drive上分享了免費的2026 Amazon AIP-C01考試題庫: <https://drive.google.com/open?id=114V8R-dBKMKU9FeVAyjvaDtUZz6v0kR>